

Belief's own metaethics? A case against epistemic normativity

By

Charles Côté-Bouchard, King's College London

Submitted
for the degree of Doctor of Philosophy
in Philosophy

September 2016

Abstract

Epistemology is widely seen as a normative discipline like ethics. Just like moral facts, epistemic facts – i.e. facts about our beliefs' epistemic justification, rationality, reasonableness, correctness, warrant, and the like – are standardly viewed as *normative* facts. Yet, whereas many philosophers have rejected the existence of moral facts, few have raised similar doubts about the existence of epistemic facts.

In recent years however, several metaethicists and epistemologists have rejected this Janus-faced or *dual* stance towards the existence of moral and epistemic facts. As recent developments in metaethics and normativity theory have made clear, objections to the existence of moral facts really are *metanormative* objections that target the existence of normative facts more generally. But since epistemic facts are no less normative than moral facts, the argument goes, the existence of the former is equally threatened by metaethical objections.

In this thesis, I argue that this rejection of the dual stance fails because epistemic facts are not normative facts. Although they imply norms, they do not imply genuine normativity since the epistemic norms of belief that they imply lack necessary normative authority or force. Unlike moral norms and just like e.g. norms of etiquette and the law, there is not automatically a normative reason to conform to epistemic norms. Therefore, even if metaethical objections target all normative facts, it does not follow that they also target epistemic facts.

I offer a two-part *abductive* argument in favour of that conclusion. First, I argue that epistemic facts lack five commonly cited features of normative facts (but not of merely norm-implicating facts). Then, I argue that this is best explained by the thesis that epistemic facts are merely norm-implicating and not genuinely normative. I end by exploring the potential consequences of this conclusion for epistemology and metaethics.

Acknowledgements

First and foremost, I must thank my supervisors Clayton Littlejohn and Maria Alvarez. I am immensely grateful for their feedback, patience, encouragements, criticisms, guidance, and support. They taught me a tremendous amount and I could not have done this without them.

I also benefited from the help and support of many members of the KCL philosophical community. Thanks especially to Bill Brewer, Mike Coxhead, Benjamin Davies, Paul Doody, Julien Dutant, Owen Englefield, Giulia Felappi, Claire Field, Will Harrison, Dave Jenkins, David Owens, David Papineau, Joseph Raz, Peter Ridley, Will Sharp, Nick Shea, Sophie Stammers, Mark Textor, Sarah Tropper, and Caspar Wilson.

Thanks also to Félix Aubé-Beaudoin, Simon-Pierre Chevarie-Cossette, Terence Cuneo, Marc-Kevin Daoust, Alexandre Duval, Davide Fassio, Elizabeth Fricker, Maud Gauthier-Chung, Alex Gajevich-Sayegh, Tanya Goodchild, Anandi Hattiangadi, Daniel Laurier, Alexandru Marcoci, Matthew McGrath, David Rocheleau-Houle, Christine Tappolet, and Teemu Toppinen for helpful feedback and discussion.

I also received useful comments on elements of this thesis from anonymous reviewers from the *International Journal of Philosophical Studies* and *Philosophical Studies*, and from audiences at the 2013 HU Berlin–KCL Philosophy Workshop, the 2013 SOPHIA conference in Salzburg, the 2014 SOPHIA conference in Ovronnaz, the 5th Northwestern-Notre-Dame Epistemology Graduate Conference, the 2014 Cracow Workshop in Analytic Philosophy, the 4th and 5th Edinburgh Graduate Epistemology Conference, the 2014 European Epistemology Network Meeting, the 2015 Spring Graduate Conference in Philosophy in London, the fourth ‘journées de métaéthique’ in Rennes, the 2015 SOPHIA Congress, the 12th London-Berkeley Philosophy Conference, the 4th Stockholm Philosophy Graduate Conference, the 2015 KU Leuven Epistemology Conference, the ‘Normativity and Rationality’ Seminar at KCL, and the fifth Understanding Value conference at the University of Sheffield.

The completion of this thesis was financially supported by the Social Sciences and Humanities Research Council of Canada (award number: 752-2012-1174) and the King’s College London graduate school (King’s Overseas Research Students award).

Finally, special thanks to my family for their love and support, and to Anne-Marie for making my PhD years some of the happiest of my life.

Table of Contents

Chapter 1: Epistemology, normativity, and metaethics	6
1.1 Epistemology as a normative discipline.....	6
1.2 Epistemology and metaethics.....	10
1.3 In defence of the Dual stance: Epistemic Non-Normativity	17
1.4 The thesis	24
PART I: EPISTEMOLOGY AND THE MISSING MARKS OF NORMATIVITY	27
Chapter 2: Value.....	28
2.1 Introduction: epistemology and value.....	28
2.2 The Triviality intuition	29
2.3 Explaining the Triviality intuition.....	32
2.4 Arguments for the ‘minimal value’ explanation	35
2.5 An objection to the ‘minimal value’ explanation	44
2.6 Summary	52
Chapter 3: Desire	53
3.1 Introduction: epistemology and desires	53
3.2 The No Desire intuition.....	56
3.3 Schroeder’s potential desire-satisfaction strategy.....	58
3.4 The aim of belief strategy.....	62
3.5 The idealization strategy	64
3.6 Kornblith’s rule-based strategy.....	65
3.7 Summary	69
Chapter 4: Autonomy	70
4.1 Introduction: epistemology and the autonomy of the normative	70
4.2 The argument from deliberation.....	72
4.3 The closed question argument	76
4.4 Summary	86
Chapter 5: Motivation	87
5.1 Introduction: epistemology and motivation.....	87
5.2 Against Epistemic Expressivism	88
5.3 Against the Epistemic Motivation Thesis (EMT).....	101
5.4 Summary	111
Chapter 6: Control	113
6.1 Introduction: epistemology and ‘ought’ implies ‘can’	113

6.2 Against the Epistemic Control Thesis (ECT)	115
6.3 The incompatibility of CP and Epistemic Normativity.....	124
6.4 Summary	127
PART II: EXPLAINING THE MISSING MARKS OF NORMATIVITY	128
Chapter 7: Two rival explanations	129
7.1 Epistemic Normativity’s explanation of the missing marks: Normative Pluralism..	129
7.2 Epistemic Non-Normativity’s explanation of the missing marks.....	131
7.3 Interlude: explaining apparent remaining marks of normativity.....	138
7.4 Summary	147
Chapter 8: Epistemic non-normativity as the best explanation.....	149
8.1 Simplicity	149
8.2 Additional explanatory criteria.....	157
8.3 Summary	161
Conclusion	162
Bibliography.....	169

Chapter I: Epistemology, normativity, and metaethics

1.1 Epistemology as a normative discipline

Imagine you are held hostage by a mad philosopher who keeps telling you that he does not exist. He threatens to kill you unless you meet a simple demand: you must believe what he says by the end of the day. That is, you must form the belief that he does not exist. If you do, he will let you go and never bother you again. He will even drug you so that you will have no memory of that ordeal. Merely pretending that you believe him won't do since he has an infallible lie detector test, which he will make you pass at the end of the day. Should you believe him?

From a *practical* or *pragmatic* standpoint, it clearly seems that you should. Not only would believing him literally save your life, it would do so without any adverse consequences. After all, you won't ever see him again and you will have no memory of the events. Yet there seems to remain another sense or standpoint according to which you should *not* believe him. For one thing, you know he exists. You can just see him in front of you and you have no evidence that you are dreaming or hallucinating. For another, he provides you with absolutely no evidence for his claim. So even if you *pragmatically* ought to believe that he exists, it also seems that you should not believe him from what philosophers call an *epistemic, theoretical, or truth-related* point of view.¹ That is, even if such a belief is permitted or required by practical norms or standards such as *prudential* or even *moral* norms, it is still forbidden by *epistemic* norms. The belief that your captor does not exist would still be, as epistemologists put it, *epistemically unjustified* and *irrational* even if it might be *pragmatically* justified and rational.²

This kind of example illustrates and motivates one of the most widely held ideas in contemporary epistemology, namely that, as Jaegwon Kim puts it,

¹ I use 'ought' and 'should' interchangeably in what follows.

² Although I associate notions like epistemic justification with epistemic norms and 'oughts', I do not assume what Alston (1988) calls a 'deontological' conception of epistemic justification i.e. a conception of epistemic justification "as having to do with obligation, permission, requirement, blame, and the like" (Alston 1988, 257). First, for all I say here, the 'ought' in 'S ought to ϕ ' could be *evaluative* and not *deontic* or *prescriptive*. That is, the epistemic 'ought' could be best seen not as an 'ought to do', but as an 'ought to be', i.e. as claims about what would be epistemically *valuable* or *good*. I return to this point below. Second, even if epistemic norms and 'oughts' are deontic – i.e. specifying what agents are epistemically required, permitted, or forbidden to believe – it is controversial to assume that they are thereby also *hypological* norms, i.e. having to do with responsibility, praiseworthiness, and blameworthiness. Many think that norms can be deontic without being hypological. See e.g. Zimmerman (2002), Srinivasan (2015), Littlejohn (Forthcoming a), and Williamson (Forthcoming).

“[e]pistemology is a normative discipline as much as, and in the same sense as, normative ethics” (Kim 1988, 383). In elucidating notions like epistemic justification, rationality, reasonableness, warrant, correctness, and knowledge, the thought goes, philosophers elucidate not merely what we *do* believe, but what we *ought* to believe. Just like moral *facts* and *claims*, *epistemic* facts and claims – e.g. the fact or claim that S’s belief is epistemically justified (unjustified), rational (irrational), warranted (unwarranted), reasonable (unreasonable), and correct (incorrect) – are *normative* facts and claims. I will take this idea that epistemology is normative to be captured by the following thesis:

Epistemic Normativity: *Epistemic facts and claims are normative facts and claims.*

Just like, say, ethics uncovers norms and facts that specify what we should or should not do, epistemology uncovers norms or facts that specify what we should or should not believe. Epistemology is, as Jonathan Adler (2002) put it, ‘belief’s own ethics’.

As the example above shows however, the idea of epistemology as belief’s ethics requires a qualification: epistemic claims like attributions of epistemic justification have to do with what we should or should not believe, but only from a distinctly epistemic, theoretical, or truth-related point of view. They have to do, in other words, with what we should or should not believe not according to moral norms, prudential norms, aesthetic norms, legal norms, and so on, but according to *epistemic norms*.

What do I mean by epistemic norms? Norms are generally seen as claims that specify conditions *C* under which one should or should not ϕ .³ Following von Wright (1963), I will refer to *C* as the norm’s *condition of application* and to ϕ as its *norm act*. As I have already indicated, what distinguishes epistemic norms from other kinds of norms – e.g. moral norms, prudential norms, aesthetic norms, etc. – is both their distinctive kind of norm act and of conditions of application. First, epistemic norms have *belief* as their norm act. They specify what we should or should not believe.⁴

³ Where, as I mentioned above, this ‘should’ or ‘ought’ can be either deontic or evaluative.

⁴ Let me make two clarifications. First, I actually take the norm act of epistemic norms to be *doxastic attitudes*, i.e. not only belief, but also disbelief and withholding or suspension of belief (or suspension of judgment as it is often put). I only mention belief here for simplicity. Second, using ‘epistemic norms’ to denote only *doxastic* epistemic norms is not strictly correct or at least not in line with the way the

However, as I mentioned above, this is not sufficient to demarcate the norms I am interested in since other kinds of norms might have belief as their norm act. Perhaps beliefs can also be e.g. immoral, imprudent, unfashionable, illegal, indecorous, unorthodox, and so on. Accordingly, what distinguishes epistemic norms is not only that they have belief as their kind of norm act, but also that they specify distinctly epistemic conditions of application. Of course, it is controversial what these exact conditions are since to propose such a condition is in part to give a substantial or first order account of what we should or should not believe. However, it is widely agreed that what marks conditions as epistemic is that they are truth-related. That is, they are conditions that have to do with notions such as truth, knowledge, and evidence. Accordingly, popular substantive accounts of epistemic norms include truth norms⁵ such as ‘you should believe that P if and only if P is true’ or ‘you should believe that P only if P is true’, knowledge norms⁶ like ‘you should believe that P only if you know that P’, and evidential norms⁷ such as ‘you should believe that P only if you have adequate evidence that P’. So in sum, epistemic norms are claims that specify distinctly epistemic or truth-related conditions under which one should or should not believe propositions.

Of course, this characterization of epistemic norms leaves several further substantive questions unanswered. Let me mention four. First, what is the *structure* of the domain of epistemic norms? How are the various epistemic norms related? Is there one fundamental or primary epistemic norm from which all other epistemic norms derive,⁸ or is there a plurality of fundamental norms⁹? If the former is true, then what is the primary norm and how does it entail the other norms? If the latter is true, then what are these fundamental epistemic norms? Second and relatedly, what is the *content* of the (fundamental and non-fundamental) epistemic norms? What

phrase is standardly used in the literature. In particular, contemporary discussions of epistemic norms include not only doxastic epistemic norms, but also epistemic norms of *assertion* and of *practical reasoning* (or of treating something as a reason). I therefore take ‘epistemic norms’ to be short for ‘doxastic epistemic norms’ or ‘epistemic norms of belief’ in this thesis. See e.g. Littlejohn and Turri (2014) for recent work on epistemic norms across all three dimensions.

⁵ David (2001), (2005a), Wedgwood (2002), Millar (2004), Lynch (2004), (2009), Shah and Velleman (2005), Shah (2006), Boghossian (2008), Whiting (2010), (2012) (2013a) (2013b) and Littlejohn (2012a).

⁶ E.g. Williamson (2000), Sutton (2007), Bird (2007), Hattiangadi (2010), McHugh (2011), Smithies (2012), and Littlejohn (2013), (Forthcoming b).

⁷ E.g. Clifford (1999), Feldman (2000), Adler (2002), and Conee and Feldman (2004).

⁸ Arguably the dominant approach. See e.g. Bonjour (1985), David (2001), (2005a), Wedgwood (2002), Boghossian (2008), Lynch (2009), Littlejohn (2013), (Forthcoming b), Steglich-Petersen (2013), and Williamson (Forthcoming). I return to this approach in chapter 4.

⁹ E.g. Alston (1993), (2005).

should we epistemically believe? As I just mentioned, popular answers include truth norms, knowledge norms, and evidential norms.

Third, are epistemic norms deontic or evaluative?¹⁰ Deontic or prescriptive ‘oughts’ are instances of ‘ought to do’, i.e. ‘oughts’ that belong to, apply to, or place demands on agents. So on a deontic reading, epistemic norms make demands or prescriptions to agents about what to believe or not to believe. Evaluative norms, on the other hand, are instances of ‘ought to be’. Instead of implying demands on agents, the evaluative sense of ‘S ought to ϕ ’ implies that it would be *good* if S ϕ -ed and thus that it ought to be the case that S ϕ -s. Hence on an evaluative reading, epistemic norms are not prescriptions, but rather claims about what it would be epistemically good to believe.¹¹

Finally, are epistemic norms objective, subjective, or both? In what way are they objective or subjective? If they can be both, how do subjective and objective epistemic norms relate? Is one kind the primary one that grounds the other kind? If so, which kind is primary?¹²

While these are crucial questions for understanding epistemic norms, it won’t be necessary for my purposes to answer them. To anticipate, my main aim in this thesis will be to examine the idea that epistemology is normative and, in particular, the thesis I labelled Epistemic Normativity. I hope to provide an evaluation of that thesis that does not depend on any particular substantive account of epistemic norms and which can be persuasive to proponents of all first-order epistemological theories.

Why focus solely on *doxastic* epistemic facts, claims, and norms? As I just mentioned, my target will be the idea that *epistemology* is normative. That is, I am interested in the normativity of the facts and claims that contemporary epistemologists are concerned with. But since contemporary epistemology is mainly the study of standards of belief that relate to knowledge, the epistemic claims that epistemologists are mostly concerned with are those having to do with the epistemic evaluation of beliefs. These include claims about whether S’s belief that P is epistemically justified, rational, reasonable, correct, warranted, and the like.

¹⁰ See e.g. Schroeder (2011) and Chrisman (2015) for discussion of this distinction.

¹¹ Most of the authors mentioned in notes 5 – 7 adopt the deontic approach. For the evaluative approach, see e.g. Alston (1993), (2005), Chrisman (2008), Fassio (2011), and McHugh (2012a). I return to this issue in chapters 2 and 6.

¹² For relevant discussion see e.g. Feldman (1988), Zimmerman (1996), (2008), Wedgwood (2002), Gibbard (2005), Gibbons (2013), Sepielli (Forthcoming), and Littlejohn (Forthcoming c).

1.2 Epistemology and metaethics

1.2.1 Moral facts and the metaethical objections

Although epistemology is widely seen as no less normative than ethics, philosophers typically have markedly different attitudes regarding the *existence* or *reality* of moral facts on the one hand and of epistemic facts on the other. Whereas few authors deny that there can be facts about whether people's beliefs are epistemically justified, rational, and the like, there is a long and rich tradition of rejecting the existence of facts about whether people's acts are morally right, wrong, good, bad, and the like. As is well-known from contemporary metaethics, philosophers have come up with numerous powerful objections to moral realism, i.e. the view that there are moral facts. The following is a very brief summary of some of the most prominent of these metaethical objections to the existence of moral facts (metaethical objections for short).¹³

Queerness. If moral facts existed, they would be both (i) normative or prescriptive, and (ii) true independently of our desires, aims, and institutions (just like, say, facts about tables, chairs, computers, black holes, etc.). That is, they would be objective or mind-independent facts that demand things from us and provide us with normative reasons. The problem is that none of the things that we know exist objectively or mind-independently are normative in that way. As Richard Garner writes:

It is hard to believe in objective prescriptivity because it is hard to make sense of a demand without a demander, and hard to find a place for demands or demanders apart from human interests and conventions. We know what it is for our friends, our job, and our projects to make demands on us, but we do not know what it is for reality to do so. A black hole swallows everything, but it demands nothing. (Garner 1990, 143)

The first problem with moral realism is therefore that if moral facts existed, they would be facts of a very strange and mysterious sort because their objective

¹³ To be clear, I am not endorsing these metaethical objections in what follows. The problems I summarize below might all turn out to have satisfactory answers and moral realism might be true. My point is only that these are some of the most pressing problems or challenges that moral realists have to deal with.

prescriptivity or normativity would make them fundamentally unlike any other facts we know exist.¹⁴

Supervenience. If there are moral facts, then they must be related to natural facts. But how? Realists answer that moral facts *supervene* on natural facts. That is, moral facts depend on natural facts in the sense that two things cannot be morally different without also being different from a non-moral natural point of view. However, it is not clear that this idea really illuminates the relation between moral facts and natural facts. As Stephen Schiffer writes: “invoking a special primitive metaphysical relation of supervenience to explain how non-natural moral properties were related to physical properties was just to add mystery to mystery, to cover one obscurantist move with another.” (Schiffer 1987, 153) Granted, if there are moral facts, then they must supervene on the natural. However, the fact that they do still calls for explanation. Why do moral facts have this particular relation with natural facts? The problem is that it is not clear how realists can answer. Hence, a second worry is that the relation between the moral and the natural seems to remain mysterious if moral realism is true.¹⁵

Disagreement. If there are facts about what is morally right and wrong, then why is there such widespread and seemingly ineliminable disagreement about moral questions? It seems that if there were moral facts, there would have been far more convergence on moral issues by now. The best explanation for the existence of widespread and seemingly ineliminable moral disagreement might very well be that there are simply no moral facts to agree on.¹⁶

Explanatory Impotence. It seems that if a kind of fact exists, then it must have explanatory power and be needed to explain other kinds of facts. The problem is that moral facts do not seem to be needed in order to explain the existence of any non-moral facts. As Gilbert Harman explains:

¹⁴ Mackie (1977), Garner (1990), Joyce (2001), and Olson (2014). See also e.g. Lillehammer (2004), (2013), Finlay (2008), and Joyce and Kirchin (2010) for discussion.

¹⁵ Blackburn (1971), (1984) and Horgan and Timmons (1992). See also e.g. Dreier (1992) for discussion.

¹⁶ Stevenson (1963), Harman (1975), Mackie (1977), Wong (1984), Loeb (1998), and Williams (1985).

You can observe someone do something, but can you ever perceive the rightness or wrongness of what he does? If you round a corner and see a group of young hoodlums pour gasoline on a cat and ignite it, you do not need to *conclude* that what they are doing is wrong; you do not need to figure anything out; you can see that it is wrong. But is your reaction due to the actual wrongness of what you see or is it simply a reflection of your moral "sense," a "sense" that you have acquired perhaps as a result of your moral upbringing? [...] [I]n neither case is there any obvious reason to assume anything about "moral facts," such as that it really is wrong to set the cat on fire or to cut up the patient in Room 306. Indeed, an assumption about moral facts would seem to be totally irrelevant to the explanation of your making the judgment you make. It would seem that all we need assume is that you have certain more or less well articulated moral principles that are reflected in the judgments you make, based on your moral sensibility. It seems to be completely irrelevant to our explanation whether your intuitive immediate judgment is true or false. (Harman 1977, 4-7)

Therefore, a fourth problem with moral realism is that moral facts do not seem to have explanatory power and to be needed to explain other kinds of facts. They rather seem explanatorily impotent and dispensable.¹⁷

Genealogy Evolutionary, sociological, and historical forces have had a tremendous influence on the content of our moral beliefs. However, such pressures have not plausibly tracked the moral truth. They have not plausibly selected, in other words, for our beliefs to match an independent moral reality. Rather, these pressures have selected for beliefs that promote reproductive fitness (in the case of natural selection) or group-selective fitness (in the case of cultural selection). Therefore, moral realists must say that these non-truth tracking influences either lead us to the truth by pure coincidence (which is extremely improbable) or pushed us *away* from the moral truth (in which case our moral beliefs are pretty much all false). Realists thus seem forced to the implausible sceptical conclusion that we cannot possibly know or justifiably believe the moral truth.¹⁸

¹⁷ Harman (1977), Williams (1985), Gibbard (1990), Wright (1992), (1995), and Leiter (2001).

¹⁸ Nietzsche (1887/2008), Ruse (1986), Kitcher (2005), (2011), Leiter (2006), Street (2006), Joyce (2006), Greene (2008), and Bedke (2009), (2014). See also Kahane (2011), Cuneo and Shafer-Laudau (2014), and Lillehammer (forthcoming) for discussion.

These objections are especially pressing if, as many realists maintain, moral facts are *non-natural* facts rather than facts that are reducible or analysable in terms of *natural* facts¹⁹ i.e. in terms of facts discovered by empirical sciences such as the facts that smoking causes cancer, that dinosaurs went extinct 65.5 million years ago, or that climate change is caused by humans. Some realists have argued that we can avoid at least some of the metaethical objections by adopting forms of *moral naturalism* according to which moral facts are reducible or analysable in terms of natural facts.²⁰ However, even if they can avoid the above objections, these naturalistic views are vulnerable to objections of their own. One prominent worry is the following:

Losing normativity. If moral facts are natural facts, then it becomes unclear how they can still keep their distinctly normative, reason-implicating, or authoritative character. David Enoch formulates this worry as follows:

Normative facts sure seem different from natural ones, different enough to justify an initial suspicion regarding reductionist attempts. This is especially clear when considered – as it should be – from the point of view of the deliberating agent. When I ask myself what I should do, it seems that just answering “Oh, pressing the blue button will maximize happiness” is a complete non-starter, it completely fails to address the question. Of course, given some background commitments it can be a better answer. If, for instance, I am already a convinced utilitarian, willing to commit myself to something like “It always makes sense to perform the action that maximizes happiness”, then “pressing the blue button will maximize happiness” seems like a reasonable answer to the question what should I do. But such background commitments are themselves paradigmatically normative, and themselves just too different from naturalist facts and beliefs. Absent such background commitments, “pressing the blue button will maximize happiness” seems just irrelevant to the question I ask myself, as does – in just the same way – any other purely naturalist answer. Rather than answering my question, such an answer simply changes the subject. It takes a normative commitment to render a naturalist answer to a normative question relevant. No natural fact by itself can have normative force. (Enoch 2011a, 107-108)

¹⁹ Moore (1903), Nagel (1986), Dworkin (1996), Hampton (1998), Scanlon (1998), (2014), Fine (2002), Shafer-Landau (2003), (2006), Strandberg (2004), Oddie (2005), Wedgwood (2007), Cuneo (2007), Fitzpatrick (2008), (2011), Parfit (2011), Enoch (2011a), and Cuneo and Shafer-Landau (2014).

²⁰ Railton (1986), Boyd (1988), Sturgeon (1988), Brink (1989), Jackson (1998), Jackson and Pettit (1990), Smith (1994), Foot (2001), Bloomfield (2001), Schroeder (2007), and Thomson (2008).

Moral facts are normative. They necessarily have normative authority or force and they imply normative reasons. However, purely natural facts do not seem to be the kind of thing that can have normative authority. Therefore, one central worry with moral naturalism is that if we reduce or analyse moral facts in terms of purely natural facts, then we lose their normativity, in which case we are changing the subject and not really dealing with genuine moral facts anymore.²¹

So the problem with moral realism can be roughly summarized as follows. If there are moral facts, then they are either natural or non-natural. If they are non-natural, then they are problematic in light of the worries regarding queerness, supervenience, disagreement, explanatory impotence, and genealogy. If they are natural, then perhaps they can escape *some* of these problems, but they seemingly do so at the cost of losing the essential normativity of moral facts.

1.2.2 Moral facts and epistemic facts: duality or unity?

As I mentioned above, even though philosophers typically take epistemic facts to be no less normative than moral ones, few have denied that there can be epistemic facts. That is, despite their wide acceptance of Epistemic Normativity, epistemologists and metaethicists have not traditionally thought that epistemic facts are vulnerable to metaethical objections like the ones above. Of course, philosophers have always been concerned with the threat of epistemological *scepticism*. However, the point of sceptical arguments in epistemology is distinct from that of the metaethical objections. The sceptic's claim is not that there cannot be facts about whether S's belief is epistemically justified or an instance of knowledge, but rather that knowledge and epistemic justification are unattainable. So although countless epistemologists have discussed the possibility that we might not know or justifiably believe anything, few have considered seriously the possibility that there are no epistemic facts, i.e. that epistemic realism is false.²²

Therefore, philosophers have traditionally occupied what I will call a *Dual stance* with respect to the existence of moral and epistemic facts:

²¹ See also Moore (1903), Johnston (1989), Scanlon (1998), (2014), Dancy (2006), Fitzpatrick (2008), (2010), and Parfit (2011). This line of thought is also pushed by many moral anti-realists. See for instance, Mackie (1977), Garner (1990), Gibbard (1990), Joyce (2001), Ridge (2014), and Olson (2014).

²² Exceptions include e.g. Field (1998), (2009), Street (2009), Hazlett (2013), and the proponents of epistemic expressivism, which I mention in chapter 5.

Dual stance: *Whether or not the metaethical objections are sound, they do not threaten the existence of epistemic facts. Even if the metaethical objections successfully established the non-existence of moral facts, it would not follow that there are no epistemic facts.*

In recent years however, a number of philosophers have argued such a Janus-faced stance is untenable. As many metaethicists and normativity theorists recently pointed out, metaethical objections such as those I outlined above are best seen as *metanormative* objections. That is, they are best seen as arguments against the existence of *normative* facts more generally. Their thought is that the allegedly objectionable features of moral facts – e.g. their objective prescriptivity, their mysterious relation with natural facts, the widespread disagreement about them, their explanatory impotence, and their problematic genealogy – are really features of normative facts more generally. Scanlon acknowledges this recent metanormative trend in the opening lines of his 2014 book:

Contemporary metaethics differs in two important ways from the metaethics of the 1950s and 1960s, and even the later 1970s, when John Mackie wrote *Ethics: Inventing Right and Wrong*. In that earlier period, discussion in metaethics focused almost entirely on morality: on the proper interpretation of claims about moral right and wrong, and other forms of moral evaluation. Today, although morality is still much discussed, a significant part of the debate concerns practical reasoning and normativity more generally: reasons for action, and, even more broadly, reasons for belief and other attitudes, which are increasingly recognized as normative, and as raising questions of the same nature as those about reasons for action. (Scanlon 2014, 1)

Similarly, Enoch writes:

Ethical or moral facts [...] are a subset of normative facts. And it is an increasingly appreciated lesson of the recent decades' metaethical literature that many of the concerns and arguments traditionally thought to be about morality are really easily and naturally generalizable to metanormative concerns and arguments. (Enoch 2011a, 2)

This metanormative trend also characterizes the work of contemporary metaethicists who *deny* the existence of moral facts and who defend the objections above. Sharon Street, for example, writes:

The right philosophical account of normativity, in my view, will be fully unified across the practical and epistemic domains. I therefore agree with those who think that antirealism about practical reasons hasn't received a full defence until antirealism about epistemic reasons has been defended as well. (Street 2009, 213-214)

The point is that the target of the metaethical objections is not simply morality, but normativity more generally. This is what I will call the Metanormative thesis:

Metanormative thesis: *metaethical objections really are metanormative objections; if any of the metaethical objections is sound, then there are no normative facts.*²³

If the Metanormative thesis is right, then there is a clear tension between the traditional Dual stance on the one hand and Epistemic Normativity on the other. If epistemic facts are normative facts and if the metaethical objections really are metanormative objections, then epistemic facts are just as vulnerable to these objections as moral facts.

Consider the queerness objection I introduced above for example. If epistemic facts are normative like moral facts, then they are no less mysterious or strange. Just like moral facts, they would have normative authority independently of our desires, aims, and institutions. In the mad philosopher example, you plausibly have no desire to avoid error and follow the evidence about the question at hand. You care about surviving and so you presumably want to believe your captor. But regardless of your desires, believing your captor remains epistemically unjustified, irrational, and incorrect. It would still be something you epistemically ought not to believe.

²³ In addition to Street (2009), Enoch (2011a), and Scanlon (2014) see e.g. Gibbard (1990), (2003), Korsgaard (1996), Scanlon (1998), Street (2006), Wedgwood (2007), Cuneo (2007), Schroeder (2007), Finlay (2010), Bedke (2010), Skorupski (2010), Parfit (2011), Ridge (2014), and Olson (2014).

Therefore, given the Metanormative thesis and Epistemic Normativity, it seems that we should give up the Dual stance in favour of what I will call a *Unitary stance*:

Unitary stance: *If any of the metaethical objections is sound, then there are no epistemic facts.*²⁴

This is what I will call the *argument from epistemic normativity* in favour of the Unitary stance and against the Dual stance:

The argument from Epistemic Normativity

1. If any of the metaethical objections is sound, then there are no normative facts; metaethical objections really are metanormative objections. (*Metanormative thesis*)
2. Epistemic facts and claims are normative facts and claims. (*Epistemic Normativity*)
3. Therefore, if any of the metaethical objections is sound, then there are no epistemic facts. (*Unitary stance*)

1.3 In defence of the Dual stance: Epistemic Non-Normativity

My goal in this thesis is to reject the argument from Epistemic Normativity. The most common strategy for doing so has been to reject the Metanormative thesis and claim that the metaethical objections only target a certain kind of normative facts. In particular, many think that such objections do not affect the kind of normativity that is grounded in our *desires*. As a result, some have defended the Dual stance by arguing that unlike moral normativity, the normative authority of epistemic norms is grounded in our desires.

²⁴ Some have recently used this idea to mount a ‘companions in guilt’ defence of moral realism against the metaethical objections. Cuneo (2007) provides the most well-developed and forceful defence of this strategy. See also Stratton-Lake (2002) and Rowland (2013). See Lillehammer (2007) and Cowie (2014), (2015) for discussion of this kind of strategy. One thing to note is that Cuneo construes epistemic facts much more widely than I do. As I explained in section 1.1, my focus in this thesis is on facts about what we should or should not *believe* epistemically speaking, i.e. the kind of facts that are at the centre of contemporary theory of knowledge and epistemic justification. Cuneo, on the other hand, takes epistemic facts to also include facts about we should *do* and *feel*.

This won't be my strategy however. For one thing, as I show in chapter 3, the normative authority of epistemic norms cannot be grounded in our desires. This is because conforming to epistemic norms does not necessarily promote the satisfaction of one's desires. For another, although it is not my aim to defend it here, I am sympathetic to the Metanormative thesis.

Instead, my strategy will be to reject the second premise – i.e. the thesis I called Epistemic Normativity – and defend the following:

Epistemic Non-Normativity: *Epistemic facts and claims are not normative facts and claims.*

The argument from Epistemic Normativity fails not because some normative facts are not vulnerable to the metaethical objections, but rather because *epistemic facts and claims are not normative facts and claims*. The following will be, for the most part, a case in favour of Epistemic Non-Normativity and against Epistemic Normativity. However, before presenting what that case will be, let me first clarify what I mean by Epistemic Normativity and Epistemic Non-Normativity.

A first thing to note is that, as is almost universally recognized in metaethics and normativity theory, we must distinguish genuine *normativity* from what we might call mere *norm-relativity*.²⁵ That is, we must distinguish between facts and claims being genuinely *normative* and them being, as I will put it, merely *norm-implying*.²⁶ Norms are one thing, but normativity is another. As John Broome explains:

[I]n one sense 'normative' simply means to do with norms, rules, or correctness. Any source of requirements is normative in this sense. For example, Catholicism is. Catholicism requires you to abstain from meat on Fridays. It is a rule and it is incorrect according to Catholicism to eat meat on Fridays. So Catholicism is normative in this sense. But I do not use 'normative' in that sense. In my sense, it means to do with ought or reasons. Given a rule or a requirement we can ask whether you ought to follow it, or whether you have reason to do so (Broome 2007, 162).

²⁵ Tiffany (2007) and Finlay (2008) can be interpreted as breaking with this consensus.

²⁶ The term 'norm-relativity' is from Hattiangadi (2007). Other labels used in the literature for the same distinction include *reason-implying* versus mere *rule-implying* normativity (Parfit, 2011), *robust* versus merely *formal* normativity (McPherson, 2011; see also Maguire and Woods, MS.), *strong* versus *weak categoricity* (Joyce, 2001), *normative requirements* versus mere *requirements* (Broome, 2013), and *irreducible* versus merely *reducible normativity* (Olson, 2014).

While all norms trivially set standards relative to which certain things can be required, permitted, forbidden, correct, incorrect, good, bad, and the like, not all norms necessarily or automatically have genuine normative authority. If a norm N forbids ϕ -ing under conditions C, then trivially, ϕ -ing under C is forbidden, incorrect, wrong, or bad *relative to* or *according to* the standard set by N. However, the question of what is required or permitted according to N is a purely descriptive and non-normative question. It is distinct from the further *normative* question of what there is *normative reasons* to do. More generally, for any norm N, you can very well recognize that N says you should or should not ϕ and still ask the distinctly normative question whether there is any normative reason to do as N says.²⁷

Plausibly, most norms are such that the answer to this normative question could be ‘no’. For most norms, in other words, there could be situations in which, even though these norms ask you to ϕ , there is no normative reason for you to ϕ . Philippa Foot famously illustrated this using the example of etiquette:

[...] one may reasonably ask why anyone should bother about what should-e, (should from the point of view of etiquette) be done, and that such considerations deserve no notice unless reason is shown. So although people give as their reason for doing something the fact that it is required by etiquette, we do not take this consideration as in itself giving us reason to act. Considerations of etiquette do not have any automatic reason-giving force, and a man might be right if he denied that he had reason to do “what's done.” (Foot 1972, 309)

Richard Joyce makes the same point using the example of gladiatorial combat:

Consider Celadus the Thracian, an unwilling gladiator: he’s dragged off the street, buckled into armor, and thrust into the arena. [...] Let’s imagine that there are various rules of gladiatorial combat: you ought not throw sand in your opponent’s eyes, for instance. [...] Imagine that things are looking bleak – his opponent is a sadistic professional fighter, and Celadus finds himself pinned down and swordless. His only hope is to throw some sand in his rival’s eyes. (Let’s stipulate, with utter implausibility, that he can get away with nobody seeing him do this, just as a way of being sure that there will be no negative repercussions in the form of punishment for breaking the rules.) The

²⁷ I borrow the label ‘normative question’ from Korsgaard (1996) and Broome (2007), (2013).

rules still say that Celadus shouldn't do it, but he doesn't care about the rules – he has no particular reason to follow them, and every reason to reject them. Given that he has never entered into any form of contract to follow the rules, and that following the rules will lead to his quick and unjust demise, I think we will all agree that Celadus ought to throw sand in his opponent's eyes. (Joyce 2001, 34-35)

Another key example is the law or legal norms. As Michael Zimmerman explains:

Suppose that, when I say that the right thing for you to do is X, what I mean, more particularly, is that doing X is what the law requires of you. And suppose that this is in fact the case. Do you then have a *good reason* to do X? Perhaps you do. If the law in question is morally well founded, then presumably you will have a *moral* reason to do X. Or if running afoul of the law puts you at risk of punishment, you will presumably have a *prudential* reason to do X. But does the law give you an *extra* reason to do X, one that *strengthens* the considerations in favor of doing X already provided by morality and prudence? Arguably not. If so, then, although moral requirements, prudential requirements, and perhaps other kinds of requirements have normative force, legal requirements as such do not. (Zimmerman 2015, 15)²⁸

Finally, Derek Parfit writes:

Certain acts are required, for example, by the law, or by the code of honour, or by etiquette, or by certain linguistic rules. It is illegal not to pay our taxes, dishonourable not to pay our gambling debts, and incorrect to eat peas with a spoon, to spell 'committee' with only one 't', and to use 'refute' to mean 'deny'. Such requirements or rules are sometimes called 'norms'. [...] When there are such rules or requirements, we may have reasons to follow them. But these reasons are mostly provided, not by the mere existence or acceptance of these rules, but by certain other facts, most of which depend on some people's acceptance of these rules. When there are no such reason-giving facts, we may have no reason to follow some rule or requirement. We may have no reason, for example, to follow some fashion, or to refrain from violating some taboo. When I was told, as a child, that I shouldn't act in certain

²⁸ This one is more controversial. Many philosophers of law maintain, for instance, that any plausible account of law must accommodate its normativity. As a result, many reject theories according to which legal facts are nothing over and above non-normative facts – e.g. legal positivism – on the grounds that they cannot account for their normative force. I have no space here to give a satisfactory response to this line of thought – see e.g. Enoch (2011b) for such a response – and, more generally, to defend Zimmerman's quote. Instead, I will simply follow Enoch, Zimmerman, Parfit, and others, and take for granted that these legal philosophers are either mistaken or not using 'normativity' in the sense that metaethicists and normativity theorists are interested in. As it happens – and as will become clear below – I think most epistemologists have made a similar mistake.

ways, and I asked why, it was infuriating to be told that such things are *not done*. That gave me no reason not to do these things. (Parfit 2011, 144-145)²⁹

We could multiply the examples. Plausibly, there is not automatically a normative reason to conform to the norms or standards set by gender stereotypes, tradition, religious sects, club rules, convention, and so on. It is not because these norms require or forbid you to ϕ that there is any normative reason for you to ϕ or not to ϕ . The point is that these are norms that lack *necessary normative authority or force*. They do not necessarily or automatically imply normative reasons, i.e. they are not, as it is often put, *reason-giving* or *reason-implying*.

One could demur at this point and reject the underlying distinction between normativity and mere norm-relativity. Perhaps being normative is nothing over and above being merely norm-implying. That is, perhaps there is nothing more to there being normative reasons to ϕ than ϕ -ing being required or correct relative to some standard.³⁰ If this were true, however, then all we would have to do in order to generate normative reasons to ϕ would be to come up with norms or standards – however arbitrary, silly, or horrible – requiring us to ϕ . But this is implausible. Normative reasons to ϕ are considerations that justify or favour ϕ -ing. But intuitively, the sole fact that ϕ -ing is required by a norm – whatever the norm – does not entail that ϕ -ing is justified or right to any extent. We can imagine a serial killing club, for instance, that would require murdering at least five innocent people without getting caught. But regardless of what the rules of this club say, there is no normative reason for anyone to pursue this horrible course of action. In any case, I will follow the general metanormative consensus in this thesis and take for granted the distinction between mere norm-relativity and genuine normativity.

Crucially, this distinction means that claims can be norm-implying without being normative. That is, a claim can imply that, relative to a particular norm or standard N, ϕ -ing is required, permitted, forbidden, correct, incorrect, good, bad, and so on *without* implying that there is any normative reason to conform to that standard, i.e. without entailing genuine normativity. This is so for claims that imply norms that

²⁹ See also Hattiangadi (2007) for a book-length defence of the claim that semantic facts and claims are merely norm-implying and not genuinely normative. As will become clear below, my project could be described as an epistemological version of Hattiangadi's project about meaning.

³⁰ As noted above, versions of this suggestion can be attributed to Tiffany (2007) and Finlay (2014).

lack necessary normative authority such as etiquette claims, legal claims, fashion claims, and the like. But this means, in turn, that claims like these are not genuinely normative claims. After all, they are not made true by normative facts. These are claims that would not cease to be true, in other words, if there were no normative facts.

To claim that S is wearing an unfashionable outfit, for example, is, trivially, to claim that she violates fashion norms or standards. However, since there is not automatically a normative reason to conform to fashion norms – i.e. since fashion norms lack necessary normative authority – to make such a fashion claim is *not* also to claim that there is a normative reason for S not to wear that outfit. Whether there is such a normative reason is a further question that is not settled by that fashion claim alone. Consequently, that fashion claim can be true even if there is no normative reason for S to be fashionable. All it takes for that claim to be true is for S's wearing that outfit to really be forbidden by fashion norms. Therefore, although they are norm-implying, fashion claims are not normative because whether something or someone is fashionable is independent from normative facts. Although they imply norm-relativity, in other words, they do not imply genuine normativity.

Similarly, ϕ -ing is legal or illegal, decorous or indecorous, traditional or unorthodox, conventional or unconventional, grammatical or ungrammatical whether or not there is any normative reason to ϕ . The point is that claims that imply norms which lack necessary normative authority are not normative claims since they do not also imply genuine normativity and so they are not made true or false by normative facts. Even if the metaethical or metanormative objections are sound and there are no normative facts, it can still be the case that certain acts are legal or illegal, decorous or indecorous, traditional or unorthodox, conventional or unconventional, grammatical or ungrammatical, and so on.

Moral claims, in contrast, seem to be paradigmatic cases of normative claims. Unlike fashion claims, etiquette claims, legal claims, and the like, moral claims imply norms (moral norms) which do have necessary normative authority. Plausibly, if morality requires you to ϕ , then there is automatically a normative reason for you to ϕ . If this is right, then the claim that an act is immoral *does* entail that there is a normative reason not to do it. So unlike the merely norm-implying claims above, moral claims are made true or false by normative facts. Whether ϕ -ing is moral or

immoral is *not* independent from the normative reality. If the metanormative objections are sound and there are no normative facts, then there are no moral facts either since normativity essential or inherent to moral facts. Hence, unlike the claims above, moral claims are normative claims.

I am now in a position to give a more precise formulation of Epistemic Normativity and Epistemic Non-Normativity. Given what I have just explained, my goal will be to reject the following:

Epistemic Normativity (extended formulation): *Epistemic facts and claims are normative facts and claims since they are reason-implying. They imply norms – epistemic norms – which have necessary normative force or authority, i.e. norms to which there is necessarily a normative reason to conform.*

And instead defend the following:

Epistemic Non-Normativity (extended formulation): *Epistemic facts and claims are norm-implying, but not normative. They imply norms – epistemic norms – that lack necessary normative force or authority, i.e. norms to which there is not necessarily a normative reason to conform.³¹*

Unlike, say, moral norms and just like fashion norms, legal norms, etiquette norms, and many others, epistemic norms are not reason-implying. There is not necessarily a normative reason to conform to them. So although epistemic claims are about what we should or should not believe epistemically speaking, they are not normative. They are merely norm-implying. Hence, even if the metanormative objections are sound and there are no normative facts, there can still be facts about whether people's beliefs are epistemically justified, rational, reasonable, warranted, and so on.

As I mentioned above, this thesis will be, for the most part, a case in favour of Epistemic Non-Normativity and against Epistemic Normativity. The argument I offer is a two-part *abductive* argument or inference to the best explanation. First, I identify five features which are commonly cited as essential marks of normative facts

³¹ Although it is by far the minority view in epistemology, metaethics, and normativity theory, versions of it can be attributed to Maffie (1990), Laudan (1990a), (1990b), Papineau (2013), Hazlett (2013), and Maguire and Woods (MS). See also Grimm's (2009) reading of Sosa (2007).

and claims – as opposed to merely norm-implicating facts and claims, which lack them – and argue that epistemic facts and claims do not bear any of these five features. Then, I argue that this is best explained by Epistemic Non-Normativity. Hence, I conclude that we should reject Epistemic Normativity and adopt Epistemic Non-Normativity instead.

1.4 The thesis

As I have just explained, my thesis will essentially be a defence of Epistemic Non-Normativity and a rejection of Epistemic Normativity via the following argument:

The abductive argument from the marks of normativity

1. Epistemic facts and claims lack five commonly cited marks of normativity, i.e. five features that merely norm-implicating facts and claims lack, but which are commonly attributed to normative facts and claims. (PART I)
2. The best explanation of premise 1 is that Epistemic Non-Normativity is true, i.e. that epistemic facts and claims are merely norm-implicating and not normative. (PART II)
3. Therefore, in all likelihood, Epistemic Non-Normativity is true and Epistemic Normativity is false.

I argue for the first premise of this argument in Part I. This first part has five chapters each corresponding to the five commonly cited marks of normativity, which, I argue, are missing in epistemic facts and claims. These are five characteristics that merely norm-implicating facts and claims lack, but which are commonly attributed to normative facts and claims. These features are (i) a necessary connection with *value*, (ii) a necessary connection with *desire*, (iii) a necessary *autonomy* from non-normative facts, (iv) a necessary connection with *motivation*, and (v) a necessary connection with *control*. I argue in chapters 2-6 that just like merely norm-implicating facts and claims, epistemic facts and claims lack these five features.³² More precisely, I argue that it is not necessarily *good* to conform to epistemic norms (chapter 2), that conforming to

³² I defend the claim that merely norm-implicating facts and claims – e.g. etiquette claims, fashion claims, legal claims, etc. – lack these five features in the second section of chapter 7. Part I focuses mainly on why *epistemic* facts and claims lack these characteristics.

epistemic norms does not necessarily promote one's desires (chapter 3), that epistemic claims are not necessarily autonomous from non-normative claims (chapter 4), that epistemic judgments are not necessarily motivating (chapter 5), and that epistemic claims do not necessarily imply control or 'can' (chapter 6).

To be clear, the point of part I is not that these actually are essential marks of normativity in general or of all normative facts.³³ There would be no need for the second premise of my argument if this were my claim. If part of my argument was that facts cannot count as normative unless they have these features, then Epistemic Non-Normativity would follow directly from each chapter of part I. But this is not part of my argument. Instead, the first premise only implies that these are five features (i) which are often attributed to paradigmatic examples of normative facts like moral and prudential facts and claims, (ii) which merely norm-implicating facts and claims lack, and (iii) which epistemic facts and claims lack as well.

I then move to the second premise in part II. This second part has two chapters that together show that the conclusions of chapters 2-6 are best explained by Epistemic Non-Normativity. More precisely, I argue that the best explanation of the fact that epistemic facts and claims do not bear the commonly cited marks of normativity listed in part I is simply *that they are not normative facts and claims*. In chapter 7, I clarify both this explanation and the rival explanation that proponents of Epistemic Normativity must give. I argue that given what I argue in part I, proponents of Epistemic Normativity can only explain part I via what I call *Normative Pluralism*. According to Normative Pluralism, epistemic facts and claims do not bear the five commonly cited marks of normativity from part I because (i) there is a plurality of kinds of normativity and (ii) the kind of normativity implied by epistemic facts and claims – i.e. epistemic normativity – lacks these five features. In chapter 8, I evaluate these two explanations according to commonly invoked explanatory virtues or criteria. I argue that Epistemic Non-Normativity clearly comes out as the best explanation given these criteria.

³³ Indeed, I consider in part II the possibility that none of them are essential features of all normative facts and claims. This is the basis of the explanation that proponents of Epistemic Normativity can give of part I. I say more about it below, but this potential explanation is, roughly, that epistemic facts just are a kind of normative facts that lack all of these features. Epistemic facts, in other words, just have a kind of normativity that does not bear any of these marks. I clarify this possible explanation in chapter 7 and argue in chapter 8 that Epistemic Non-Normativity's explanation is better.

Since the results of part I are best explained by Epistemic Non-Normativity, I conclude that we should reject Epistemic Normativity and adopt Epistemic Non-Normativity instead. Epistemic facts and claims are best seen as merely norm-implicating and not genuinely normative. Therefore, I also conclude that the argument from Epistemic Normativity against the Dual stance and in favour of the Unitary stance is not sound. Since they are not genuinely normative, we cannot conclude from the Metanormative thesis that epistemic facts are also vulnerable to the metaethical objections.

PART I: EPISTEMOLOGY AND THE MISSING MARKS OF
NORMATIVITY

Chapter 2: Value

2.1 Introduction: epistemology and value

A first feature often associated with normativity, but not with mere norm-relativity is a necessary connection with *value* or *goodness*. It is the idea, in other words, that unlike merely norm-implying facts and claims, normative facts and claims necessarily imply or have to do with value. I will take this first commonly cited mark of normativity to be captured by the following principle:

Value principle (VP): *there is a normative reason to ϕ if and only if it is good to ϕ .*³⁴

VP is most readily associated with *value-based* or *teleological* accounts of normativity according to which normative reasons come from or depend on facts about value. When there is a normative reason to ϕ , in other words, this is ultimately because or in virtue of the fact that ϕ -ing is good or valuable.³⁵

It is important that this feature be understood in terms of a *necessary* connection with value and not, for instance, as the weaker claim that it must be very *often* or *almost always* good to ϕ . This is because such a weaker principle would be easily met by many merely norm-implying claims and by norms that clearly lack necessary normative authority. After all, it is very often or almost always good to conform to e.g. the law, etiquette, or fashion. However, it is not *necessarily* good to do so. Paradigmatic examples of normative claims, on the other hand, do seem to have such a necessary connection with value. It is hard to imagine, for example, cases where there would not be anything good in being moral or prudent.

If VP is true and if epistemic facts and claims are normative – as Epistemic Normativity maintains – then epistemic facts and claims must also have a necessary connection with value. More precisely, if normative reasons are necessarily connected

³⁴ Here and in what follows, I take ‘good’ and ‘valuable’ to mean *pro tanto* or *prima facie* good and not good *all things considered* or *overall*. I also take ‘good’ in VP to mean either *finally* (or *intrinsically*) or *instrumentally* good. I discuss additional distinctions such as epistemic versus non-epistemic value, and attributive versus predicative value (or good ‘simpliciter’) below.

³⁵ See e.g. Moore (1903), Stich (1990), Väyrynen (2006), Reisner (2009a), (2015), Steglich-Petersen (2011), and Maguire (2016). See e.g. Scanlon (1988, chapter 2), Way (2013), and Wiland (2012, chapter 4) for discussion.

with value and if there is necessarily a normative reason to conform to epistemic norms, then the following must also be true:

Epistemic Value Thesis (EVT): *it is necessarily good to conform to epistemic norms.*

The route from VP to EVT can be summarized as follows:

1. There is a normative reason to ϕ if and only if it is good to ϕ . (VP)
2. There is necessarily a normative reason to conform to epistemic norms. (Epistemic Normativity)
3. Therefore, it is necessarily good to conform to epistemic norms. (EVT)

In this chapter, I argue that the Epistemic Value Thesis (EVT) is false. It is not necessarily good to conform to epistemic norms. Therefore, the commonly cited mark of normativity captured by VP is missing in epistemic facts and claims.

The starting point of my argument against EVT, which I introduce in sections 2.2 and 2.3, is what I call the *Triviality intuition*. According to the Triviality intuition, there are cases where, intuitively, conforming to epistemic norms does not matter at all. There can be situations, in other words, where there seems to be *nothing* good in e.g. believing the truth, avoiding error, following the evidence, and the like. The problem with EVT, I will argue, is that it cannot give a satisfactory explanation of the Triviality intuition. It cannot, in other words, meet the explanatory challenge that the Triviality intuition poses to the idea that it is necessarily good to conform to epistemic norms. This is because (i) none of the most prominent arguments in support of EVT's explanation succeed and (ii) it is vulnerable to a fundamental objection. I argue for (i) and (ii) in sections 2.4 and 2.5 respectively.

2.2 The Triviality intuition

According to the Epistemic Value Thesis (EVT) it is necessarily good to conform to epistemic norms. However, an obvious initial worry with this suggestion is that we can easily imagine cases where there is intuitively nothing good in believing the truth, avoiding error, following the evidence, and the like.

First, we can imagine cases where, intuitively, believing in conformity with epistemic norms would be overwhelmingly bad and would not lead to anything good

(and where violating epistemic norms would be overwhelmingly good and would not lead to anything bad). Here is one such example:

Illness. After making several tests, Rita's doctor has bad news. She has a very serious illness and only has a 5% chance of survival. This particular illness is very sensitive to patients' anxiety and stress levels. Since Rita is very anxious and stressed about dying, following the evidence and believing the truth about her prognosis will lower her chances to almost 0. On the other hand, believing (falsely and unjustifiably) that she will almost certainly survive will dramatically increase her chances.

Second, there can be cases where, intuitively, (i) there is nothing bad and a lot of good in believing what is epistemically *unjustified* and (ii) nothing good in conforming to epistemic norms given the completely inconsequential character of the matter at hand. Here is one example:

Even Stars. Myriam is deeply fascinated by even numbers. She thinks they are the most harmonious and aesthetically pleasing thing in the universe. Realizing that something is in an even quantity always fills Myriam with deep awe and joy. One night, she looks at the sky and suddenly realizes something incredible: the number of stars could very well be even. This thought fills her with such excitement and wonder that she decides to start pretending and assuming that the number of stars really is even. Why not, she thinks? It is not as if being wrong about the number of stars could have any adverse consequences. Weeks pass by and her pretence unconsciously and gradually turns into a genuine, but epistemically unjustified belief.

Finally, there can be cases where there is seemingly nothing good in conforming to epistemic norms as well as nothing bad in violating them because of the complete insignificance or triviality of the matter at hand. For example:

Dream Ten years ago, Vincent had a dream in which he saw a historian on TV saying that a pub in London named The Red Lion closed its doors on February 1st 1748. Today, something randomly reminds him of the content of

that dream. However, it has been so long since he had that dream that he cannot remember if he really saw that on TV or if it was a dream. But since he does not care at all about such a trivial and insignificant issue he just assumes that it really happened without really thinking about it. He thus forms the false and epistemically unjustified belief that a pub named the Red Lion closed its doors on February 1st 1748 as a result. He then goes on to believe that for the rest of his life without it ever occurring to him again.

We could multiply the examples. Intuitively, there is (at least sometimes) nothing good in believing the truth, avoiding error, or following the evidence concerning e.g. the number of blades of grass on your neighbour's lawn at 11:59 am on 11 May 2016, how long it took the fourteenth customer served today at the third most popular coffee shop in New Jersey to drink her coffee, the average number of threads in the carpets of all the blue houses in Ireland, your grandfather's favourite sexual position when he was 39 years old, and so on. More abstractly, there does not seem to be necessarily something good in conforming to epistemic norms about, as Marian David puts it, "[...] every conjunction of every two truths I already believe (trivial or non-trivial), including [...] complex conjunctions whose conjuncts are themselves conjunctions of truths I already believe (trivial or non-trivial)" or about "every disjunction of (trivial or non-trivial) truths I already believe with any propositions you like." (David 2005b, 297-298)

The point is that agents can be in situations where the demands of epistemic norms seem, on the face of it, completely trivial and insignificant, i.e. situations where, intuitively, it does not matter at all that they violate epistemic norms. This is what I will call the *Triviality intuition*:

The Triviality intuition: *there can be cases where there is intuitively nothing good in conforming to epistemic norms.*³⁶

To see the significance of the Triviality intuition for Epistemic Normativity, note that it does not seem to extend to the two paradigmatic examples of norms with necessary normative authority that I have already mentioned, namely *moral* and

³⁶ See e.g. Whiting (2013c) for discussion.

prudential norms. While we can easily imagine cases where there seems to be nothing good in believing the truth, avoiding error, following the evidence, and the like, it is much harder to imagine cases where there is nothing good in being prudent or moral and nothing bad in being imprudent or immoral. Intuitively, to find out that ϕ -ing is moral or prudent just is to find out that ϕ -ing is good to some extent.

Of course, the Triviality intuition alone does not suffice to show that the Epistemic Value Thesis (EVT) is false. It does, however, pose a pressing explanatory challenge or question to EVT. The challenge is this: if, as EVT claims, it is necessarily good to conform to epistemic norms, then why are there cases where there seems to be nothing good in conforming to epistemic norms? Since taking the Triviality intuition at face value would mean rejecting the necessary goodness of conforming to epistemic norms and thus directly challenges the Epistemic Value Thesis, I take it that EVT is only viable if it can plausibly meet this explanatory challenge. In the rest of this chapter however, I will argue that it cannot plausibly meet this challenge.

2.3 Explaining the Triviality intuition

Why do we have the Triviality intuition? Why does it seem, intuitively, that there is sometimes nothing good in conforming to epistemic norms? Of course, there is a straightforward answer if we are willing to reject EVT: it is simply because sometimes, there *really* is nothing good in conforming to epistemic norms. This is how Epistemic Non-Normativity explains it. But how can proponents of Epistemic Normativity and EVT explain the Triviality intuition? If you think that conforming to epistemic norms is necessarily good, then why are there cases where, intuitively, it seems that there is nothing good in e.g. believing the truth, avoiding error, following the evidence, and the like?

Before examining what is generally seen as the most promising answer, let me first rule out two initially tempting, but ultimately inadequate answers. The first is what I will call the *overall value* explanation:

The ‘overall value’ explanation: *The intuition is due to a failure to distinguish between pro tanto or prima facie value and overall or all things considered value. We mistake the fact that it is not necessarily good all things considered to conform to epistemic norms with the stronger claim that it is not necessarily pro tanto good to do so.*

The problem is that the Triviality intuition is already about pro tanto value. The cases from section 2.2 suggest not just that conforming to epistemic norms is sometimes not good all things considered, but that there is sometimes *nothing* good in conforming to epistemic norms, i.e. that it is sometimes not even pro tanto good to do so.

EVT's second initially tempting but ultimately inadequate explanation of the Triviality intuition is what I will call the 'epistemic value' explanation:

The 'epistemic value' explanation: *The intuition is due to a failure to distinguish between epistemic and non-epistemic or practical value. Once we make this distinction, we see that although it is not practically or non-epistemically good to conform to epistemic norms in the cases above, it is still epistemically good.*

The problem with this second explanation is that merely establishing the necessary epistemic value of conforming to epistemic norms is not sufficient for vindicating EVT. To see this, note first that when it comes to value and value claims, we need to make a distinction that is analogous to the one I introduced in chapter 1 between genuine normativity and mere norm-relativity. In the case of value, this distinction is the one between what I will call *genuine* or *normativity-grounding* value on the one hand and mere *domain-relative* value on the other. The former is goodness which, if VP is true, entails normativity and normative reasons. It is the kind of value, in other words, which is such that if ϕ -ing is good in that sense, then there is a normative reason to ϕ .

On the other hand, mere domain-relative goodness only implies the trivial claim that something is good *relative to* or *according to* or *given* the norms, standards, or goals of a particular domain. Unlike genuine goodness, mere domain-relative goodness alone does not entail genuine normativity. The sole claim that something is 'domain-relatively' good does not entail that there is a normative reason to ϕ and hence does not entail genuine goodness either. Trivially, there is necessarily something *legally good* in respecting the law and something *legally bad* in breaking it. Conformity to the law is good from a legal point of view and legal violations are bad from a legal point of view. However, it is a further question whether it is necessarily

genuinely good to do what is legally good. It is a further question, in other words, whether legal value necessarily constitutes genuine, normativity-grounding value. Similarly, unfashionable outfits are, trivially, bad from a fashion point of view and fashionable outfits are good from a fashion point of view. They respectively have fashion disvalue and value. However, it is a further question whether it is necessarily genuinely good to do what is good from a fashion point of view. It is a further question, in other words, whether fashion value necessarily constitutes genuine, normativity-grounding value.

In the same way, it is trivial to claim that there is necessarily something *epistemically* good in conforming to epistemic norms. Obviously, violating epistemic norms is necessarily bad from an epistemic point of view. However, it is a further question whether it is necessarily genuinely good to do what is epistemically good. It is a further question, in other words, whether epistemic value necessarily constitutes genuine, normativity-grounding value. The point I am making at this stage is not (yet) that epistemic value does not necessarily constitute genuine value, but rather (i) that EVT is only true if it is necessarily *genuinely* good to conform to epistemic norms (since the kind of goodness that is relevant to EVT and Epistemic Normativity is normativity-grounding goodness and not mere domain-relative goodness) and (ii) that merely associating epistemic norms with *epistemic* value leaves it open whether it is necessarily genuinely good to conform to epistemic norms. To put the point differently, what the Triviality intuition suggests is precisely that it is not necessarily genuinely good to believe what is epistemically good to believe. This is why the 'epistemic value' explanation of the Triviality intuition is also a non-starter.

Fortunately for proponents of EVT, these two explanations are not the only ones available to them. A more promising avenue is to go instead for what I will call the 'minimal value' explanation

The 'minimal value' explanation: *We have the Triviality intuition because (i) the above cases are ones where conformity to epistemic norms is only very minimally good and (ii) it is easy to mistake very little pro tanto value with no value at all.*

This explanation is more promising. As I have suggested, moral and prudential value are plausible candidates for kinds of value that necessarily constitute genuine or

normativity-grounding value. However, we can probably imagine situations where the value of conforming to moral or prudential norms is sufficiently low that some might mistake it for a situation where there is nothing good in conforming to those norms.

Simply offering the ‘minimal value’ explanation is not sufficient for getting EVT off the hook however. Proponents of EVT must also show that using this explanation is justified or appropriate. This is because one could invoke the same kind of explanation in the case of any norms, including those to which it is clearly not necessarily genuinely good to conform. Once again, consider the example of legal norms or norms of etiquette. Confronted to examples where there seems to be nothing genuinely good in conforming to the law or to etiquette – e.g. cases of immoral, absurd, or obscure laws or etiquette rules – one could give a ‘minimal value’ explanation in response and say that these are simply cases where conformity to the law or etiquette is only very minimally good, which is easy to confuse with no value at all. Hence, in addition to invoking the ‘minimal value’ explanation, proponents of EVT must show us that in the case of epistemic norms, and unlike in the case of e.g. laws and etiquette, invoking that explanation is appropriate.

In the rest of this chapter, I will argue that the ‘minimal value’ explanation is *not* defensible or warranted in the case EVT. That is, there is no plausible case for the claim that examples like those from section 2.2 involve a minimal amount of genuine value rather than no genuine value at all. This is because (i) the most prominent arguments in support of the EVT and its ‘minimal value’ explanation are all unconvincing and (ii) it is vulnerable to a fundamental objection. I argue for these claims in sections 2.4 and 2.5 respectively.

2.4 Arguments for the ‘minimal value’ explanation

In this section I examine and ultimately reject five arguments in favour of EVT and its ‘minimal value’ explanation of the Triviality intuition.

2.4.1 The argument from Lynch’s thought experiments

One of the most explicit proponents of EVT is Michael Lynch. In particular, Lynch argues in his 2004 book *True to Life* that believing the truth and avoiding error is, as he puts it, worth caring about for its own sake.

One of Lynch’s claims is that, at the very least, we *believe* that truth is worth caring about for its own sake. We are committed, in other words, to the idea that

believing the truth and avoiding error is valuable for its own sake. This is because, he argues, we actually *do* care about the truth for its own sake and not just for the sake of something else. He supports this claim with the following two thought experiments:

Experience machine. Some super neuroscientists give you the choice between continuing to live normally, or having your brain hooked up to a supercomputer that will make it *seem* as if you are continuing to live normally (even though you're really just floating in a vat somewhere). When in the vat, you will continue to have all the same experiences you would have in the real world. Because of this, you would believe that you are reading a book, that you are hungry, and so on. In short, your beliefs and experiences will be the same, but most of your beliefs will be false. If we didn't really prefer true beliefs to false ones, we would be simply ambivalent about this choice. Vat, no vat; who cares? But we don't say this. We don't want to live in the vat, even though doing so would make no difference to what we experience or believe. This suggests that we have a basic preference for truth. (Lynch 2004, 17)

Russell world. Suppose that, unbeknownst to us, the world began yesterday—it *seems* older, but it isn't. If I really lived in a Russell world, as I'll call it, almost all my beliefs about the past would be false. Yet my desires would be equally satisfied in both worlds. This is because the future of both worlds unfolds in exactly the same way. [...] In other words, whatever plans I accomplish now, I would also accomplish if the world had begun yesterday, despite the fact that in that case, my plans would be based on false beliefs about the past. Yet, given the choice between living in the actual world and living in a Russell world, I strongly prefer the actual world. Of course, once "inside" that world, I wouldn't see any difference between it and the real world; in both worlds, after all, events crank along in the same way. But that is beside the point. The fact remains that thinking about the worlds only insofar as they are identical in instrumental value, *there is difference right now* between the two worlds that matters to me. Even when it has no effect on my other preferences, I – and presumably you as well – prefer true beliefs to false ones. (Lynch 2004, 18)

If Lynch is right and we actually think that the truth is valuable for its own sake, then we have the basis for a promising argument in favour of EVT. This is because the fact that we do value something is often taken to be strong evidence that it really is valuable.³⁷ The argument could be summarized as follows:

³⁷ See e.g. Whiting (2013c) for discussion.

1. We value the truth for its own sake.
2. If we value the truth for its own sake, then in all likelihood, the truth really is a valuable for its own sake.
3. Therefore, in all likelihood, the truth is a valuable for its own sake.
4. Therefore, in all likelihood, it is necessarily good to conform to epistemic norms.³⁸

There are two main problems with this argument, which both have to do with premise 1, i.e. with Lynch's initial claim that we do value the truth for its own sake.

First, even if we admit that Lynch's thought experiments provide evidence that we value the truth for its own sake, our data or evidence on that matter also includes the Triviality intuition. While we have to look at Lynch's thought experiments, we must also remember the cases from section 2.2, which suggest precisely that we do *not* value the truth for its own sake. Hence, even if Experience Machine and Russell World do provide indication that we value the truth for its own sake, they do not conclusively do so since there is also the Triviality intuition.

Second and more importantly, Lynch's thought experiments do not convincingly show what he thinks they show. Regarding Experience Machine, I agree with Lynch that most of us would choose not to hook up to the machine. However, the reason is not plausibly that it would prevent us from knowing the truth, but rather that it would rob us of many of the things we value most. If I accepted the neuroscientists' offer I would not, for instance, be in contact with the people I love anymore, I would stop being involved in their lives and they would stop being involved in mine, I would not help anyone ever again, I would stop making the world a better place for others, I would not achieve anything anymore, I would never finish my PhD, I would not witness any of the great historical events of my lifetime, I would stop travelling, going to concerts, sporting events, and so on. Granted, many of these things probably entail or require knowing the truth. However, that does not mean that I value knowing the truth for its own sake. What I value is being in contact with loved ones, helping others, finishing my PhD, travelling, and so on. Not *knowing* that I am.

³⁸ Here, I take for granted that if it is necessarily good to believe the truth, then it is necessarily good to conform to epistemic norms. This is because I assume, for the sake of argument, that conforming to epistemic norms is the best means to the goal of believing the truth and hence that if believing the truth is necessarily good, then so is conforming to epistemic norms.

Regarding the Russell World example, I should say first that it is not clear to me that we would all prefer the actual world to a world that – unbeknownst to us – just came into existence. After all, a mind-boggling amount of horrible things happened up until yesterday in the actual world. If we were in a Russell World however, none of these horrible things would actually happen. No war, no murder, no plague, no genocide, no torture, no humiliation, no injustice. Wouldn't that be great? Moreover, since we would still believe that these horrible things happened, we would be motivated to prevent them from happening 'again' and better prepared to cope with them if they happened. Clearly some of us might prefer that.

But in any case, my main worry with Lynch's second thought experiment is that even if we admit that we would prefer the actual world, it is not clear that this preference would essentially have to do with knowing the truth. I would prefer the actual world because the other option would mean that most of the things that give value and meaning to my life never actually happened. It would mean, for instance, that I never finished my MA thesis, that I never travelled around the world, that I never built a relationship with my spouse, that the Montreal Canadiens did not win 24 Stanley Cups, that I never helped or improved the life of anyone, and so on.

The point is that in both Experience Machine and Russell World, our preference for the actual world can easily be explained without invoking our valuing the truth for its own sake. Of course, that alone does not entail that we do *not*, in fact, value the truth for its own sake. What it does mean however, is that these thought experiments do not show that we do. Hence, this first defence of EVT and its 'minimal value' explanation is not persuasive.

2.4.2 Lynch's argument from intellectual integrity and sincerity

Lynch has a second argument in his book that could be used to defend EVT and its 'minimal value' explanation of the Triviality intuition. *Intellectual integrity* and *sincerity*, he argues, are necessarily good because they are partly constitutive of two further necessarily good things, namely *flourishing* and *respecting persons*. You cannot live a flourishing life and respect persons, in other words, without having intellectual integrity and sincerity respectively. Crucially however, Lynch argues, part of what it is to have intellectual integrity and sincerity is to care about the truth for its own sake. Therefore, since intellectual integrity and sincerity are necessarily good, Lynch concludes that it is also necessarily good to care about the truth for its own sake and

hence that truth really is worth caring about for its own sake. But since conforming to epistemic norms is plausibly our best way to get to the truth and since the truth is worth caring about for its own sake, doesn't it follow that it is always good to conform to epistemic norms?

The main problem with this second argument lies in the way in which Lynch argues for his conclusion that the truth is worth caring about for its own sake. The way he reaches his conclusion is *not* by showing that *believing the truth* is itself good, but rather by showing that *caring about the truth* is good. This is problematic because what the Triviality intuition denies is not that *caring* about the truth is always good, but rather that *believing the truth* is necessarily valuable. Therefore, Lynch's argument does not address the Triviality intuition since he reaches the conclusion that valuing truth for its own sake is good without ever showing that *true belief itself* is good for its own sake. To put the point differently, an argument for EVT and its 'minimal value' explanation requires *object-given* or *right kind* of reasons for valuing the truth. However, Lynch only provides *state-given* or *wrong kind* of reasons to value the truth for its own sake. He does not give evidence that truth is good, but rather that it is good to value the truth. Hence, Lynch's second argument is not sufficient to establish EVT either.

2.4.3 Kvanvig's cognitive ideal argument

Another defender of the idea that true belief is valuable for its own sake is Jonathan Kvanvig. In a 2008 paper, he argues that we can see this when we reflect on what a *cognitively ideal* agent would be like in an environment where no practical needs are left unmet. As Kvanvig explains:

[I]magine a world where no practical needs are left unmet and where no limitation of cognitive power creates any need for informational content to trump any value for truths with little or no content. [...] We should ask ourselves, regarding possible individuals in such a cost-free environment, what the cognitive ideal would involve. [...] Part of the cognitive ideal, whatever else it may involve, is knowledge of all truths; omniscience, for short. But for omniscience to be part of the ideal, no truth can be pointless enough to play no role at all in the story of what it takes to be cognitively ideal. (Kvanvig 2008, 18)

Kvanvig's argument can be summarized as follows:

1. A cognitively ideal agent in a cost free environment would know all truths.
2. The cognitive ideal under these circumstances would not consist in knowing all truths if true belief was not valuable for its own sake.
3. Therefore, true belief is valuable for its own sake.

This argument is unconvincing for at least two reasons.

First, as Carter (2011, 290) remarks, it is not clear why the cognitive ideal must have to do with perfect epistemic *successes* and not with perfect cognitive *abilities* or *capacities*. Why, in other words, would a cognitively ideal agent have to be omniscient and not, instead, one who is able to e.g. never make reasoning mistakes, know whatever she wants to know, and the like?

Second, Kvanvig's argument is either invalid or question begging. What does 'cognitive' in 'cognitive ideal' refer to exactly? If it means 'epistemic' – i.e. what is best or ideal from the epistemic point of view – then the argument is invalid. This is because as I have argued in 2.3, merely showing that something is *epistemically* good is not sufficient to show that it is genuinely good and not merely domain-relatively good. It might be true that the epistemic ideal would involve omniscience. However, it is a further question whether it is always genuinely good to do what the epistemic ideal being would do, i.e. to believe what is epistemically best to believe. Compare this with the idea of the *etiquette* ideal. Trivially, the ideal agent from the point of view of etiquette would always conform to every rules of etiquette. But that does not mean that it is always genuinely good to behave in the way that the ideal agent from the point of view of etiquette would behave. In the same way, no matter how we characterize the epistemic ideal, the crucial question is whether it is always genuinely good to do what the epistemically ideal agent would do. What the Triviality intuition suggests is precisely that it is not. In Dream, for example, it does seem that Vincent would suspend judgment if he were an epistemically ideal agent. But the point is that it does not seem to matter at all that Vincent diverges from the epistemic ideal.

But perhaps 'cognitive ideal' does not merely mean 'epistemically' ideal, but instead ideal as far as cognition is concerned *simpliciter* and not just epistemically. In other words, perhaps the first premise of Kvanvig's argument should be interpreted as the claim that an ideal agent *simpliciter* – i.e. an agent who always does what is genuinely best – always behaves as the epistemically ideal agent would behave.

However, if this is what Kvanvig means, then although his argument is valid, it is question begging since it essentially relies on what needs to be shown, namely the Epistemic Value Thesis (EVT). Rather than *establishing* that epistemic value necessarily constitutes genuine normativity-grounding value, this second reading of Kvanvig's argument uses this claim as its starting point. But this is precisely the claim that is cast into doubt by the Triviality intuition. That is, what the examples from section 2.2 suggest is precisely that there can be situations where being an ideally good agent would not require doing what is epistemically ideal. So for all these reasons, Kvanvig's cognitive ideal argument is not a promising defence of EVT either.

2.4.4 The argument from the aim of belief

A third idea that philosophers sometimes use or allude to in order to defend the necessary value of conforming to epistemic norms is the popular thesis that belief *constitutively aims at truth* (the aim of belief thesis for short).³⁹ According to the aim of belief thesis, part of what it is for a state or attitude to be a belief is for it to be, in some sense, directed or regulated towards truth.⁴⁰ One popular suggestion is that the aim of belief thesis is best interpreted as a metaphor for an evaluative claim. In particular, many take it to mean that part of what it is for something to be a belief is for it to be *good* or *well-functioning* (i.e. to be as it is supposed to be *qua* belief) if it is true and *bad* or *defective* (i.e. not as it is supposed to be *qua* belief) if it is false. Necessarily, in other words, true beliefs are good or well-functioning *qua* beliefs and false beliefs are bad or defective *qua* beliefs.⁴¹ Another closely related interpretation of the aim of belief thesis is that part of what it is to be a belief is for it to be *correct* or *fitting* if and only if it is true.⁴² But since conforming to epistemic norms is our best

³⁹ This sub-section draws in part from Côté-Bouchard (Forthcoming), especially sections 3 and 5.

⁴⁰ Or some other truth-related or epistemic goal such as knowledge, evidential support, coherence, and the like. I remain neutral regarding what belief aims at exactly. I only use truth as my example here for simplicity and because it is the most popular candidate. See for instance Williams (1973), Railton (1994), (1997), Velleman (2000), Wedgwood (2002), Burge (2003), (2010), Millar (2004), Gibbard (2005), Shah and Velleman (2005), Vahid (2006), Boghossian (2008), Steglich-Petersen (2009), Whiting (2010), (2012), (2013a), (2013b) and Littlejohn (2012a). Knowledge is the most popular alternative. See for instance Peacocke (1999), Williamson (2000), Adler (2002), Bird (2007), Sutton (2005), (2007), Huemer (2007), Hattiangadi (2010), McHugh (2011), and Littlejohn (2013), (Forthcoming b). For an overview of issues surrounding the aim of belief thesis, see McHugh and Whiting (2014) and Fassio (2015).

⁴¹ See for instance Velleman (2000), Burge (2003), (2010), Bird (2007), McHugh (2011), Littlejohn (2013), and Nolfi (2015).

⁴² Wedgwood (2002), Gibbard (2005), Shah and Velleman (2005), Boghossian (2008), Fassio (2011), Engel (2013), McHugh (2014), and McHugh and Way (Forthcoming).

way to attain the aim of belief – i.e. to have well-functioning or correct beliefs and not defective or incorrect ones – doesn't this mean that it is necessarily good to conform to such norms?

The problem is that the required conclusion does not follow from this aim of belief thesis alone. Showing that conforming to epistemic norms leads to beliefs that are good, well-functioning, correct, or fitting *qua* beliefs is not sufficient for showing that it is necessarily *genuinely* good to do so. This is because goodness or correctness *qua* member of a kind is not necessarily genuine or normativity-grounding goodness. That is, the mere fact that X is well-functioning or correct *qua* member of kind K does not entail that it is genuinely good to promote X or to be X since it does not entail that it is genuinely good to be a good member of kind K. For example, part of being a good torturer is to make one's victims suffer intensely. This is the correct or fitting way for a torturer to behave *qua* torturer. However, this does not make it genuinely good to make people to suffer intensely. Similarly, part of being a good hired killer is to murder one's victims quickly without leaving a trace. This is the correct or fitting way for a hired killer to behave *qua* hired killer. Yet this does not make killing people in that way genuinely good. Even if we admit that this is necessarily what is correct or fitting for good torturers or hired killers to do, that claim is silent on whether it is genuinely good. It is silent on whether there is any normative reason to be a good or torturer or a good hired killer in the first place.

Therefore, it is not enough for proponents of EVT to determine what a good, well-functioning, correct, or fitting belief consists in. One must also show that it is necessarily genuinely good to have well-functioning or correct beliefs and to avoid defective or incorrect beliefs. After all, even if we admit that belief aims at truth in the sense outlined above, the Triviality intuition suggests precisely that there is sometimes nothing genuinely good in having beliefs that are good or fitting *qua* beliefs and in avoiding defective or incorrect beliefs.

2.4.5 The argument from potential usefulness

A final possible argument in favour of EVT and its 'minimal value' explanation of the Triviality intuition invokes the potential usefulness of true and epistemically justified beliefs. Even if we admit that conforming to epistemic norms does not always *actually* lead to anything genuinely good, one might argue that it always *potentially* does so. For example, it might be claimed that conforming to epistemic norms necessarily

provide us with potential premises for successful good actions. It might also be the case that violating epistemic norms is necessarily potentially harmful because, for example, false and epistemically unjustified beliefs are all potential obstacles to successful good action. But if this is right, then doesn't it follow that conforming to epistemic norms does necessarily have at least a minimal amount of value?

First, it is not obvious that something being a *potential* means to good ends makes that thing good and not merely *potentially* good. If this is so, then, the potential usefulness of epistemically justified beliefs supports not EVT, but rather the weaker claim that it is always *potentially* good to conform to epistemic norms. But this would entail, at best, that there is always *potentially* a normative reason to conform to epistemic norms, which falls short of Epistemic Normativity.

Second, the argument seems to prove too much. Even if we concede that conforming to epistemic norms is always potentially useful, any false and epistemically unjustified belief is also potentially useful. Even if my belief that my train is leaving at noon is false and epistemically unjustified, my journey could be postponed to noon at the last minute. Similarly, any true or epistemically justified belief is also potentially harmful. Therefore, if the potential usefulness of conforming to epistemic norms sufficed to establish EVT, it would also establish, implausibly, that it is also always good to violate epistemic norms and bad to conform to them.

Perhaps this problem can be solved by replacing 'potentially' with 'very likely' since false and epistemically unjustified beliefs are not very likely to be useful and true beliefs are not very likely to be harmful. The problem, however, is that it is just not the case that *any* true or epistemically justified belief is likely to be useful and that any false belief is likely to be harmful. In the Dream example, for instance, Vincent's false and epistemically unjustified belief is extremely unlikely to be harmful given its trivial and inconsequential character.

Finally, it is simply not the case that conforming to epistemic norms is necessarily potentially good. Imagine that in Dream, there is an evil demon that is prepared to kill Vincent instantly if he conforms to epistemic norms and suspends his judgment about the question at hand. In such a situation, it is not even potentially useful for Vincent to conform to epistemic norms. And yet, it still seems that suspending judgment about the question at hand is the epistemically justified attitude

to have for Vincent. Therefore, this fifth argument in favour of EVT is not persuasive either.⁴³

2.5 An objection to the ‘minimal value’ explanation

In the previous section, I considered and rejected five prominent arguments in favour of EVT and its ‘minimal value’ explanation of the Triviality intuition. In this section, I argue that there is a fundamental problem with the idea that it is necessarily genuinely good to conform to epistemic norms. The problem, in a nutshell, is that there is no sense ‘good’ in which EVT follows. On each of the main possible senses of ‘good’, in other words, it is not plausible to conclude that it is necessarily genuinely good to believe the truth, avoid error, follow the evidence, and the like.

2.5.1 Three senses of ‘good’

Philosophers standardly distinguish between three main senses of ‘good’. First, something can be good in the sense of being good *for* someone. That is, something can be good in the sense that it makes one’s life go well or that it contributes to one’s *well-being* or *interest*. Second, ‘X is good’ is sometimes used to mean that X is *plainly* good or good *simpliciter* even though it might not be good *for* anyone. This is commonly referred to as the *predicative* sense of ‘good’ since in this sense, ‘good’ is a predicative adjective like ‘red’. In the same way that ‘X is red’ means that X has the property of redness, which all red things have in common, the predicative sense of ‘X is good’ means that X has the property of goodness, which all good things have in common. Finally, ‘good’ can also be used to mean that something is good or well-functioning qua member of a kind K. Similarly, ‘bad’ can also be used to mean that something is bad or defective qua member of K. To be a good or bad member of K in this sense is to do or to be unable to do what members of K are *supposed* to do. It is, in other words, to achieve or to be unable to achieve the constitutive, essential, or defining aim, function, or telos of members of that kind. This is commonly referred to as the *attributive* sense of ‘good’ since in this sense, ‘good’ is an attributive adjective like ‘big’ or ‘tall’. We are using this attributive sense of ‘good’ when we judge, for instance, that X is a good toaster, a good football player, a good teacher, and so on.

⁴³ See also Wrenn (2010) for discussion.

2.5.2 EVT and good *for*

According to EVT, it is necessarily good to conform to epistemic norms. But in which of these three senses ‘good’ is conforming to epistemic norms supposed to be necessarily good? As the examples from 2.2 already indicate, it is not always good *for* us to believe the truth, avoid error, follow the evidence, and the like. In *Dream*, the matter at hand is so trivial and disconnected from any of Vincent’s concerns and interests that conforming to epistemic norms and suspending his judgment about whether the event in question really happened would have had absolutely no impact on his well-being. Similarly, in *Even Stars*, Myriam’s violating epistemic norms and believing that the number of stars is even only has positive impacts on her well-being and interests. Having that epistemically unjustified belief seems to only make her life go better. Finally, in examples like *Illness* and the one at the beginning of chapter 1, not only would violating epistemic norms would not only be overwhelmingly in the agent’s interests, conforming to epistemic norms would be highly detrimental to their well-being. Believing what is epistemically justified would only make their lives go dramatically worse.

Moreover, there is empirical evidence that some violations of epistemic norms are actually good for us. In particular, there is evidence that some of the cognitive biases and reasoning mistakes that we constantly make have a clear positive impact on our well-being and flourishing. This is best illustrated by the phenomenon of *positive illusions*. Positive illusions are, roughly, beliefs which are incorrectly and unjustifiably optimistic. They are overestimations of things like one’s capacities, one’s control over one’s life, and how good one’s future will be.⁴⁴ There is evidence not only that such epistemically unjustified and irrational beliefs are extremely common, but also that they promote people’s well-being and flourishing. As Taylor and Brown explain, “[t]he mentally healthy person appears to have the enviable capacity to distort reality in a direction that enhances self-esteem, maintains beliefs in personal efficacy, and promotes an optimistic view of the future.” (Taylor and Brown 1988, 34) Similarly, Bortolotti writes:

Not only do positive illusions make an agent feel better about herself, but they can also enhance her health and have other evolutionarily relevant benefits, contributing to her chances of reproduction and survival. Shelley Taylor and

⁴⁴ See e.g. Taylor (1989) and Bortolotti (2015, 130-133) for discussion.

Jonathon Brown (1988) have shown that there are strong links between positive illusions about health prospects and the promotion of psychological health (in terms of creativity and productivity), and physical health (in terms of prolonged longevity). In the context of relationships and parenting, when an agent has an inflated conception of the good qualities of herself and of her loved ones, the cohesion of her familial and social relationships becomes stronger, leading to increased mutual support. (Bortolotti 2015, 132-133)⁴⁵

Since positive illusions involve violating epistemic norms and since there is evidence that at least some such illusions are essential to our well-being, it is highly doubtful that conforming to epistemic norms is necessarily good for us and that violating such norms is necessarily bad for us.

2.5.3 EVT and good *simpliciter*

Even if it is not necessarily good *for* us, could conforming to epistemic norms be necessarily good *simpliciter*? When it is not good for anyone to conform to epistemic norms, does it still necessarily lead to plainly good states of affairs? This seems to be what many proponents of EVT have in mind, especially those, such as Lynch and Kvanvig, who claim that true belief and knowledge are *intrinsically* good or good *for their own sake*.⁴⁶ This second option is also problematic however.

For starters, some philosophers are sceptical of the very idea that things can be good *simpliciter* or plainly good. In particular, some like Thomson (2008) have argued that 'good' cannot function as a predicative adjective like 'red'. To say that something is a red X is to say that it has two properties: being X and being red. So as Thomson explains, from the fact that X is e.g. a red car and that X is a Mercedes, we can infer that X is a red Mercedes. However, she points out that we cannot make the same kind of inference from the premise that something is a good X. For example, from the fact that X is a good tennis player and that X is a chess player, we *cannot* infer that X is a good chess player. Examples like these, she argues, show that 'good' is not a predicative adjective like 'red' and thus that we should be sceptical of the claim that there is a property of goodness that all good things have.⁴⁷

⁴⁵ See also Bandura (1989), Taylor et al. (1992), (2003), Taylor and Brown (1994), McKay and Dennett (2009), and Hood (2012).

⁴⁶ See also Zagzebski (2003) and Horwich (2006).

⁴⁷ Thomson (2008, 3-5). See also Geach (1956). See Almotahari and Hosein (2015) for discussion.

Moreover, this second approach is problematic even if we put Thomson's worry aside and assume that things can be good simpliciter. First, if knowledge and true beliefs were good simpliciter, it would mean that the more instances of knowledge and true belief there are in the world, the better things are other things being equal. That is, it would mean that even when it does not make things better for anyone, more knowledge and more true beliefs necessarily makes the world a better place, other things being equal. Similarly, it would mean that *losing* knowledge and true beliefs necessarily makes the world a worse place other things being equal, and that missing out on possible knowledge and true belief is necessarily plainly bad or regrettable, even when it does not make things worse *for* anyone. But this is implausible. As Littlejohn (Forthcoming d) writes:

We don't think that it's regrettable that someone knows p when that knowledge precludes knowing some other things. (For example, if you know that you are a star pupil, you know that you'll never know that you don't know much about history, that you'll never know that you don't know much about biology, that you'll never know that you don't know much about science books, etc.). (Littlejohn Forthcoming d, 15)

Similarly, suppose I have the choice between two equivalent routes A and B to go to work but that by taking A, I got to know one more truth than if I had chosen B, namely:

(T) that there are two leaves on one of the branches of the third biggest tree on the block.

Let us suppose, plausibly, that this knowledge is too trivial and inconsequential to be in any way good *for* me. More generally, let us suppose that apart from this additional knowledge, the world would have been exactly the same if I had chosen B.

Is the world really a better place than it would have been just because I got this additional piece of trivial knowledge? Would it have been at all regrettable if I had chosen the other option and missed out on this single additional piece of useful knowledge? Relatedly, will it make the world a worse place if I end up forgetting this information and never remember it? It is hard to see why anyone would answer positively to any of these questions.

Second, the claim that conforming to epistemic norms is necessarily good simpliciter leads to at least three absurd or deeply counter-intuitive conclusions. That is, it is vulnerable to at least three *reductio* arguments. First, if conforming to epistemic norms were necessarily good simpliciter – i.e. if necessarily led to plainly good states of affairs – then given the Value Principle (VP), there would necessarily be a normative reason to go change the world to make it fit our beliefs and turn our beliefs into knowledge. In particular, there would necessarily be such a normative reason precisely and only because it would make our belief true or an instance of knowledge. As Joseph Raz (2011) puts it:

[I]magine that in all cases, if we have a belief about a certain matter then it is *pro tanto* better to have a true rather than a false belief, just because it is true. Consider an example: A month ahead of time I believe that Red Rod will win the Derby or that the Social Democrats will win the elections in Denmark. There may be ways to increase the likelihood that my belief is true. Perhaps I could give valuable advice to Red Rod's jockey, or lend my expertise to the Social Democrats. Is the fact that that will make it more likely that my beliefs are true a reason to do so? If there is value in one's beliefs being true as such then there should be no difference between making reality conform to the belief and making the belief conform to how things are. Raz (2011, 45)

But this is a deeply counter-intuitive result. It does not seem like there is necessarily a normative reason to act so as to make our beliefs true *just because* it would make them true and give us knowledge.

Second, the 'good simpliciter' option seems to lead to the conclusion that there is necessarily a normative reason for anyone to do things that they will *regret*. This second *reductio* argument stems from the factivity of regret, i.e. from the fact that part of what it is to regret that you ϕ -ed is *knowing* that you ϕ -ed. You cannot count as regretting that you ϕ -ed, in other words, unless you know that you ϕ -ed. So if knowing that P is necessarily plainly good, it follows that there is necessarily something plainly good in regretting that you ϕ -ed, namely the knowledge that you ϕ -ed. What this means, however, is that doing things that you will in fact regret necessarily brings about something good, i.e. knowledge. Two seemingly absurd and counter-intuitive conclusions seem to follow. First, doing things that you will in fact regret is necessarily (instrumentally) good. Second, given the Value Principle (VP), it

follows that there is necessarily a normative reason to do things that you will in fact regret.

A third seemingly unacceptable conclusion that the present approach leads to is that there is necessarily a normative reason for anyone to do *anything*. This third reductio argument starts from the widely held thesis that you cannot count as ϕ -ing intentionally if you do not know that you are ϕ -ing. As Kieran Setiya explains:

[...] if I have no idea that in humming Beethoven's Ninth I am driving my wife crazy, I simply cannot be driving her crazy *intentionally* – at least not so far as my humming goes. And if I am ignorant of the impatient tapping of my foot, as I pore over a draft of these pages, it too must be unintentional. (Setiya 2007, 25)

So part of what it is to ϕ intentionally is *knowing* that you are ϕ -ing. Once again, if knowledge is necessarily plainly good, it follows that there is necessarily something good in ϕ -ing intentionally, namely the knowledge that you are ϕ -ing. What this means, however, is that ϕ -ing intentionally necessarily brings about something good, i.e. knowledge. Two seemingly absurd and counter-intuitive conclusions follow. First, it is necessarily (instrumentally) good to intentionally do *anything*. Second, given VP, there is necessarily a normative reason to intentionally do anything.

Of course, one could reply to these last two reductio arguments by rejecting the claim that regret and intentional action both essentially involve knowledge. But even if these two premises are mistaken, knowledge is still part of *most* cases of regret and intentional action. But this weaker claim is sufficient to generate an almost equally counterintuitive conclusion, namely that there is *almost always* a normative reason for anyone to things they will regret and to do anything.

Another tempting reply is to bite the bullet and claim that there is, in fact, necessarily a normative reason to make our beliefs true, do things you will regret, and intentionally do anything. However, even if these bullets can be bitten, they remain deeply counter-intuitive conclusions. So the fact that EVT's 'minimal value' explanation entails such counterintuitive conclusions still constitutes an important cost or disadvantage of that explanation. This is important because EVT is up against an alternative explanation of the Triviality intuition, which does not have these important intuitive costs, i.e. the explanation according to which we have the Triviality

intuition because there really are cases where it is not genuinely good to conform to epistemic norms. The problem for the 'minimal value' explanation is that given what we have said so far, it is hard to see why it would be preferable to this other, more straightforward explanation.

So for all these reasons, it is not plausible that conformity to epistemic norms is necessarily good simpliciter either.

2.5.4 EVT and *attributive* goodness

Given the implausibility of EVT in terms of 'good for' and 'good simpliciter', we are left with the possibility that conforming to epistemic norms is necessarily *attributively* good. But the problem is that as I have already indicated in section 2.4.4, attributive goodness does not suffice to vindicate EVT and Epistemic Normativity. To see this, consider the possible ways in which it could be necessarily attributively good to conform to epistemic norms. I have already examined one in section 2.4.4, namely the aim of belief thesis according to which beliefs that conform to epistemic norms are necessarily good, well-functioning, correct, or fitting *qua* beliefs and that false or epistemically unjustified beliefs are necessarily bad, defective, or incorrect *qua* beliefs. However, as I already explained, even if we concede this, EVT and Epistemic Normativity do not follow since it does not follow that it is always genuinely good to have well-functioning beliefs and genuinely bad to have defective beliefs. What the Triviality intuition suggests is precisely that it is not the case. The same problem arises if we claim instead that agents that conform to epistemic norms are necessarily good *qua* *believers*, *qua* *inquirers*, *qua* *epistemic agents*, and the like. These claims might be correct, but they only entail domain-relative goodness since they leave open the question whether it is always genuinely good to be a good believer, a good inquirer or a good epistemic agent. What the examples from 2.2 suggest is precisely that it is not necessarily genuinely valuable to be a good believer, inquirer, epistemic agent, and the like.

One possible reply is that we can avoid this problem by claiming that conforming to epistemic norms is necessary for being a good or well-functioning *agent* or *human being* and that violating epistemic norms necessarily makes us a bad or defective agents to some extent. This alternative claim might avoid the previous problem because it might be implausible or incoherent to doubt that it is necessarily

genuinely good to be a well-functioning agent and genuinely bad to be a defective agent. This reply is also problematic however.

First, it is not at all clear why the question whether to be a well-functioning agent must be incoherent or implausible. One common claim is that it is incoherent or implausible to ask whether there is any normative reason *to be an agent* since, one might claim, normative questions like this one already presuppose agency or can only be asked within agency. If you are asking normative questions, then you are already an agent in the first place and so, one might claim, the question whether you should be an agent cannot be a coherent or live one.⁴⁸ However, the same kind of story cannot be told about good or well-functioning agency. Even if you need to be an agent in order to ask normative questions, you do not have to be a good or well-functioning agent to do so. Hence, even if we admit that conforming to epistemic norms is necessary for being a good or well-functioning agent, this still seems to entail only domain-relative value and not genuine value since it still leaves open the question whether it is always genuinely good to be a well-functioning agent and always genuinely bad to be a defective agent.

Second, even if we concede that it is necessarily genuinely good to be a well-functioning agent and bad to be a defective agent, it is simply implausible that violating epistemic norms necessarily makes people defective qua agents and that conforming to epistemic norms is always necessary for being a good or well-functioning agent. As we have already seen, to say that *X* is a good member of a kind *K* is to say that it does what members of that kind are supposed to do. It is to say, in other words, that it can achieve the constitutive, essential, or defining aim, function, or *telos* of members of that kind. It can fulfil, in other words, the function that makes it a member of that kind. But what is the essential aim, function, or *telos* of agency supposed to be? When are agents defective or well-functioning qua agents?

One natural answer is that agents are well-functioning insofar as they *flourish* and they are defective insofar as they do not flourish. It is not plausible, however, that conforming to epistemic norms is always necessary for human flourishing. As we have already seen, conforming to epistemic norms is not necessarily good *for us* and violating epistemic norms is not necessarily bad for us. But it is hard to see how ϕ -ing could be necessary for our flourishing without being always good for us. So the

⁴⁸ Although this claim is also, I think, unpersuasive for reasons raised most clearly by Enoch (2006), (2011c).

reasons why it is not necessarily good for us to conform to epistemic norms are also plausibly reasons why it does not necessarily contribute to our flourishing.⁴⁹

Another natural answer is that agents are well-functioning insofar as they are *autonomous* or can intentionally do what they think they ought to do. But one can very well violate some epistemic norms and still be autonomous in that way. Failing to, say, follow the evidence and avoid error about particularly trivial and insignificant questions does not necessarily interfere with one's autonomy or ability to act as one thinks one ought.

So for all these reasons, the claim that it is always attributively good to conform to epistemic norms is of no help to the Epistemic Value Thesis (EVT) and to Epistemic Normativity. Therefore, there is no sense of 'good' in which it EVT follows, i.e. no sense of 'good' in which it is plausible to conclude that it is necessarily genuinely good – and not merely domain-relatively good – to conform to epistemic norms.

2.6 Summary

A first commonly cited mark or feature of normative facts and claims is a necessary connection with value. This is captured by what I called the Value Principle (VP) according to which there is a normative reason to ϕ if and only if it is good to ϕ . As I have explained, if VP is true and if, as Epistemic Normativity maintains, there is necessarily a normative reason to conform to epistemic norms, then it must be necessarily good to conform to epistemic norms. This is what I called the Epistemic Value Thesis (EVT). I argued that EVT is not acceptable unless it can plausibly explain the Triviality intuition and that the most promising avenue for EVT is the 'minimal value' explanation. However, I have shown that there is no plausible case for accepting the 'minimal value' explanation of the Triviality intuition. This is because (i) the most prominent arguments in its support are all unconvincing and (ii) it is vulnerable to a fundamental problem or objection, namely that there is no sense of 'good' in which EVT follows. Therefore, I conclude that the EVT is false and therefore that epistemic facts and claims do not bear the commonly cited mark of normativity captured by VP.

⁴⁹ See also Hazlett (2013, part I).

Chapter 3: Desire

3.1 Introduction: epistemology and desires

A second feature that is commonly associated with normativity, but not with mere norm-relativity is a necessary connection with *desire*. Unlike merely norm-implying facts, the thought goes, normative facts essentially have to do with what promotes the satisfaction of goals or aims that agents desire or care about.⁵⁰ I will take this second commonly cited mark of normativity to be captured by the following principle:

Instrumental Principle (IP): *There is a normative reason for one to ϕ if and only if ϕ -ing promotes the satisfaction of one's desires.*⁵¹

This second feature is most readily associated with what is sometimes called *desire-based, internalist, neo-Humean, or subjectivist* accounts of normativity and normative reasons.⁵²

According to desire-based accounts, normativity is grounded or explained by facts about our desires. Roughly, whenever there is a normative reason for S to ϕ , this is because or in virtue of the fact that ϕ -ing would promote the satisfaction of a desire that S has. Suppose I want to sing karaoke and tonight happens to be karaoke

⁵⁰ I use verbs like 'desire', 'want', and 'care about' interchangeably in what follows. I also take 'goals', 'ends', and 'aims' to be synonymous. I use 'S's desire(s)' to refer to ends that S desires or wants to achieve. It is important to note that the ends that are relevant to this second commonly cited feature of normativity are those that we desire or possess. Travelling to the best boomerang shop in Australia promotes the achievement of the end or goal of getting a top quality boomerang, but there is no normative reason for me to travel to that shop since I do not possess that end. I have absolutely no desire to get a top quality boomerang. Things would be different, however, if I desired to get a top quality boomerang or if getting a top quality boomerang would lead to the satisfaction of another of my desires.

⁵¹ It is important that this second feature of normativity be understood in terms of a *necessary* connection with desire satisfaction and not as the weaker claim that ϕ -ing must e.g. *very often* or *almost always* promote the satisfaction of our desires. A weaker condition like the latter would be easily met by many merely norm-implying facts and claims. After all, conforming to norms of e.g. etiquette and fashion very often promotes the satisfaction of our desires. It does not *necessarily* do so however.

One obvious worry, however, is that it is far from clear that even morality is normative according to IP. This is a well-known problem for accounts of normativity associated with IP such as what I call *desire-based* or *neo-Humean* theories in the next paragraph. However, I set this problem aside in this chapter since (i) this is a prominent approach in contemporary metaethics and normativity theory and (ii) several authors have suggested that the normative authority of epistemic norms might be desire-based.

⁵² See for instance, Foot (1972), Williams, (1979), Smith (1994), Tiberius (2000), Joyce (2001), Schroeder (2007), Street (2008), Goldman (2009), Manne (2014), and Markovits (2014). Note that many of these authors put constraints on which of S's desires can ground normative reasons. In particular, many think that it is not our actual desires that matter for normativity, but rather the ones that we would have under certain idealized circumstances. I return to this issue below.

night at my local pub. In such a situation, it seems that the fact that tonight is karaoke night at the pub is a normative reason for me to go to the pub. At the same time, since my wife hates karaoke and wants to avoid karaoke nights, it is natural to conclude that there is no normative reason for her to go to the pub. If someone were to ask why the fact that tonight is karaoke night at the pub is a normative reason for me to go to the pub and for my wife not to go, we would naturally respond that this is because I want to sing karaoke and my wife does not.

Given IP, if epistemology is normative and there is necessarily a normative reason to conform to epistemic norms – as Epistemic Normativity maintains – then conforming to epistemic norms must necessarily promote the satisfaction of our desires. That is, if IP and Epistemic Normativity are both true, then the following must also be true:

Epistemic Instrumental Thesis (EIT): *conforming to epistemic norms necessarily promotes the satisfaction of one's desires.*

The route from IP to EIT can be summarized as follows:

1. There is a normative reason for one to ϕ if and only if ϕ -ing promotes the satisfaction of one's desires. (*IP*)
2. There is necessarily a normative reason to conform to epistemic norms. (*Epistemic Normativity*).
3. Therefore, conforming to epistemic norms necessarily promotes the satisfaction of one's desires. (*EIT*)

EIT is closely related to what epistemologists and philosophers of science call *instrumentalist* conceptions of epistemic rationality.⁵³ According to epistemic instrumentalists, epistemic rationality is a distinctly epistemic species of instrumental or *means-end* rationality.⁵⁴ To be epistemically rational and conform to epistemic

⁵³ See, for instance, Foley (1987), (1993), Maffie (1990), Laudan, (1990a), (1990b), Nozick (1993), Kornblith, (1993), (2002), Papineau (1999), and Leite (2007). See Kelly (2003), (2007), Cuneo (2007, ch.7), Lockard (2013), and Côté-Bouchard (2015) for discussion.

⁵⁴ As Lockard (2013, 1702) notes although most instrumentalists only mention the notion epistemic rationality, they typically use 'rational' as a generic term for positive epistemic evaluation and conformity to epistemic norms.

norms, the thought goes, is essentially to take the means to one's epistemic goals like believing the truth, avoiding error, following the evidence, and the like.⁵⁵

The relation between epistemic instrumentalism and EIT is not always clear however. In particular, it is often unclear whether instrumentalists are concerned with vindicating Epistemic Normativity – i.e. with providing an instrumental or desire-based account of the necessary normative authority of epistemic rationality and epistemic norms – or only with giving a first-order account of epistemic rationality and epistemic norms. On the former interpretation, epistemic instrumentalism is essentially a defence of Epistemic Normativity on the basis of both IP and EIT.⁵⁶ On the latter, it is in principle compatible with Epistemic Non-Normativity and the falsity of EIT. This is so if, for instance, being epistemically rational is a matter of taking the means to cognitive goals *period*, whether or not you *have* or *desire* these goals. An instrumentalist in the latter sense could claim, in other words, that being epistemic rational necessarily promotes the achievement of cognitive goals, but that we do not necessarily want to achieve those goals, and thus that there is not necessarily a normative reason to be epistemically rational.⁵⁷

In any case, whether or not epistemic instrumentalism goes hand in hand with EIT, I will only be concerned with the latter in this chapter since it is the thesis that matters for Epistemic Normativity given IP. I will argue that EIT is false. Conforming to epistemic norms does *not* necessarily promote the satisfaction of our desires. S's belief that P can very well be epistemically justified, rational, and the like, in cases where believing that P would *not* promote the satisfaction of any of S's desires. Therefore, this second commonly cited mark of normativity is also missing in epistemic facts and claims.⁵⁸

My strategy against EIT is similar to the one I used against the Epistemic Value Thesis (EVT) in chapter 2. My starting point, which I introduce in section 3.2, is what I call the *No Desire intuition*. Intuitively, there seems to be cases where conforming to epistemic norms does not promote the satisfaction of any of our desires. Since this intuition directly contradicts EIT, the latter is only viable if it can plausibly explain away the former. The problem is that it cannot do so. I show this by examining and

⁵⁵ This is roughly Kelly's (2003) characterization of the approach.

⁵⁶ This seems to be Kelly's (2003) and Lockard's (2013) interpretation. This is also how I interpret epistemic instrumentalism in Côté-Bouchard (2015).

⁵⁷ This seems to be what Maffie (1990) and Laudan (1990a), (1990b) claim.

⁵⁸ The content of this chapter draws in part from Côté-Bouchard (2015).

rejecting the most promising or prominent strategies for meeting the challenge posed by the No Desire intuition in sections 3.3 – 3.6.

3.2 The No Desire intuition⁵⁹

According to the Epistemic Instrumental Thesis (EIT), conforming to epistemic norms necessarily promotes our desires. However, an immediate problem for EIT is what I will call the *No Desire* intuition.

The No Desire intuition: *there seem to be cases where, intuitively, conforming to epistemic norms does not promote the satisfaction of any of our desires.*

That is, there are cases where, intuitively, S's having a particular doxastic state D does not promote S's desires, but where epistemic norms still clearly require S to have D – i.e. where D is still epistemically justified, rational, and the like.⁶⁰

Kelly (2003) suggests two kinds of problematic cases for theses like EIT. First, there seems to be cases where conforming to epistemic norms would not promote agents' desires because they are utterly indifferent about the matter at hand. Call these cases of *epistemic indifference*. Second, there are also cases in which conforming to epistemic norms does not promote agents' desires because they want to *avoid* learning the truth about P. Call these cases of *truth-avoidance*. Here is an example of epistemic indifference:

Delaware's Beverage Like most people, Nancy has absolutely no desire to know what the official beverage of the state of Delaware is. Unbeknownst to her however, her friend Brett, whom she knows to be very reliable, has recently developed a deep obsession for Delaware. One day he comes up to her, grabs her by the shoulders and says: 'Listen to me, Nancy. I've got to tell you something. I just found out that the state beverage of Delaware is milk! Isn't that amazing?'

⁵⁹ This section draws in part from section 2 of Côté-Bouchard (2015).

⁶⁰ This is similar to Tom Kelly's (2003) argument against instrumentalist conceptions of epistemic rationality. See also Cuneo (2007, ch.7) and Lockard (2013).

Here is an example of truth-avoidance:

Spoiler Alert Max missed the finale of his favourite television series, which aired last night and revealed whether character X ends up dying or surviving. Given his passion for the show, Max really wants to avoid learning the answer to that question before he watches the rerun tonight after work. But to his dismay, one of his colleagues is convinced that Max did watch the show last night and comes up to him in the morning shouting: ‘I knew it! I told you character X would die!’

Nancy and Max are both epistemically justified in believing the proposition at hand – i.e. that milk is Delaware’s beverage and that character X dies respectively. It would also be epistemically unjustified for them to disbelieve the relevant proposition or to suspend judgment about the matter. Yet conforming to epistemic norms would not, on the face of it, promote any of Nancy and Max’s desires. After all, Nancy is completely indifferent about the state beverage of Delaware and Max has a clear desire not to have his favourite show spoiled for him.

The cases from chapter 2 also support the No Desire intuition. In *Illness*, Rita’s conforming to epistemic norms and believing that she will probably not survive her illness would plausibly only frustrate her desires since it would most likely lead to her death. In *Even star*, it seems that Myriam’s avoiding the epistemically unjustified belief that the number of stars is even would not promote any of her desires. It would be completely useless for her, after all, to have the epistemically justified doxastic attitude regarding the number of stars. In fact, it would plausibly only frustrate her desires given the happiness she gets from her epistemically unjustified belief. Finally, in *Dream*, Vincent believes falsely and unjustifiably that he really once saw a historian on TV saying that a pub in London named The Red Lion closed its doors on February 1st 1748. But given the complete triviality and unimportance of the issue at hand, it clearly seems that not forming that belief and instead suspending judgment would not promote any of Vincent’s desires. Similarly, forming the epistemically unjustified doxastic attitude plausibly won’t lead to the frustration of any of his desires.

So there are clearly cases where, intuitively, conforming to epistemic norms does not promote any of one’s desires. The Epistemic Instrumental Thesis (EIT) therefore faces a pressing challenge: it must explain away these cases and show that

contrary to what the No Desire intuition suggests, it really does necessarily promote the satisfaction of our desires to, say, believe the truth, avoid error, and follow the evidence. This is a pressing challenge because the intuition suggests precisely the falsity of EIT i.e. that conforming to epistemic norms does not necessarily promote the satisfaction of one's desires. Therefore, in order to have any plausibility, EIT must be able to satisfactorily explain away this intuition.

In the rest of this chapter, I argue that there is no plausible way for EIT to do so. More precisely, I reject all of the most promising or prominent strategies for explaining away the No Desire intuition.

3.3 Schroeder's *potential desire-satisfaction strategy*

As we have seen in chapter 2, one tempting argument in favour of the necessary goodness of conforming to epistemic norms is that any true or epistemically justified belief, no matter how apparently trivial, is *potentially* useful (and any false or epistemically unjustified belief is *potentially* hurtful). In the same way, one might suggest, any true or epistemically justified belief could *potentially* lead to the satisfaction of our desires even though they do not always actually do so. Conversely, any false or epistemically unjustified belief could potentially lead to the frustration of one's desires. Hence, one might conclude that, given a sufficiently weak understanding of what it is for desire satisfaction to be promoted, conforming to epistemic norms does, in fact, necessarily promote the satisfaction of one's desires.

Mark Schroeder (2007) has provided the clearest articulation of this strategy.⁶¹ He argues that because of the interconnectedness of our beliefs, believing falsely any proposition can potentially start a chain of errors that could end up frustrating any of one's desires. But since conforming to epistemic norms is one's best means to avoid error and believe the truth, doing so necessarily promotes *all* of one's desires to some extent since any such desires might potentially be frustrated as a result of *any* error. Any of one's desires, in other words, is promoted by avoiding error with respect to *any* proposition. To use one of Schroeder's examples, take Mary's desire to buy a new pair of shoes on the one hand, and the question of the number of moons that Jupiter has on the other. According to him:

⁶¹ See especially Schroeder (2007, 113–15). This section draws in large part from section 3 of Côté-Bouchard (2015).

Being in error about [how many moons Jupiter has] might lead to being in error about other things, such that being in error about them might lead to being in error about other things, and so on until something might lead to Mary having trouble getting new shoes. If this is right, then for any proposition, Mary's desire to get a new pair of shoes will serve to explain why there is a reason for Mary to believe it only if it is true. (Schroeder 2007, 114)

Schroeder's striking claim is not only that Mary's desire to get new shoes is promoted by avoiding error with respect to the number of Jupiter's moons, but that it is promoted by avoiding error with respect to *any* proposition. Moreover, avoiding error regarding the number of Jupiter's moons promotes not only Mary's desire to get new shoes, but also each and every one of her desires.

As I mentioned above, the other key component of such a strategy is an extremely weak understanding of the promotion relation, i.e. of what it takes for ϕ -ing to promote the satisfaction of one's desires. Here is how Schroeder construes the promotion relation:

X's doing A promotes p just in case it increases the likelihood of p relative to some baseline. And the baseline, I suggest, is fixed by the likelihood of p conditional on X's doing nothing – conditional on the status quo. (Schroeder 2007, 113)

This strategy is problematic for two reasons.

First, even if we admit that this strategy successfully explains away the cases in 3.2, we can modify them to make them immune from this explanation. Imagine, for instance, that in Delaware's beverage, an evil demon will kill Nancy instantly unless she *disbelieves* what Brett tells her. Even given Schroeder's weak account of the promotion relation, believing the truth or avoiding error in such a situation would *not* promote the satisfaction of any of Nancy's desires. Only disbelieving Brett would allow her desires to be satisfied. Crucially, however, it would still be epistemically unjustified for her to disbelieve him. She would still count as violating epistemic norms.

If you find this example too outlandish, consider instead cases of agents forming beliefs the instant before they die. Imagine, for instance, that Brett tells Nancy about Delaware's beverage right before they hit the ground after having jumped off a plane without a parachute. Even if we grant the premises of Schroeder's strategy,

Nancy's conforming to epistemic norms won't promote the satisfaction of any of her desires since she will certainly die the second after. Yet she would still be epistemically justified in believing him and epistemically unjustified in disbelieving him or suspending judgment.

Second, even if there is a way for people like Schroder to deal with these modified cases, this strategy for explaining away the No Desire intuition has at least two deeply counter-intuitive implications. One is that it would entail, implausibly, that there is also necessarily a normative reason to *violate* epistemic norms. Although it is true that any chain of errors can potentially end up frustrating one's desires, any such chain can also potentially end up *satisfying* one's desires. False and epistemically unjustified beliefs can, after all, lead to the satisfaction of our desires. Imagine the following elaboration of Schroeder's Mary case:

Coincidence Mary ends up believing, incorrectly and without any evidence, that Jupiter has 35 moons. This fascinates her because she is 35 years old and she is obsessed with coincidences. Because she also believes – incorrectly and unjustifiably – that such coincidences mean something deep and important about her life, she then goes outside to take a walk and think about what it means and to see if she can spot other meaningful signs from the universe. While taking her walk, she wanders into a part of town that she did not know about and discovers a newly opened shoe shop which happens to have the brand of shoes she could not find anywhere else.

Similarly, any epistemically unjustified or incorrect belief can potentially lead to further *true* beliefs which might, in turn, lead to the satisfaction of one's desires. For example, imagine that in Coincidence, Mary's false belief leads her to form the true belief that *Jupiter does not have 34 moons* and that one of the questions at her pub quiz later is whether Jupiter has 34 moons. Finally, any true or epistemically justified beliefs also has the potential to start a chain of beliefs that will end up *frustrating* one's desires. In the Illness case from chapter 2, for example, Rita's following her evidence and believing that she will probably die is very likely to frustrate her desire to survive her illness.

The other counter-intuitive implication of this strategy is what McPherson (2012) calls an *explosion* of normative reasons. That is, it would entail, implausibly, that there is a normative reason for anyone to do anything. As McPherson explains:

Consider again the desire to get home on time. For just about anything that I might do – scratching my ear, running the other way, trying to drive my car off a cliff – there is some possible scenario in which this act leads me to get home on time. (McPherson 2012, 447)

For any act A and desired end E, in other words, doing A would increase the likelihood of E being satisfied since we can always come up with a possible scenario (however remote) in which A leads to E being realized. In McPherson's example, his desire to get home on time thus makes it the case that he has a normative reason to scratch his ear, run the other way, try to drive his car off cliff, and so on.

Why is such an explosion be problematic? For one thing, the fact that there are things that there is no reason for us to do seems like a truism or an intuitive datum that any account of normativity should try to accommodate. Is there really a reason for me to e.g. jump out of my window while eating two bananas and singing happy birthday to John Travolta in a batman costume? More straightforwardly, is there really a normative reason for anyone to torture, humiliate, disrespect, or murder friends and relatives? This is obviously highly counter-intuitive. This also trivializes normativity. If it is so easy for considerations to have normative force, then why do normative reasons matter? Relatedly, explosion is also incompatible with the idea, which is at the centre of this thesis and of much of contemporary metaethics and normativity theory, that most norms lack necessary normative authority. If there is necessarily a reason for anyone to do anything, then there is necessarily a normative reason to conform not only to norms of morality, prudence, and epistemic norms, but also to those of fashion, etiquette, tradition, grammar, and so on.

Of course, one could bite the bullet and concede that there is necessarily a normative reason to violate epistemic norms and to do anything. This, however, would make this defence of EIT unattractive on balance. This is because adopting it would mean giving up these two intuitions *in addition* to the initial No Desire intuition. If we can only explain away the No Desire intuition by making these additional counter-intuitive claims, then why not simply take the No Desire intuition at face value in the first place and avoid these two further counter-intuitive implications? This

question is made all the more pressing by the fact that these two implications – i.e. explosion and the idea there is necessarily a reason to violate epistemic norms – are markedly *more* counter-intuitive than the claim that conforming to epistemic norms does not necessarily promote the satisfaction of our desires i.e. the denial of EIT. After all, whereas many – if not most – metaethicists and normativity theorists would deny EIT, very few would deny that there are many things that there is just no normative reason to do and that there is not necessarily a normative reason to violate epistemic norms.

Thus, for all these reasons, the strategy of appealing to potential desire satisfaction is not a promising way to explain away the No Desire intuition.

3.4 *The aim of belief strategy*

As I noted in chapter 2, philosophers sometimes use the idea that belief constitutively aims at truth to support the necessary goodness of conforming to epistemic norms. This is because the aim of belief thesis is often interpreted as an evaluative claim according to which beliefs are good, correct, or fitting qua beliefs if they are true and bad, defective, or incorrect qua beliefs if they are false.

However, another possible way to interpret the aim of belief is to view it instead as a *personal* aim i.e. as an aim that agents necessarily have or want to achieve in forming and revising beliefs.⁶² So on the personal interpretation, to say that belief constitutively aims at truth is to say, roughly, that we necessarily aim at believing the truth and not what is false about P whenever we form or revise beliefs about whether P. If this personal version of the aim of belief thesis is true, then it provides a straightforward way to explain away the No Desire intuition and establish the Epistemic Instrumental Thesis (EIT). Conforming to epistemic norms is plausibly our best way to achieve the aim of believing what is true and avoid error. But if that aim is one that we necessarily have or desire whenever we form or revise beliefs, then conforming to epistemic norms does necessarily promote our desires.

The problem with this strategy is that the personal version of the aim of belief thesis is untenable. First, more often than not, the formation and revision of our beliefs is something that happens to us automatically, unconsciously, and sub-personally. In Spoiler alert, Max's hearing his colleague's testimony forces him to

⁶² This section draws in part from section 4 of Côté-Bouchard (Forthcoming) and from section 4 of Côté-Bouchard (2015).

automatically believe that the character dies despite his desire to avoid learning the truth about the show's ending. Similarly, if I turn around and see a football rapidly coming towards my face, I will immediately form the belief that a football is coming at me. However, the formation of this belief will happen automatically at the sub-personal level before I can even realize it. In many cases of perceptual beliefs like this one, my believing is best seen as something that happens to me whether I like it or not, much like e.g. my digestion or my blood circulation.

Those sub-personal belief-forming processes do seem to be somehow directed or aimed towards producing true beliefs and responding to the evidence. Yet that does not entail or support the claim that we thereby have that truth aim. To the contrary, how can *I* plausibly be said to have wanted the truth in forming that perceptual belief if it happened at the sub-personal level before I could even be aware of it? Perhaps there is a metaphorical sense in which our sub-personal belief-forming systems 'want' to produce true beliefs and avoid producing false ones. But even if that is true, it does not follow we want it too. We do not necessarily want what our sub-personal mechanisms aim at. In comparison, all human agents have bodies that continually replace their skin cells, but that does not mean that all humans continually want to replace their skin cells. So despite their apparent truth-directedness, it is hard to see why these automatic and sub-personal belief-forming processes would necessarily involve my desiring to believe what is true.

Second, these sub-personal and unconscious belief-forming processes also include epistemically irrational ones like wishful thinking, delusion, and self-deception. These are problematic for the present defence of EIT not only because they are unconscious, but also because they are characterized precisely by a lack of concern for the truth. Suppose you are ill and despite your conclusively evidence to the contrary, you become convinced that you will certainly recover from your illness. You form that belief not as the result of deliberating about whether you will recover, but instead unconsciously as the result of your fear of dying and your inability to cope with the thought of not recovering. It is hard to see how you can be said to have wanted to believe the truth about your health in forming that belief. What you wanted was rather that the world be a certain way and your strong desire then caused you to believe that things really are that way. More generally, to engage in things like wishful thinking is precisely to (unconsciously) disregard the evidence and believe what we want to be true. Therefore, it is hard to see how engaging in these sorts of

sub-personal and epistemically irrational belief-forming processes could ever – let alone necessarily – involve a desire to believe the truth and avoid error.

Therefore, the thesis that belief constitutively aims at truth won't be of any help to the Epistemic Instrumental Thesis (EIT).

3.5 *The idealization strategy*

As I mentioned in 3.1, the Instrumental Principle (IP) is most notably associated with desire-based or internalist accounts of normativity and normative reasons. However, many internalists put constraints on what kind of desires can ground normative reasons. In particular, some argue that to determine what there is normative reason for S to do, we should look not at S's *actual* desires, but rather at the desires S would have under some idealized circumstances. We should look, for instance, at what S would want if she were coherent, fully informed about the situation, not prone to reasoning mistakes, and the like.⁶³ This suggests the following alternative strategy for explaining away the No Desire intuition. Although the agents in the cases from 3.2 do not *actually* want anything that would be promoted by conforming to epistemic norms, perhaps their relevantly idealized selves would. That is, perhaps conforming to epistemic necessarily promotes the satisfaction of desires that anyone's would have under relevantly idealized circumstances. However, this idealization strategy is not promising for two main reasons.

First, people's idealized selves would also be indifferent about sufficiently trivial questions. In Delaware's beverage, it is not clear why Nancy's idealized self would not also be completely indifferent about which beverage is the state beverage of Delaware. Moreover, people's idealized selves would also be in a better position to know the potential consequences (or lack thereof) of knowing the truth about a certain question. So in situations where the matter at hand is as inconsequential as the ones in Delaware's beverage, Dream, and Even Stars, it is plausible that the agents' idealized selves would also lack any desire that would be promoted by conforming to epistemic norms given that their knowledge of the trivial or inconsequential character of the matter at hand. It seems more plausible that at least many people's idealized selves would instead focus on knowing important truths and avoiding consequential errors.

⁶³ See for instance Williams (1979), Smith (1994), and Street (2008).

Second, no matter how exactly we construe agents' idealized circumstances, the idealization strategy cannot explain away cases like the modified versions of Delaware's beverage in which Nancy will die the instant after forming her doxastic attitude. Similarly, in the case where Nancy will be immediately killed by an evil demon unless she disbelieves Brett, looking at her idealized desires makes no difference. Since conforming to epistemic norms would get her killed immediately, doing so would not promote the satisfaction of her idealized desires either. Therefore, appealing to agents' idealized desires won't help explaining away the No Desire intuition either.

3.6 Kornblith's rule-based strategy⁶⁴

Finally, Hilary Kornblith (1993) has suggested a different kind of strategy to explain away the No Desire intuition. It differs from the ones I examined so far in one fundamental respect. So far I have only focused on what we might call *case-based* strategies or arguments. Case-based arguments attempt to show that in every particular case where epistemic norms intuitively apply to S, conforming to epistemic norms in that particular instance would promote the satisfaction of S's desires.

However, some like Kornblith have suggested an alternative kind of strategy which focuses less on what to believe in particular cases and more on what general *policies* or *rules* of belief formation to adopt.⁶⁵ The point of such *rule-based* strategies is that it would be sufficient to explain away the No Desire intuition and defend EIT if we could show that adopting the policy of *systematically conforming to epistemic norms* best promotes the satisfaction of one's desires in the long run.⁶⁶ The advantage of rule-based strategies is that they concede that there can be *particular* cases in which conforming to epistemic norms does not promote the satisfaction of any of one's desires, but adds that even if we admit that, it remains true that systematically conforming to epistemic norms, as a matter of general rule or policy, promotes the satisfaction of one's desires.

Kornblith's key idea is that the policy of systematically conforming to epistemic norms is the one that best promotes the satisfaction of *all* of our desires over time. The starting point of his argument is that all of us, trivially, want to achieve

⁶⁴ This section draws in part from section 5 of Côté-Bouchard (2015).

⁶⁵ See also Leite (2007).

⁶⁶ I borrow the labels 'case-based' and 'rule-based' from Lockard (2013).

our ends. But since achieving our desired ends requires successful practical reasoning and since successful practical reasoning requires true beliefs, adopting a belief-forming policy that maximizes true beliefs and minimizes false beliefs will best promote the achievement of our ends and thus the satisfaction of our desires over time. As he puts it:

It seems that someone who cares about acting in a way which furthers the things he cares about, and that includes all of us, has pragmatic reasons to favor a cognitive system which is effective in generating truths, whether he otherwise cares about truth or not. We should thus adopt a method of cognitive evaluation which endorses truth-conducive processes. (Kornblith 1993, 371-372)

Since systematically conforming to epistemic norms is arguably the best policy available to beings like us for maximizing true belief and minimizing error, we can conclude that there is necessarily a normative reason for us to do conform to such norms.

There are several problems with this strategy. First, it is not clear how Kornblith's story is supposed to explain away the No Desire intuition and establish EIT. For one thing, it is not clear why a policy of *systematically* conforming to epistemic norms would promote desire-satisfaction any more than a policy of always conforming to them *minus a few rare exceptions*. Take cases of trivial true beliefs for example. As we have already seen, some true or epistemically justified beliefs are too trivial and inconsequential to be of any use for future practical reasoning. Thus, sticking to the policy and conforming to epistemic norms in such cases of trivial truths would not promote our desires, even in the long run. So the policy of always conforming to epistemic norms *minus those very few triviality exceptions* promotes our desires just as much – if not more – than *systematically* following them. If that is so, then there remains no clear sense in which sticking to epistemic norms even in those very rare triviality cases promotes our desires. Since the policy of systematic conformity to epistemic norms does no better than the policy minus exceptions with respect to desire-satisfaction over time, it is unclear why there is any normative reason to conform to epistemic norms in problematic cases like triviality cases.

Second, even if we admit that Kornblith's strategy can plausibly explain away the cases from 3.2, we can imagine additional cases that cannot be explained away using Kornblith's strategy. Consider the following.

Mistaken Maddie Maddie does not want to know the truth about anything. To the contrary, the only thing she really wants is to maximize her false beliefs and minimize true ones. It is the only thing she really cares about. She is aware that she probably constantly forms many true beliefs thereby frustrating her goal and that makes her miserable. She also hates that she cannot help but believe what she thinks is true. She just hopes that she is wrong. Maddie is aware of how difficult and dangerous it is to pursue that aim (although she hopes she is wrong about that too). But she is fine with that. She thinks it is worth it and she is willing to die for the cause if that is what it entails.

Kornblith's story does not seem apply to Maddie's situation since systematically conforming to epistemic norms would clearly not be optimal for her desire-satisfaction. To the contrary, it would systematically frustrate her desires given her fundamental and strongest goal of maximizing false beliefs and minimizing true ones.⁶⁷

Similarly, consider cases of agents forming beliefs the instant before they die. Recall, for instance, the second modified version of Delaware's beverage from section 3.3 in which Brett tells Nancy about Delaware's beverage right before they hit the ground after having jumped off a plane without a parachute. Since Nancy will certainly die the instant after, sticking to the policy of systematically conforming to epistemic norms at that particular time certainly won't promote the satisfaction of any of her desires, even in the long run. Yet she would still be epistemically justified in believing him and epistemically unjustified in disbelieving him or suspending judgment.

Finally, it is far from clear that the optimal belief-forming policy with respect to desire-satisfaction of beings like us would be to systematically conform to

⁶⁷ Cases of agents who want nothing else but to believe *everything* or, alternatively, to believe *nothing*, are also problematic for Kornblith's strategy. For the former kind of agent, the optimal belief-forming policy for promoting desire satisfaction would presumably allow contradictory beliefs and forbid suspension of judgment. Yet, such doxastic attitudes would obviously still be epistemically unjustified. For the latter, the optimal belief-forming policy would allow for at least many instances of suspension of judgment and disbelief, which are, intuitively, epistemically unjustified.

epistemic norms or to maximize true beliefs and minimize error. Instead, the optimal policy might very well include or allow for at least *some* non-truth-conducive or unreliable cognitive processes, and thus for epistemically unjustified beliefs.

To see this, note first that as I mentioned above, successful practical reasoning does not necessarily require true or epistemically justified beliefs. In fact, for any of our ends, we can imagine scenarios in which a false or epistemically unjustified belief would actually *improve* our chances of satisfying that desire. Suppose I want to get from London to Edinburgh today and my only way to do so is to take a train that exceptionally departs at three o'clock today. Of course, one way to reach my goal would be to check the train schedule and form the true belief that the train departs at three today. But an even more straightforward and cost-effective way would be to believe that the train departs at 3pm as the result of wishful thinking.

The fact that we can come up with examples of that kind for any of our ends casts doubt on the idea that systematically conforming to epistemic norms is the optimal policy for long term desire-satisfaction. Couldn't the optimal policy for beings like us include at least *some* unreliable belief-forming mechanisms such as wishful thinking, cognitive biases, delusions, self-deception, and so on? After all, there is ample evidence that everyone constantly forms and revises beliefs via such unreliable processes. Couldn't they be so widespread in part because they help us satisfy our desires?

Of course, the question of which belief-forming system would be the optimal one for the desire satisfaction of beings like us is ultimately an empirical question (in fact, one worry with Kornblith's proposal is that he gives an answer to this empirical question without backing it with empirical evidence). But as I already mentioned in chapter 2, the actual empirical evidence does not seem not support Kornblith's story. Phenomena like positive illusions rather suggest that the optimal belief-forming policy would *not* be to systematically conform to epistemic norm, but would instead include at least some biases. Recall that positive illusions are beliefs which are incorrectly and unjustifiably optimistic. They are overestimations, for instance, of one's capacities, of one's control over one's life, and of how good one's future will be. As I mentioned above, there is evidence not only that such unjustified beliefs and optimistic biases are extremely common, but also that they promote well-being and thus, plausibly, the satisfaction of our desires. But since positive illusions involve violating epistemic norms, it is highly doubtful that the optimal belief-forming policy for desire-

satisfaction would be devoid of any biases and violations of epistemic norms. So for all these reasons, rule-based strategies like Kornblith's cannot explain away the No Desire intuition either.

3.7 Summary

A second commonly cited mark of normativity is a necessary connection with desires. I took this second feature to be captured by the Instrumental Principle (IP) according to which there is a normative reason for one to ϕ if and only if ϕ -ing promotes the satisfaction of one's desires. Given IP, if Epistemic Normativity is true and there is necessarily a normative reason to conform to epistemic norms, then the Epistemic Instrumental Thesis (EIT) must also be true: conforming to epistemic norms must necessarily promote the satisfaction of one's desires.

In this chapter, I argued that EIT is false because it cannot plausibly explain away the No Desire intuition according to which there seem to be cases where conforming to epistemic norms does not promote any of one's desires. In particular, I have shown that it cannot be explained away by invoking (i) potential desire-satisfaction, (ii) the aim of belief thesis, (iii) our idealized desires, and (iv) the optimal belief-forming policy for desire-satisfaction. Since EIT is false, I conclude that epistemic facts and claims do not bear the commonly cited mark of normativity captured by IP.

Chapter 4: Autonomy

4.1 Introduction: epistemology and the autonomy of the normative

A third commonly cited mark of normativity, but not of norm-relativity, is a necessary *autonomy* from the non-normative. Unlike merely norm-implying claims, in other words, it is often seen as an essential feature of normative claims that there is an unbridgeable gap between them and non-normative truths. This idea is most often introduced via this famous passage from Hume's *Treatise*:

In every system of morality, which I have hitherto met with, I have always remarked, that the author proceeds for some time in the ordinary ways of reasoning, and establishes the being of a God, or makes observations concerning human affairs; when all of a sudden I am surprised to find, that instead of the usual copulations of propositions, *is*, and *is not*, I meet with no proposition that is not connected with an *ought*, or an *ought not*. This change is imperceptible; but is however, of the last consequence. For as this *ought*, or *ought not*, expresses some new relation or affirmation, 'tis necessary that it should be observed and explained; and at the same time that a reason should be given, for what seems altogether inconceivable, how this new relation can be a deduction from others, which are entirely different from it. (Hume 1978, 469).

We cannot, in other words, derive an 'ought' from an 'is'. Normative conclusions cannot follow from non-normative premises alone. I will thus take this third commonly cited mark of normativity to be captured by the following principle:

Autonomy Principle (AP): *no normative claim can be derived from non-normative claims alone.*⁶⁸

AP is most readily associated with *non-naturalist* accounts of normativity, according to which normative facts are not analysable or reducible to natural facts. For one thing, non-naturalists typically invoke AP to motivate their view and to reject normative naturalism.⁶⁹ If there is an unbridgeable 'is'-'ought' gap and we cannot infer normative

⁶⁸ For discussion, see e.g. Prior (1960), Karmo (1988), Pigden (1989), (2010), Scanlon (2014), Van Roojen (2015), and Maguire (2015).

⁶⁹ See for instance Moore (1903), Johnston (1989), Scanlon (1998), (2014), Dancy (2006), Fitzpatrick (2008), (2010), Enoch (2011a), and Parfit (2011).

conclusions from purely natural non-normative premises, then this must mean that normative facts are of a fundamentally different kind from natural facts. But if they are so different, then it seems that any reduction or analysis of normative facts in terms of natural facts will inevitably entail changing the subject and *losing* the distinctly normative character of these facts. For another, even if the autonomy of the normative turned out not to entail or support non-naturalism, the *non*-autonomy of the normative would plausibly render non-naturalism unmotivated and even implausible. If it turned out that we *can* infer normative conclusions from natural facts alone, then it would be unclear why such normative truths could not be natural truths. Hence, AP is plausibly a central tenet of normative non-naturalism.

If AP is true and if Epistemic Normativity is also true, then the epistemic must also be autonomous from the non-normative. More precisely, if there is necessarily an unbridgeable gap between normative claims and non-normative claims and if epistemic claims are normative claims, then there must also be an unbridgeable gap between epistemic claims and non-normative claims. That is, the following thesis must also be true:

Epistemic Autonomy Thesis (EAT): *no epistemic claim can be derived from non-normative claims alone.*

The route from AP to EAT can be summarized as follows:

1. Normative claims cannot be derived from non-normative claims alone. (AP)
2. Epistemic claims are normative claims (*Epistemic Normativity*)
3. Therefore, epistemic claims cannot be derived from non-normative claims alone. (EAT)

In this chapter, I argue that the Epistemic Autonomy Thesis (EAT) is false. At least some epistemic claims can be derived from non-normative premises alone, which means that epistemic claims are *not* necessarily autonomous from non-normative claims.

My case against EAT has two parts. In section 4.2, I argue that although first person normative deliberation cannot be settled by non-normative facts alone (given AP), such facts *can* settle epistemic deliberation from a first person perspective. This is what I call the argument from epistemic deliberation against EAT. In section 4.3, I

argue that questions of the form ‘P, but is the case that E?’ where P is a non-normative claim and where E is an epistemic claim (*descriptive-epistemic* questions for short) are not necessarily *open* questions. At least some such questions are *closed* or trivial questions. But since descriptive-epistemic questions would be necessarily open the epistemic were autonomous from the non-normative, I conclude that EAT is false. Therefore, the commonly cited mark of normativity that is captured by AP is missing in epistemic facts and claims.

4.2 *The argument from deliberation*

4.2.1 AP and deliberation

A first way in which the autonomy of the normative manifests itself is in the phenomenon of first person deliberation. It seems that from the first person point of view, non-normative considerations alone cannot settle for you the normative questions you ask yourself in deliberation such as ‘should I ϕ ?’ or ‘what should I do?’. This idea was recently used by David Enoch to argue against the idea that normative truths can be identified or reduced to natural truths. He writes:

Because only normative truths can answer the normative questions I ask myself in deliberation, nothing less than a normative truth suffices for deliberation. And because the kind of normative facts that are indispensable for deliberation are just so different from naturalist, not-obviously-normative facts and truths, the chances of a naturalist reduction seem rather grim. [...] The gap between the normative and the natural, considered from the point of view of a deliberating agent, seems unbridgeable. (Enoch 2011, 80)

Van Roojen summarizes the thought as follows:

[W]hen one is deliberating and asking oneself what one should do, no purely naturalistic answer will provide an answer. [...] [A] naturalistic answer would change the subject from something normative to something else. Pointing out that one option for action satisfies all the interests involved won’t be an answer unless one already thinks one should satisfy all the interests involved. (Van Roojen 2015, 257-258)⁷⁰

⁷⁰ See also Rosati (1995).

If this is right, then a consequence of the Autonomy Principle (AP) is that it is an essential feature of normative deliberative questions that they cannot be settled by normative facts alone from the first person point of view. First person normative deliberation, in other words, is necessarily autonomous from the non-normative. AP would be false if normative deliberation were not autonomous from the non-normative since it would entail that at least some normative claims – i.e. answers to normative deliberative questions such as ‘I should ϕ ’ or ‘I should not ϕ ’ – can be derived from non-normative premises alone.

As I will argue however, the kind of deliberation that is relevant to epistemology and epistemic claims – i.e. the kind of deliberation that concludes with a first person epistemic claim of the form ‘I epistemically should believe that P’ – *can* be settled by non-normative facts alone. That is, *epistemic* deliberative questions of the form ‘epistemically speaking, should I believe that P?’ *can* be settled by non-normative claims alone from the first person point of view. Therefore, at least some epistemic claims – namely first person epistemic claims of the form ‘I epistemically should believe that P’ – can be derived from non-normative claims alone. This first argument against the Epistemic Autonomy Thesis (EAT) can be summarized as follows:

The argument from epistemic deliberation

1. If epistemic claims cannot be derived from non-normative claims alone (if EAT is true), then *first person epistemic deliberation* cannot be settled by non-normative facts alone.
2. First-person epistemic deliberation *can* be settled by non-normative claims alone.
3. Therefore, EAT is false; epistemic claims can be derived from non-normative claims alone.

4.2.2 Epistemic deliberation

To deliberate is to ask oneself deliberative questions like ‘what should I do?’ or ‘should I ϕ ?’. However, a well-known complication arises when ϕ -ing is a propositional attitude and not an action. In cases of such attitudinal deliberation, the question ‘should I ϕ ?’ admits both an *object-given* and a *state-given* interpretation. Suppose I ask

myself whether I should admire Wayne Rooney. On the one hand, he has impressive abilities and has had an outstanding career with many notable achievements. On the other hand, a group of Manchester City fans threaten to beat me up if I end up admiring Rooney. Both sets of facts seem to answer my deliberative question, but they do so differently. While the former set answers it by settling the question whether Rooney is *admirable* (irrespective of the consequences of admiring him), the latter answers it by settling the question whether admiring him would be *good* or *desirable* (regardless of whether Rooney is admirable). Hence, the question ‘should I admire Wayne Rooney?’ can be read either as ‘is Wayne Rooney admirable?’ or ‘would it be desirable to admire Wayne Rooney?’.

Another attitude we can deliberate about is, of course, belief. Since belief is an attitude and not an action, the same state-given/object-given ambiguity arises with *doxastic* deliberative questions like ‘should I believe that P?’ But although this question can be read as the (state-given) question ‘would it be *desirable* to believe that P?’, it is more commonly or naturally read in its *object-given* guise, namely ‘is P *true*?’ Suppose I ask myself whether I should believe that my wife is cheating on me. Even though I can conceivably answer this question by determining whether it would be desirable or beneficial to have that belief, I can also answer it by settling for myself the (object-given) question whether she really is cheating on me.

One important thing to note, however, is that the former, state-given version of doxastic deliberation is not strictly speaking relevant for epistemology and epistemic facts. That is, answers to state-given versions of doxastic deliberative questions – i.e. whether believing that P would be desirable, regardless of the truth of P – are not epistemic claims. Only *object-given* doxastic deliberation – the kind of doxastic deliberation that can be answered by settling the question ‘is P is true?’ – results in an epistemic claim or conclusion. This is because, as I have explained in the chapter I, epistemic facts and claims are not claims about the desirability of beliefs, but rather about what we should or should not believe from a truth-related point of view. That is, epistemic norms specify what we should or should not believe precisely from the point of view of things like truth, knowledge, and evidence, and not practical benefits of beliefs. Consequently, the distinctly *epistemic* kind of doxastic deliberation – i.e. the only kind of doxastic deliberation that we need to be concerned with as far as epistemic facts and claims go – is *object-given* doxastic deliberation. So to ask oneself object-given doxastic deliberative questions is to ask oneself the *epistemic* version of

the doxastic deliberative question, i.e. it is to ask oneself ‘what should I believe, *epistemically speaking*’. Hence, I will use ‘epistemic deliberation’ and ‘object-given doxastic deliberation’ interchangeably in what follows.

4.2.3 The non-autonomy of epistemic deliberation

Crucially however, such epistemic deliberation is not autonomous from the non-normative. It can be settled by non-normative facts alone from the first person perspective. As we have seen, epistemic or object-given doxastic deliberative questions of the form ‘should I believe that P?’ can be settled by answering the question ‘is P true?’ Crucially however, non-normative truths alone *can* settle for one the question ‘is P true?’ Once again, suppose I ask myself ‘should I believe that my wife is having an affair?’ On the epistemic or object-given interpretation, this question is settled by answering the question ‘is my wife really having an affair?’. Crucially however, non-normative claims alone can suffice to settle that question for me. Suppose I know the following:

(W) Several times in the past few weeks, my wife has come home late smelling of cheap cologne. She has been very distant and distracted lately. Some of my friends saw her with another man at a bar. I found incriminating romantic messages in her phone. I just saw her kissing another man at a restaurant. When I confronted her and told her what I saw her, she admitted that she was having an affair.

The facts in (W) would undoubtedly suffice to settle for me the question ‘is my wife really having an affair?’ But (W) only includes non-normative facts.

Similarly, suppose I ask myself whether I should believe that Nick is a bachelor and then think:

(N) Nick is an adult, Nick’s gender is male, Nick is not married, and bachelors just are unmarried adult males.

Although it only includes non-normative facts, (N) would obviously suffice to settle for me the question ‘is Nick a bachelor?’, which would in turn suffice to settle for me the epistemic or object-given version of the question ‘should I believe that Nick is a

bachelor?', i.e. the question 'should I believe that Nick is a bachelor, epistemically speaking?'

Finally, suppose I wonder whether I should believe that I am in Hong Kong right now and then think:

(L) I am in London right now.

Although (L) is a non-normative fact, it suffices to settle for me the question 'am I actually in Hong Kong right now?' The proposition that I am in Hong Kong right now is obviously false. This, in turn, suffices to settle for me the question whether I epistemically should believe that I am in Hong Kong.

So as these examples show, non-normative facts alone can suffice to settle epistemic doxastic deliberation from the first person perspective. All I need in order to settle the question whether I should *epistemically* believe that P is to settle the question whether P is true and non-normative facts alone *can* settle this kind of question for me. I therefore conclude that at least some epistemic claims can be derived from non-normative premises alone, namely first person claims of the form 'I epistemically should/should not believe that P'. Hence, the Epistemic Autonomy Thesis (EAT) is false.

4.3 The closed question argument

4.3.1 Open questions and the Autonomy Principle

A second phenomenon that is intimately connected with the Autonomy Principle (AP) is the one behind G.E. Moore's famous open question argument (OQA) in favour of the unanalysability of 'good' and against naturalistic accounts of value. According to Moore (1903), evaluative or normative claims X cannot be analysed in terms of non-normative properties N since for any such putative analysis, questions of the form 'X is N, but is it good?' (descriptive-normative questions for short) will always be *open* or *substantive* rather than *closed* or *trivial* questions. To use a common example, even if we accept a simple form of act-utilitarianism, the question 'ϕ-ing maximizes happiness, but is it morally right?' still seems open or substantive. Roughly put, the idea is that it is a question that someone who understands moral concepts could still reasonably ask. In comparison, if someone sincerely asked 'S is an unmarried adult

male, but is S a bachelor?', we would likely conclude that she does not understand or master the concept of bachelor in the first place.

Most contemporary metaethicists argue that the OQA is unsuccessful as an objection against naturalistic accounts of normativity. Nevertheless, as Darwall, Gibbard, and Railton write: "[h]owever readily we now reject as antiquated his views in semantics and epistemology, it seems impossible to deny that Moore was on to something." (Darwall, Gibbard, and Railton 1997, 3) In particular, many take Moore's observation about the openness of descriptive-normative questions to reveal something deep and distinctive about normative claims. In particular, one thing that seems to be behind the openness of descriptive-normative questions is precisely the Autonomy Principle (AP). If there were no unbridgeable gap between normative claims and non-normative claims, then such questions would not be open in the first place.

I will remain neutral about whether the openness of these questions supports or entails AP. What does seem clear, however, is that if AP is true, then descriptive-normative questions must be open. According to AP, normative claims cannot be derived or inferred from non-normative claims alone. But if some questions of the form 'X is N, but is it good?' were closed or trivial, then this would mean that some normative claims can be derived from non-normative premises alone. This link between open questions and the Autonomy Principle can be summarized as follows:

Open Question Principle (OQP): *If a claim of the form 'X is Y' is autonomous from the non-normative and if 'N' is a non-normative property, then questions of the form 'X is N, but is it Y?' must be open.*

Autonomy from the non-normative, in other words, leads to open descriptive-normative questions. Given the Open question principle (OQP), if the Epistemic Autonomy Thesis (EAT) is true – i.e. if epistemic claims are autonomous from non-normative claims – then the following must also be true:

Open Epistemic Questions (OEQ): *any descriptive-epistemic questions – i.e. questions of the form 'P, but is the case that E?' where P is a non-normative claim and E is an epistemic claim – must be open.*

In this section, I argue that OEQ is false. At least some descriptive-epistemic questions are not open or substantive, but rather closed or trivial.⁷¹ Therefore, at least some epistemic claims can be derived from non-normative claims alone and EAT is false.

4.3.2 Descriptive-epistemic *correctness* questions

A first kind of descriptive-epistemic questions which can be closed are ones involving epistemic *correctness* claims. Here are four examples:

- (a) P is true, but is S's belief that P correct, epistemically speaking?
- (b) P is true, but is S correct, epistemically speaking, in believing that P?
- (c) P is false, but is S's belief that P incorrect/mistaken, epistemically speaking?
- (d) P is false, but is S incorrect/mistaken, epistemically speaking, in believing that P?

The answer to (a)-(d) is, trivially, 'yes'. Once we find out that a proposition P is true, it is not a substantial or open question whether the belief that P is correct and whether anyone who believes that P is correct, at least from an epistemic standpoint. We can immediately infer or conclude that it is. Similarly, once we find out that P is false, we can immediately conclude that the belief that P is epistemically incorrect or mistaken, and that anyone who believes that P is epistemically incorrect or mistaken. Therefore, at least some epistemic correctness claims can be derived from non-normative premises alone.

One might object at this point that attributions of correctness are not really distinctly epistemic claims because 'correct' and 'incorrect' just means 'true' and 'false'. If this is right, then the triviality or closedness of (a)-(d) does not count against EAT since it is simply due to the fact that 'correct' and 'incorrect' are just other words for 'true' and 'false'. However, this is mistaken. 'Correct' and 'incorrect' are not just other words for 'true' and 'false' since things can be correct or incorrect without being true or false. As Rosen writes:

'[C]orrect' has application where 'true' does not. You can play a sonata incorrectly. You can dance the Mambo incorrectly. A recitation of 'Gunga

⁷¹ See also Heathwood (2009) for a similar view.

Din', an attempt to spell 'chiaroscuro', a move in chess, the placement of a fork: a vast range of performances can be correct or incorrect, most of which are quite incapable of being true. (Rosen 2001, 619)

Moreover (and relatedly), we can evaluate beliefs as correct or incorrect from non-epistemic points of view. Although a true belief is obviously correct from an epistemic point of view, it might not be the *prudentially* correct thing to believe if, for example it would cost you your life. Similarly, the truth might not be the *morally* correct thing to believe if, for instance, it would humiliate others or cause people to suffer.

A second potential reply to the closedness of some descriptive-epistemic correctness questions is that even if it is not synonymous with truth, correctness is not the kind of epistemic claims that are relevant to Epistemic Normativity. Instead, normative epistemology is concerned with things like epistemic justification, epistemic rationality, and epistemic 'oughts', which do not lead to closed descriptive-epistemic questions. There are two problems with this response.

First, this reply is of no help to proponents of Epistemic Normativity who also accept a popular *veritist* picture of the epistemic domain according to which, very roughly, correctness or truth is the fundamental or primary epistemic norm, goal, or value, from which all other epistemic norms derive. On such a picture, all epistemic norms besides the fundamental norm of correctness or truth are secondary ones that only have *instrumental* normative force or value. More precisely, according to this kind of picture, nonfundamental epistemic norms or values – e.g. evidential norms, epistemic justification, epistemic rationality, etc. – only have normative force or are only valuable insofar as they are means to the primary goal of believing what is correct and not what is incorrect. For example, Laurence Bonjour writes:

What then is the differentia which distinguishes epistemic justification, the species of justification appropriate to knowledge, from these other species of justification? The answer is to be found, I submit, by reflecting on the implicit rationale of the concept of knowledge itself. What after all is the point of such a concept, and what role is epistemic justification supposed to play in it? Why should we, as cognitive beings, *care* whether our beliefs are epistemically justified? Why is such justification something to be sought and valued? Once the question is posed in this way, the following answer seems obviously correct, at least in first approximation. What makes us cognitive beings at all is our capacity for belief, and the goal of our distinctively cognitive endeavors is *truth*!. We want our beliefs to correctly and accurately depict the world.

[...] The basic role of justification is that of a *means* to truth, a more directly attainable mediating link between our subjective starting point and our objective goal. [...] If epistemic justification were not conducive to truth in this way, if finding epistemically justified beliefs did not substantially increase the likelihood of finding true ones, then epistemic justification would be irrelevant to our main cognitive goal and of dubious worth. It is only if we have some reason for thinking that epistemic justification constitutes a path to truth that we as cognitive beings have any motive for preferring epistemically justified beliefs to epistemically unjustified ones. Epistemic justification is therefore in the final analysis only an instrumental value, not an intrinsic one. (BonJour 1985, 7-8)

Similarly, Lynch writes:

[W]e take it to be correct to believe what is based on evidence *because* beliefs based on evidence are likely to be *true*, and thus the value of truth [...] is more basic than the value of believing what is based on evidence, (Lynch 2009, 229)⁷²

So according to such veritist accounts of epistemic norms and their normativity, there is only a normative reason to e.g. follow the evidence because it is our best means to meeting the fundamental norm of correctness and to attain the fundamental goal of believing what is true and not what is false.

Why does this matter for our purpose? If AP is true – i.e. if no normative claim can be derived from non-normative claims alone – and if the popular veritist picture sketched above is right, then showing that correctness claims are not autonomous from non-normative claims is sufficient to cause trouble for EAT and Epistemic Normativity. First, if AP is true, then the fact that epistemic correctness claims lead to closed questions means that they are not normative and that there is not necessarily a normative reason to conform to the truth norm of correctness. Second, according to the veritist picture, nonfundamental epistemic norms are means to the end of achieving the primary epistemic norm of correctness. But there is a normative reason to take the means to an end only if there is a normative reason to achieve that end. If there is no normative reason to achieve an end, then there is no normative reason to take the means to that end. Therefore, there is a normative

⁷² See also e.g. Goldman (1999a), David (2001), (2005a), Wedgwood (2002), Olsson (2007), Sylvan (2012), Steglich-Petersen (2013), and Ahlstrom-Vij (2013a). For criticism, see e.g. Stich (1990), DePaul (2001), Kvanvig (2003), and Littlejohn (Forthcoming b).

reason to conform to nonfundamental epistemic norms only if there is a reason to conform to the fundamental truth norm of correctness. But since there is not necessarily a normative reason to conform to the fundamental truth norm of correctness, it follows that there is not necessarily a normative reason to conform to nonfundamental epistemic norms either. So if the common veritist picture of the epistemic domain is the right one, then showing that correctness claims are merely norm-implying and not normative is sufficient to cast doubt on the normative character of all epistemic claims.

In any case, even if you do not accept the veritist conception or epistemic norms and normativity, epistemic correctness claims are not the only ones that can be involved in closed questions.

4.3.3 Other closed descriptive-epistemic questions

First, as the discussion of first person epistemic deliberation from section 4.2 already indicated, a second kind of descriptive-epistemic questions that can be closed are those involving first person claims about what we should or should not believe, epistemically speaking. Here are three examples:

- (e) P is false, but should I believe that P, epistemically speaking?
- (f) P is true, but should I believe that P, epistemically speaking?
- (g) P is true, if P then Q, but should I believe that Q, epistemically speaking?

Questions (e)-(g) are closed. If I am convinced that P is false (true), then the first-personal question whether I epistemically should believe that P is not an open or substantial question for me. I will immediately conclude that I epistemically should not (should) believe that P.

There can also be closed descriptive-epistemic questions involving third-personal epistemic claims. For instance, descriptive-epistemic questions involving attributions of epistemic justification and epistemic rationality are not necessarily open. Here are three examples:

- (h) S believes that P only because S wants P to be true and P is false, but is S's belief that P epistemically rational/justified?

- (i) S consciously believes that P and $\neg P$, but is S's belief that P and $\neg P$ epistemically rational/justified?
- (j) S believes that $\neg Q$ only because she believes 'if P then Q' and ' $\neg P$ ', but is S's belief that $\neg Q$ epistemically rational/justified?

The answer to (h)-(j) is, trivially, 'no'. Once we find out that S's belief has the characteristics described in the left-hand side of (h)-(j), it is not a substantial or open question whether that belief is epistemically justified or rational. We can immediately infer that it is epistemically unjustified and irrational. Therefore, we have another example of epistemic conclusions – i.e. claims of the form 'S's belief that P is epistemically unjustified/irrational' – that can be derived from non-normative premises alone.

One possible reply to examples (h)-(j) is that although notions like epistemic justification and rationality are normative, they might not be *purely* normative. That is, perhaps they are not *thin* normative concepts like 'should' or 'ought', but rather thick concepts that necessarily involve a descriptive or non-normative elements and perhaps the closedness of (h)-(j) is only due to the non-normative elements of epistemic justification and rationality. But the problem with this reply is that (h)-(j) remain closed even if we reformulate them in terms of epistemic 'shoulds' or 'oughts':

- (k) S believes that P only because S wants P to be true and P is false, but should S believe that P epistemically speaking?
- (l) S consciously believes both that P and that $\neg P$, but should S believe P and $\neg P$ epistemically speaking?
- (m) S believes that $\neg Q$ only because she believes 'if P then Q' and ' $\neg P$ ', but should S believe $\neg Q$ epistemically speaking?

Trivially, the answer to (k)-(m) is 'no'. It is not a substantial or open question whether S ought to believe these propositions from an epistemic point of view.

At this point, some might be tempted to concede that the above kinds of claims are not autonomous from the non-normative, but add that the kind of epistemic claims that matter for Epistemic Normativity – the kinds of epistemic claims that are the real purely normative epistemic claims – are those involving *reasons for belief*. After all, many see the concept of a normative reason as the fundamental

normative concept or unit, so perhaps reasons for belief are the fundamental *epistemic* normative concept or unit.

However, even if we concede that questions involving reasons for belief *simpliciter* are necessarily open, epistemic claims have to do *not* with reasons for belief *simpliciter*, but with *epistemic* reasons for belief. This is because attributions of non-epistemic reasons for belief – i.e. prudential, moral, or aesthetic reasons for belief – are not epistemic claims in the sense we are interested in. The fact that there are reasons to believe that P from a moral, prudential, or aesthetic point of view is irrelevant to the epistemic status of the belief that P. So attributions of reasons for belief do not count as epistemic claims unless they are attributions of *epistemic* reasons.

One likely reply to this point is that according to several epistemologists and normativity theorists, reasons for beliefs *just are* epistemic reasons because there are no such things as practical or non-epistemic reasons for belief.⁷³ According to them, there can only be practical reasons to *bring about* a belief. While I am sceptical of that claim, there is no space to properly evaluate it here.⁷⁴ However, even if we concede that normative reasons for belief can only be epistemic reasons, it does not follow that attributions of epistemic reasons are necessarily normative claims. It does not follow, in other words, that epistemic reasons necessarily constitute normative reasons or that they necessarily have normative force. This is because just like ‘should’, ‘ought’, ‘good’, and the like, ‘reasons’ can be used both to denote genuine normativity or mere norm-relativity. Suppose I wonder whether there is any reason for me to do as e.g. the law or etiquette require. One could very well respond that, trivially, there are *legal reasons* and *etiquette reasons* to do so. However, such legal claims and etiquette claims only entail norm-relativity. That is, they only entail that there is a reason to do so *relative to* or *according to* these norms, which in turn does not entail that there is any genuinely normative reason to do so. They do not entail, in other words, that these legal and etiquette reasons have any genuine normative force. Hence, it might be that just like legal reasons, epistemic reasons do not

⁷³ This position is often referred to as *evidentialism* in the ‘ethics of belief’ debate initiated by Clifford (1999). See e.g. Kelly (2002), Shah and Velleman (2005), Shah (2006), Skorupski (2010), and Way (Forthcoming a).

⁷⁴ See Foley (1991), James (1995), Owens (2000), (2003), Steglich-Petersen (2006), (2008), Reisner (2008), (2009b), (2013), (2014), (Forthcoming), and Littlejohn (2012a) for criticism.

necessarily have normative force.⁷⁵ And this is a possibility even if we admit that there cannot be normative practical reasons for belief and that normative reasons for beliefs can only be epistemic. In other words, even if all normative reasons for beliefs are epistemic, it might still be the case that not all epistemic reasons are genuinely normative.

In any case, the main point I want to make about epistemic reasons is that even epistemic reasons claims can be involved in closed descriptive-epistemic questions. That is, descriptive-epistemic reasons claims are not necessarily open questions. Take, for instance, questions involving claims about evidence or likelihood on the descriptive side:

- (n) It is highly likely that P, but is there any (defeasible) epistemic reason to believe that P?
- (o) P is evidence that Q, but is P a (defeasible) epistemic reason to believe that Q?
- (p) P is true and the probability of Q given P is much higher than the probability of Q given not-P, but is P a (defeasible) epistemic reason to believe Q?

The answer to (n)-(p) is, trivially, 'yes'. Once we know that it is highly likely that P, it is not an open question whether there is a defeasible epistemic reason to believe that P. We immediately conclude that there is. Similarly, once we find out that P is evidence that Q, it is not a substantive question whether P is a defeasible epistemic reason to believe that Q. We immediately infer that it is.

Of course, my opponents will be quick to respond that many philosophers view evidence and probability as epistemic notions and that this is why questions like (n)-(p) are trivial.⁷⁶ Perhaps the concept of evidence just is the concept of a kind of epistemic reason or perhaps the notion of likelihood or chance is best understood in terms of e.g. epistemically rational credence. There is no space to properly evaluate these conceptions of evidence and likelihood here, but I will make three points in response.

⁷⁵ I return to this issue in chapter 7.

⁷⁶ See e.g. Greco (2015).

First, it won't help my opponents if notions of evidence and likelihood are best analysed in terms of epistemic rationality or epistemic justification since, as we have seen with (h)-(j), descriptive-epistemic questions involving these kinds of epistemic claims are not necessarily open. Second, even if evidential and likelihood claims are epistemic claims, they too can be involved in closed descriptive-epistemic questions. For example, the following seem closed:

- (q) P is necessarily false, but is it likely that P?
- (r) P is necessarily true, but is it likely that $\neg P$?
- (s) P is false and no one believes that P, but is P evidence that Q?
- (t) P is true and P entails Q, but is it likely that $\neg Q$?

Finally, there can be closed descriptive-epistemic reasons questions that do not involve evidence and probability on the descriptive side. For one thing, the first personal questions like (e)-(g) can be put in terms of epistemic reasons.

- (u) P is true, but is there any (defeasible) epistemic reason for me to believe that P?
- (v) P is false, but is there any (defeasible) epistemic reason for me to believe that $\neg P$?
- (w) P is true, if P then Q, but is there any (defeasible) epistemic reason for me to believe that Q?

The answer to (u)-(w) is, trivially, 'yes'. Second, there can also be closed third personal descriptive-epistemic reasons questions that do not involve evidence or likelihood:

- (x) P is a necessary truth and P entails Q, but is there any (defeasible) epistemic reason to believe that Q?
- (y) P is necessarily false and Q is only true if P is true, but is there any (defeasible) epistemic reason to believe $\neg Q$?

Once again, the answer to (x) and (y) is, trivially, 'yes'. Once we find out that the descriptive side of (x) and (y) is true, it is not a substantial or open question whether the epistemic reasons claim on the right-hand side is true.

4.3.4 Taking Stock

As I explained at the beginning of this section, one plausible upshot of the Epistemic Autonomy Thesis (EAT) is that any descriptive-epistemic question – i.e. questions of the form 'P, but is the case that E?' where P is a non-normative claim and E is an epistemic claim – must be open. This is what I called the Open Epistemic Questions thesis (OEQ). In this section, I argued that OEQ is false. There can clearly be closed or trivial descriptive-epistemic questions. In particular, I gave examples of closed descriptive-epistemic questions involving correctness claims, first-personal 'epistemic should' claims, attributions of epistemic justification and rationality, third-personal 'epistemic should' claims, 'epistemic reasons' claims, evidential claims, and likelihood claims. Since OEQ is implied by EAT and OEQ is false, I conclude that EAT is false as well.

4.4 Summary

A third commonly cited mark of normativity is that unlike merely norm-implying facts and claims, normative facts and claims are autonomous from the non-normative. That is, no normative conclusion can be derived from non-normative premises alone. This is what I called the Autonomy Principle (AP). If AP is true and if, as Epistemic Normativity maintains, epistemic claims are normative claims, then it must also be the case that no epistemic claim can be derived from non-normative claims alone. This latter thesis is what I called the Epistemic Autonomy Thesis (EAT).

I argued that EAT is false. At least some epistemic claims *can* be derived from non-normative premises alone. This is because (i) non-normative premises alone *can* settle one's epistemic deliberation from the first person perspective, and (ii) that at least some descriptive-epistemic questions are closed or trivial and not open or substantive. Hence, I conclude that epistemic claims do not bear the mark of normativity that is captured by AP.

Chapter 5: Motivation

5.1 Introduction: epistemology and motivation

A fourth commonly cited mark of normativity, but not of mere norm-relativity is a necessary connection with *motivation*. In particular, many see it as an essential feature of normative *judgment* or *thought* – in particular, of normative judgments of the form ‘I should ϕ ’ – that unlike merely norm-implying judgments, they necessarily motivate us to some extent to behave in accordance with them. If you are not motivated at all to ϕ , the thought goes, then you do not sincerely think that you ought to ϕ .⁷⁷ I will take this fourth commonly cited feature of normativity to be captured by the following principle:

Motivation Principle (MP): *making a normative judgment of the form ‘I should ϕ ’ necessarily motivates one to behave in accordance with that judgment.*⁷⁸

Given MP, if epistemic judgments are normative judgments – as proponents of Epistemic Normativity maintain – then epistemic judgments of the relevant form must also be necessarily motivating. That is, if MP and Epistemic Normativity are both true, then the following thesis must also be true:

Epistemic Motivation Thesis (EMT): *making an epistemic judgment of the form ‘I epistemically should believe that P’ necessarily motivates one to behave in accordance with that judgment.*

The route from MP to EMT can be summarized as follows:

1. Making a normative judgment of the form ‘I should ϕ ’ necessarily motivates one to behave in accordance with that judgment. (MP)
2. Epistemic judgments are normative judgments (*Epistemic Normativity*)

⁷⁷ See, for instance, Korsgaard (1986), Smith (1994), Jackson and Pettit (1995), and Wedgwood (2007). This idea is also held by most proponents normative *expressivism*. I discuss expressivism below.

⁷⁸ Let me make two clarifications. First, MP does not mean that we necessarily do or try to do what we think we ought to do, but rather that these judgments necessarily give us *some* motivation to ϕ , i.e. only *pro tanto* motivation. It allows that we might not end up ϕ -ing or trying to ϕ . Second, normative judgments of the form ‘I ought to ϕ ’ are not the only kinds of judgments that are necessarily motivating according to MP. However, I focus on first-personal ought judgments because if any normative judgments are motivating, then it is them.

3. Therefore, making an epistemic judgment of the form ‘I epistemically should believe that P’ necessarily motivates one to behave in accordance with that judgment. (EMT)

In this chapter, I argue that the Epistemic Motivation Thesis (EMT) is false. Epistemic judgments – including those of the form ‘I epistemically should believe that P’ – do not necessarily motivate us to behave in accordance with them. My case for this thesis has two parts.

In section 5.2, I reject *expressivist* accounts of epistemic judgments. Rejecting these accounts is necessary because as I will explain, epistemic expressivism would plausibly provide a straightforward route to EMT. Then in section 5.3, I move away from expressivism and ask whether EMT might be true independently of expressivism. In response, I examine various possible interpretations of EMT and argue that none of them is plausible. That is, there is no plausible sense in which epistemic judgments are necessarily – or even typically – motivating. Thus, I conclude that EMT is false and that this fourth commonly cited mark of normativity is also missing in epistemic judgments.

5.2 *Against Epistemic Expressivism*

5.2.1 From normative expressivism to epistemic expressivism, to EMT

According to normative expressivism, normative judgments are not ordinary representational beliefs, but rather states that necessarily involve non-cognitive motivating states. Normative claims, in other words, necessarily express desire-like, non-representational states.⁷⁹ For expressivists, the mark of the normative is precisely the expression of the relevant kind of non-representational state. Expressivism is opposed to cognitivist or descriptivist accounts according to which normative judgments or thoughts are ordinary representational beliefs. Since they typically accept Epistemic Normativity, most expressivists think that an adequate version of their theory must extend to epistemic thought and discourse. Consequently, some of them have recently formulated versions of *epistemic* expressivism according to

⁷⁹ Stevenson (1937), Ayer (1949), Hare (1952), Urmson (1968), Blackburn (1984), (1993), (1998), Gibbard (1990), (2003), Horgan and Timmons (2006), Schroeder (2008), Sinclair (2009), Ridge (2014), and Toppinen (2015).

which epistemic judgments are not mere ordinary representational beliefs, but rather, at least in part, non-cognitive motivating states.⁸⁰

Adopting Epistemic Expressivism is a first possible route to establish the Epistemic Motivation Thesis (EMT). This is because the Motivation Principle (MP) is typically entailed by normative expressivism. For most expressivists, the non-cognitive states expressed by normative claims are precisely the kind of states that we are in when we are motivated to ϕ . If this is right, then MP is obviously true if expressivism is true because the latter entails that normative judgments *just are* – at least in part – motivating states.⁸¹ But if epistemic judgments are normative judgments, then they are also – at least in part – motivating states. Hence, epistemic expressivism plausibly entails that epistemic judgments are necessarily motivating (EMT).

However, in the rest of this section, I argue that epistemic expressivism is implausible. Epistemic judgments are not best seen as states that necessarily involve non-cognitive motivating states. My case for this thesis has two parts. First, I reject the three most well-developed versions of epistemic expressivism in the literature, namely Allan Gibbard's *plan-reliance expressivism*, Michael Ridge's *procedure-expressivism*, and Christos Kyriacou's *habits-expressivism*.⁸² Second, I offer a more general objection against all versions of epistemic expressivism.

5.2.2 Gibbard's plan-reliance expressivism

In his 2003 book *Thinking how to live*, Allan Gibbard argues that since normative claims are essentially claims about what to do, they express non-cognitive 'contingency

⁸⁰ Gibbard (1990), (2003), Blackburn (1996), Chrisman (2007), (2012), Ridge (2007), Kappel (2010), (2011), Carter and Chrisman (2011), Kyriacou (2012), Kappel and Moeller (2013), and Alhstrom-Vij (2013b).

⁸¹ In fact, MP is standardly seen as a central reason to adopt normative expressivism and reject cognitivism. That is, many have maintained that the motivating character of normative judgments (MP) means that such judgments cannot be ordinary beliefs because, the argument goes, ordinary beliefs alone cannot motivate.

⁸² Simon Blackburn (1996) has also suggested a version of epistemic expressivism. However, I do not consider it below because Ridge's proposal is essentially an elaboration of Blackburn's. As Ridge writes:

Blackburn's account is fine as far as it goes, but it does not go far enough. As an expressivist about knowledge attributions, he is committed to taking attributions of knowledge to express distinctive attitudes. Precisely what kinds of attitudes are involved here, though? What is it, on Blackburn's account, for a speaker to take a certain improvement in his acquaintance with the facts to *justify* abandoning a given belief? More to the point, what is it for a speaker to judge that no such improvements would justify abandoning a given belief? Without answers to these questions, it is hard to know just how plausible an expressivist account of knowledge attributions really is. [...] To this significant extent, Blackburn's account is incomplete. Because I think Blackburn's basic approach was on the right track, I shall try to remedy this incompleteness. (Ridge 2007, 93-94)

plans', i.e. one's plans about what to do for expected or hypothetical scenarios. In the same book, Gibbard extends this account to epistemic claims. Focusing on knowledge attributions, he argues that such claims express *plans of reliance*. More precisely, he argues that claims of the form 'S knows that P' necessarily express a plan *to rely on S's judgment with respect to P*. So you cannot count as sincerely attributing knowledge to someone unless you intend to rely on S's judgment regarding the question at hand. The problems with Gibbard's plan-reliance expressivism have been articulated most clearly by other epistemic expressivists.

First, according to Ridge (2007, 91-92), plan-reliance expressivism faces a dilemma. According to Gibbard, my claim that S knows that P expresses plans to rely on S's judgment that P. But this raises a question: can these plans take into account the peculiarities of my own circumstances or must they abstract from such circumstances? The problem with the former option is that it fails to do justice to the fact my knowledge attribution is a judgment about S and not about me. It makes my judgment that S knows that P too much of a judgment about myself. Moreover, Ridge thinks that going this first route makes plan-reliance expressivism vulnerable to clear counter-examples:

[S]uppose I judge that S is reliable with respect to p, and not subject to defeaters. I do not, however, plan to rely on S with regard to p because I judge it to be morally distasteful even to take a view about whether p is true. Moreover, I do not even form contingency plans for what to believe should I come to abandon this moral evaluation (of taking a view with respect to p), because I believe that making plans for what to do should I become so corrupt is itself a sign of moral corruption. Intuitively, this should be consistent with my judging still judging that S knows whether p is the case, for I might know that S is in as good a position as anyone to form a reliable judgement about p. However, on the view I am now considering, I cannot count as attributing this knowledge to S; in so far as I do not plan to rely on S (or indeed, anyone, under any circumstances) specifically with regard to p, I cannot on Gibbard's account count as attributing knowledge to S. (Ridge 2007, 91)

This leaves the second approach, according to which the plans expressed in knowledge attributions are plans which abstract from the attributor's own peculiar circumstances.

But this option is no more tenable since ordinary speakers do not seem to form plans of this sort in everyday instances of knowledge attributions. After all, it is

unclear why we would even need to make plan that completely abstract from our own situation. As Ridge points out “[s]uch planning behind a sort of ‘veils of ignorance’ might serve a useful theoretical role in certain contexts, but it does not seem to be part of our ordinary everyday planning about on whom we shall rely.” (Ridge 2007, 92) But since ordinary speakers do constantly attribute knowledge to others, these everyday epistemic judgments are plausibly not constituted by such abstract contingency plans.

Second, according to Kyriacou (2012), the notion of planning is alien to ordinary epistemic discourse. It just does not seem like we express antecedently chosen plans when we make epistemic claims. Instead, as Kyriacou writes, “[w]e make judgments rather unreflectively and ignore the possibility of antecedently chosen plans constraining our judgments.” (Kyriacou 2012, 218) Plan-reliance expressivism thus seems to over-intellectualize epistemic judgments. It implies that for every epistemic claim, the speaker has antecedently decided on a plan for what to do or believe in future and hypothetical situations. But it does seem that most of us most of the time go about making epistemic judgments without having gone through such a process.

Third, as Kyriacou (2012, 214-216) points out, Gibbard’s plan-reliance expressivism is at best incomplete because it is ill suited for non-factive epistemic claims. It is possible to judge that S is e.g. epistemically rational or reasonable in believing that P while also knowing that P is false. But if I know that S is wrong about whether P, then I won’t rely on S’s judgment. Suppose Myriam believes that the cinema will be open on Wednesday because she went there last Wednesday and checked the schedule on the website before going. Unbeknownst to her, however, I bought the cinema the day after and changed the opening hours just to prove Myriam wrong. Although she is now mistaken, she is still epistemically rational in believing that the cinema will be open on Wednesday. Yet since I know that it is false, I would not plan to rely on her judgment regarding the cinema’s opening hours. Hence, Gibbard’s account is at best incomplete as an account of epistemic judgments.⁸³

⁸³ Note however, that the same worry does not extend to ‘epistemic ought’ judgments in general. For one thing, if I sincerely judge that S knows that P, then I am also convinced that P is true, which means that I will also judge that I should believe that P epistemically speaking. For another, if I think S knows that P, then I also think that S believes as she epistemically ought. Hence, if Gibbard is right about knowledge attributions, then we can reasonably expect his account to cover at least some ‘epistemic ought’ judgments as well.

Finally, we routinely attribute knowledge to children and animals. Yet we obviously do not plan to rely on their judgment about the matter at hand whenever we make such attributions. My dog knows that someone is at the door right now, but I have no plan to rely on his judgment on that matter. So for all these reasons, Gibbard's version of epistemic expressivism is not plausible.

5.2.3 Ridge's procedure-expressivism

Another leading proponent of epistemic expressivism is Michael Ridge. According to him, claims of the form 'S knows that P' express a complex state which includes the following two states:

- (1) A non-cognitive state of "[e]pistemic endorsement of certain procedures for deciding what to believe." (Ridge 2007, 103)
- (2) "The [ordinary representational] belief that S's judgement that p is causally regulated by either (a) those procedures (anaphoric reference back to those procedures the speaker endorses in [1]) or (b) procedures which are close enough to those procedures, so far as p goes or (c) more fully informed successors to those procedures." (Ridge 2007, 103)

Several components of this account call for explanation.

First, what is it to endorse procedures of belief-formation and revision? According to Ridge, such a state of endorsement "is just a commitment to follow those procedures [...]." (Ridge 2007, 102) Moreover, in order to be committed to follow these procedures, one must genuinely accept these procedures and not - as Gibbard (1990) puts it - merely be 'in their grip'. Simplifying somewhat, to genuinely accept something like a procedure - rather than merely being in its grip - is, according to Gibbard, in part to be disposed to avow it in normative discussion (Gibbard 1990, 74). Therefore, whenever I attribute knowledge to S, I express my commitment to follow a certain belief-forming procedure, where I have this commitment because I am disposed to avow such a procedure in normative discussion and not merely because I am 'in its grip'.⁸⁴

Second, what is it for an endorsement to be epistemic? According to Ridge, it is for it to be "taken on at least in part because the speaker takes it that following such procedures will reliably [...] track the descriptive truth. Other motives for

⁸⁴ I return to this distinction below.

taking on the commitment must also be recognizably epistemic.” (Ridge 2007, 101) Therefore, whenever I attribute knowledge to S, I express a commitment to follow a certain epistemic procedure and I have this commitment because of epistemic or truth-related motives. That is, I have taken on this commitment to follow this procedure because, for instance, I take it to be reliable.

Finally, we need condition (2) because when we attribute knowledge to S, we do not necessarily presume that S uses exactly the procedures which we are committed to follow. For one thing, we often attribute knowledge to S without having any idea what S’s procedure is. Ridge illustrates this point as follows:

I might have no idea what procedures the physicist uses to determine that the material is radioactive. Whatever procedure he uses, though, I judge to be more reliable at tracking the truth than my own methods (if any I have!) would be for such purposes (putting to one side the possibility of my just asking the physicist, which is itself parasitic on his own methods anyway). Because I view procedures which track the truth (with respect to p) as better than my own (with respect to p, anyway), I can on the proposed account still attribute knowledge to the scientist without the absurd suggestion that he uses my own primitive methods of belief formation to answer his scientific question. (Ridge 2007, 104)

Since, I have no idea what the physicist’s procedure is for determining whether X is radioactive, I cannot be committed to follow it. But I can still attribute knowledge to him. According to Ridge, this is because I still believe that his procedure is more reliable than the one I am committed to follow.

This last example raises a first problem for Ridge’s proposal. Suppose I have extremely limited knowledge of science but nevertheless judge that the physicist knows that X is radioactive. Given Ridge’s account, this means I believe that the physicist’s belief that X is radioactive is causally regulated by belief-forming procedures that are more reliable than the ones I am committed to follow. The problem, however, is that I am not committed to follow any procedure to decide whether X is radioactive. This is because I have absolutely no idea how to settle that question. But since I have no procedure to endorse, how can I, on Ridge’s account, really attribute knowledge to the physicist?

A second problem is that it is not clear that knowledge attributions always refer to belief-forming *procedures*. First, take attributions of propositional knowledge

that do not mention any specific proposition. Suppose I say, for instance, 'Stephen Hawking knows a lot more things than the average Briton'. I am obviously attributing knowledge to Hawking in making such a claim. But it is hard to see which belief-forming procedure I could be endorsing in making such a claim. Second, according to many epistemologists, a lot of our knowledge is knowledge 'by acquaintance', i.e. direct, unmediated, and noninferential knowledge. But knowledge by acquaintance seems to be - perhaps by definition - the kind of knowledge that is not causally regulated by procedures for deciding what to believe. Suppose I know by acquaintance that I am now typing on a keyboard. Plausibly, I do not believe this as a result of following any procedure for deciding what to believe. Rather, I just know it as a result of my direct and unmediated acquaintance with my typing on the keyboard.

Finally, Ridge's proposal is also vulnerable to Kyriacou's alienation worry. Recall what Ridge says when he explains why the relevant state of endorsement must be epistemic:

In order for a commitment to a set of such procedures to count as epistemic, that commitment must be taken on at least in part because the speaker takes it that following such procedures will reliably (how reliably may vary from speaker to speaker) track the descriptive truth. Other motives for taking on the commitment must also be recognizably epistemic. (Ridge 2007, 101)

This suggests that we can only genuinely judge that S knows that P if we have actually taken on a commitment to follow the relevant procedures for deciding what to believe. But what is it to take on a commitment to follow a certain procedure? A natural answer is that it is to deliberately choose or decide to follow the procedure in the future. This interpretation is supported by Ridge's claim that the commitment must have been taken on for a particular kind of motive, namely epistemic motives. This suggests that I must have thought about the procedure, judged it e.g. reliable, and as a result decided to follow it in future belief-forming endeavours.

The problem is that we rarely, if ever, go through such a process of deliberately choosing or deciding what belief-forming procedures to follow in the future. Very often, we just form our beliefs via whatever procedure we are compelled to follow without having ever thought about that procedure. We could even imagine someone spending her entire life without ever even thinking about procedures for

deciding what to believe. Even such an agent could make at least some genuine epistemic judgments.

5.2.4 Kyriacou's habits-expressivism

A third version of epistemic expressivism was recently proposed by Christos Kyriacou (2012). Unlike the previous two, Kyriacou's focuses on attributions of epistemic justification. According to him, claims of the form 'S's belief that P is epistemically justified' express endorsement not of procedures, but of *habits* of belief-formation. More precisely, they are in part non-cognitive states of "[e]ndorsement of the habits of belief-fixation in virtue of which S believes that p, where these habits of belief-fixation employ and are constrained by certain epistemic norms." (Kyriacou 2012, 229). Kyriacou illustrates his account via the following example:

Suppose I say 'Mary justifiably believes that her next pottery class will take place this weekend'. According to the expressivist semantic picture we have painted so far, what I, the attributor, mean with this sentence is that Mary believes that her next pottery class will take place this weekend and endorse Mary's habits of belief-fixation in virtue of which she has formed this belief. Maybe the habit of belief-fixation responsible for the formation of Mary's belief is her habit to trust what reliable sources of information (like the pottery class website) say. (Kyriacou 2012, 222)

There are at least three problems with this proposal.

First, some epistemically justified beliefs are not formed in virtue of any *habit*.

Consider the following example:

Sarah and the scientists. For all her life, Sarah has had the extremely firm habit of not trusting scientists. Whenever a scientist would affirm something, she would automatically take this as counting against the veracity of that statement. She did however break this habit once in her life. One day, she met with a zoologist who told her she was making her dog sick by giving it chocolate. This time, and only this time, she immediately believed what the scientist told her. But right after that, she went back to systematically distrusting scientists and kept that habit for the rest of her life.

Clearly, Sarah was epistemically justified in believing what the zoologist told her. But since we know that her trusting a scientist was a one-time thing and thus clearly not a habit, it is not clear how we could be endorsing any habit in attributing epistemic justification to Sarah. Here is a second example:

One-time belief: Jo has a rare syndrome that prevents her from having full beliefs. She can imagine, suppose, assume, guess, suspect, or partially believe that P, but she cannot fully believe anything. One day, she is invited at a lab to try a newly developed pill that is supposed to treat her syndrome. She goes to the lab, takes the pill, and after a few minutes, forms her first full belief, namely the belief that she is at a lab. However, the instant after the pill starts working, there is a massive explosion that knocks Jo unconscious and kills all the scientists who were developing the pill. It also destroys all the pills they were about to give her as well as the hard drives with all the research used to develop the pill. When she wakes up, the effects of the medicine have worn off and she is once again unable to fully believe anything. Unfortunately for Jo, the pharmaceutical company then abandons the project and so she remains unable to have full beliefs for the rest of her life.

We know that Jo does not have any habit of belief-formation since the belief she formed in the lab was the only one she ever had in her life. Yet her single belief was still epistemically justified.

Second, we can very well judge that a belief is epistemically justified despite knowing that it was formed in virtue of a belief-fixation habit that we do not endorse, but rather *disavow*. Consider the following case:

Godless René. René has the extremely firm habit of believing only what he could not possibly doubt. So while he believes that he exists, René suspends his judgment about everything else. After all, he thinks, everything else can be doubted. For all he knows, he could be systematically deceived by an evil demon.

No one would endorse the ridiculously stringent belief-forming habit in virtue of which René formed his belief that he exists. But at the same time, no one would hesitate to claim that René is epistemically justified in believing that he exists.

So for all these reasons, endorsement of belief-forming habits is not necessary for sincere attributions of epistemic justification.

5.2.5 The problem of epistemic psychopaths

I have argued that the versions of epistemic expressivism offered by Gibbard, Ridge, and Kyriacou are all implausible. Of course, this does not mean that epistemic expressivism in general is hopeless. For one thing, I have not mentioned two additional versions of epistemic expressivism that were recently proposed. One is Kristoffer Ahlstrom-Vij's (2013b) proposal that epistemic claims express commitments not to procedures or habits, but to *goals of inquiry*. Another one builds instead on Gibbard's earlier *norm-expressivist* account of normative claims, which views normative judgments as endorsement or acceptance of *norms*. Accordingly, some have argued that epistemic claims might express endorsement of epistemic norms. Matthew Chrisman, for instance, suggested that "knowledge claims could be understood as expressing our acceptance of particular epistemic norms, which when applied to a particular person's belief entitle or don't entitle the belief." (Chrisman 2007, 241)⁸⁵ Although there is no space here to evaluate these proposals individually, I will reject them by raising a problem that affects all versions of epistemic expressivism. In particular, I will argue that the kind of mental state that is posited by epistemic expressivist accounts is not essential or necessary for genuine epistemic judgments.

Although they disagree on what non-cognitive states epistemic claims express exactly, epistemic expressivists all seem to share the idea that such states must involve a kind of *endorsement, commitment, or acceptance*. Moreover, they take this endorsement or commitment to be endorsement of something that is distinctly epistemic, truth-related, or conducive to some sort of epistemic success. As we have seen, this thing can be a belief-forming procedure, a belief-forming habit, someone's judgment, a goal, a norm, and so on. But what matters to epistemic expressivism is that we endorse this thing in light of or because of its conduciveness to truth or some

⁸⁵ See also Kappel (2010).

related epistemic success. This is part of what marks epistemic judgments as distinctly epistemic and not, say, prudential or moral. As a shorthand to denote this kind of state, I will use the label *endorsement of epistemic norms*.

As Gibbard and Ridge explain, someone only counts as endorsing a norm if one genuinely *accepts* that norm and is not merely 'in its grip'. Gibbard (1990) uses the example of Milgram's famous experiment to illustrate the difference between accepting a norm and being in its grip. For Gibbard, typical subjects in Milgram's experiment were faced with a conflict between two norms. One is a norm against intentional harm, which forbids administering potentially lethal electric shocks to people. The other is a norm of cooperativeness and of doing one's job, which requires following the experimenter's instructions. In this experiment, conforming to the latter norm meant administering potentially lethal shocks to other subjects. As is well known, typical subjects in Milgram's experiment ended up following the cooperativeness norm at the expense of the non-harm norm. Yet it seems wrong to say that they genuinely accepted or endorsed that norm which they ended up following. As Gibbard explains:

We, as judges, accept a norm against infliction of harm, and accept that this norm, in the situation of Milgram's subjects, overrides norms of cooperativeness and doing one's job - norms that we also accept. The subjects, on the other hand, do not genuinely accept that in their situation, norms of cooperativeness and the like override all other norms. Rather we might say, they are in the grip of these norms. [...] It is accepting norms that matters here, not being in their grip. What, after all, does [the] subject [...] think it rational to do? [H]e thinks that it makes no sense to cooperate, but finds himself cooperating nevertheless. [...] [W]hat he actually does, in this case, is a matter of the norms that have him in their grip - norms of politeness and cooperativeness. What he think it rational to do, on the other hand, is what is required by norms against inflicting pain and danger - and these are the norms he accepts as having most weight in his situation. (Gibbard 1990, 60-61)

As I explained above, Gibbard suggests that to accept a norm is in part to be disposed to avow that norm in normative discussion. Plausibly, typical subjects in Milgram's experiment were not disposed to avow what they ended up doing. They would certainly agree with external observers that following the cooperativeness norm at the expense of the non-harm norm is indefensible. But since they did end up following

the cooperativeness to the point of inflicting harm, they are best characterized as having been *in the grip* of that norm rather than having endorsed and accepted it in their situation.

What does this mean for epistemic expressivism? As I explained above, epistemic expressivists share the basic idea that epistemic judgments necessarily involve what I called a non-cognitive state of endorsement of epistemic norms. Given Gibbard's explanations, this means that epistemic judgments require genuinely accepting epistemic norms and not merely being in their grip. However, consider the following case:

Pat, the epistemic psychopath. Pat despises epistemology and epistemologists. Epistemologists always go on about how to get things like truth and knowledge. But for Pat, promoting knowledge is a profoundly evil thing to do. Pat is not a sceptic. She is convinced that we know a ton of things. But this, according to her, is a tragedy. For Pat, the world is such a horrible and terrifying place that all the suffering in the world is ultimately due to our accurately representing reality. For Pat, all of humanity's problems and suffering would be eliminated if only we could all manage to live in our own fantasy, safely disconnected from reality.

Pat is what I will call an *epistemic psychopath*. Epistemic psychopaths do not endorse epistemic norms. To the contrary, they readily and emphatically *disavow* them given their complete aversion to true belief and knowledge.

Of course, most of the time, epistemic psychopaths would be unable to avoid conforming to the epistemic norms that they disavow. For example, it is very difficult – if not impossible – for beings like us to avoid believing what we take to be conclusively shown by the evidence. More generally, we have very limited control over how we form our beliefs and, as I mentioned in earlier chapters, beliefs might very well aim at truth in some sense. However, these considerations are perfectly compatible with the possibility or conceivability of epistemic psychopaths. They simply mean that such agents would constantly be *in the grip* of epistemic norms without genuinely accepting them, much like most of us would be in the grip of the cooperativeness norm in Milgram's experiment without accepting it.

According to epistemic expressivism however, epistemic psychopaths like Pat cannot make epistemic judgments and sincerely attribute things like knowledge, epistemic justification, and epistemic rationality. But imagine you met Pat and she told you: “Sceptics are mistaken: people do have knowledge and epistemically justified beliefs. What epistemologists don’t understand is that this is a tragedy. This is what makes everyone miserable: they keep representing this horrible reality correctly! Take me for instance. I know that I exist. I cannot help it. Heck, I even know that I know that I exist. I keep trying to get rid of all this knowledge, but it won’t go away. Isn’t that awful?” Would you deny that she really is attributing epistemic justification, correctness, and knowledge to herself and others in uttering these words? Surely not. It is hard to see why we would deny that she makes genuine epistemic claims in uttering things like ‘I know that P and this is terrible’ or ‘Every person knows that she exists and that is a tragedy’. Knowing that she completely disavows epistemic norms and the aim of truth or knowledge does not change this.

So in sum, the argument is the following. According to epistemic expressivism, epistemic judgments necessarily involve what I called endorsement of epistemic norms. However, if epistemic psychopaths can make epistemic judgments, then epistemic judgments do not necessarily involve such endorsements. Since epistemic psychopaths can, on the face of it, make sincere epistemic judgments, I conclude that epistemic judgments do not necessarily involve such endorsements.

5.2.6 Taking stock

I have done two main things in this section. First, I examined and rejected three of the most well-developed versions of epistemic expressivism in the literature. Second, I raised what I called the problem of epistemic psychopaths for the fundamental claim shared by all versions of epistemic expressivism. I argued that endorsement of epistemic norms is not essential for epistemic judgments since agents who disavow epistemic norms can still, on the face of it, make such judgments. For all these reasons, I conclude that epistemic expressivism, which plausibly implies the Epistemic Motivation Thesis (EMT), is not tenable.⁸⁶

⁸⁶ There is an additional problem for epistemic expressivism, which derives from the conclusions of chapter 4. Many normative expressivists invoke the Autonomy Principle (AP) as a central motivation for their view. This is because, roughly, expressivism appears to offer an attractive option for those who (i) agree with nonnaturalists that naturalistic accounts of normative facts are implausible given AP, (ii) want to avoid the error theory, but (iii) reject the existence of nonnatural facts. Hence, it

5.3 Against the Epistemic Motivation Thesis (EMT)

Although epistemic expressivism plausibly entails the Epistemic Motivation Thesis (EMT), EMT does not necessarily entail expressivism. After all, some of the most prominent advocates of the Motivation Principle (MP) are cognitivists.⁸⁷ So even if epistemic expressivism is implausible, EMT might still be true. In the rest of this chapter, I set expressivism aside and turn to EMT directly to evaluate it independently of the dispute between cognitivists and expressivists. I will consider the two main possible ways to interpret EMT more precisely – i.e. as the claim that epistemic judgments necessarily motivate us to *believe* or to *act* in accordance with those judgments – and argue that none of them is plausible. Hence, I conclude that EMT is false. Even if we accept the MP, there is no plausible sense in which epistemic judgments – including those of the form ‘I epistemically ought to believe that P’ – are necessarily motivating.

5.3.1 Doxastic EMT

Since epistemic judgments are about what we epistemically should or should not believe, the most natural interpretation of EMT is a *doxastic* interpretation according to which making an epistemic judgment of the form ‘I epistemically should believe that P’ necessarily motivates one to *believe that P*. However, I will argue that doxastic EMT is untenable given what being motivated to ϕ plausibly involves.

What is it to be motivated to ϕ ? A common suggestion is that it is to *intend* or to be inclined to *try* to ϕ .⁸⁸ If this is right, then the Motivation Principle (MP) entails that if you make a normative judgment of the form ‘I should ϕ ’, then you intend or

promises to accommodate the unbridgeable gap between ‘is’ and ‘ought’ while positing neither mysterious nonnatural facts nor massive and systematic normative error on our part. But if AP does not extend to epistemic facts and claims (as I have argued in chapter 4), then this argument for expressivism does not get off the ground in the epistemic case. There is no need to accommodate any unbridgeable gap between the descriptive and the epistemic in the first place.

⁸⁷ For instance, Korsgaard (1986), Smith (1994), Jackson and Pettit (1995), and Wedgwood (2007).

⁸⁸ See for instance Mele (2003), Wedgwood (2007), and Cuneo (2007). Another common suggestion is that to be motivated to ϕ is to *want* to ϕ . See for instance Jackson and Pettit (1995). If this is right, then EMT entails that we necessarily *want* to behave in accordance with our epistemic judgments. I won’t consider this suggestion in the present chapter since I have already argued in chapter 3 that conforming to epistemic norms does not necessarily promote the satisfaction of our desires. In any case, intending is plausibly a better candidate. After all, we often seem to intend to do things that we do not want to do. And in many cases, this is precisely because we think that this is what we ought to do. If I received a call telling me I must go to the morgue to identify the corpse of a close friend, I would judge that I ought to do it and would form the intention to go to the morgue as a result. Yet I would not want to do it.

have the inclination to try to behave in accordance with that judgment. Similarly, then, the doxastic version of EMT entails that if you judge that you epistemically should believe that P, then you intend or are inclined to try to *believe that P*.

An initial problem with this version of EMT is that although we often make epistemic judgments, we very rarely try or intend to believe things. The formation and revision of our beliefs is typically something that automatically happens to us whether we like to not. This morning, I saw a fox in my garden and I immediately formed the belief that there was a fox in my garden as a result. Believing that there is a fox in my garden is not something I intended or tried to achieve, but rather an unintended and automatic consequence of my looking through the window. Most of our beliefs seem to come about in this way. So more often than not, the process of forming and revising our beliefs neither requires nor involves intending or trying to believe. Therefore, if there are cases where epistemic judgments do lead us to intend or try to believe that P, then they would have to be much rarer than epistemic judgments.

Second and relatedly, when we judge 'I epistemically should believe that P', it is typically because we take the case for the truth of P to be conclusive. Call these *conclusive epistemic judgments*. After all, as I explained in chapter 4, the epistemic doxastic 'should' is, at least typically, the *object-given* doxastic 'should', i.e. the sense of 'should believe' which is such that the question 'should I believe that P?' is settled, from the first person perspective, by an answer to the question 'is P true?' Consequently, it seems that an agent who sincerely makes what I called a conclusive epistemic judgment will *thereby* believe in accordance with that judgment. As Joseph Raz explains:

[O]ne who believes that there is a conclusive case for the truth of a proposition cannot but believe that proposition [...]. There is no gap, no extra step in reasoning, between believing that the case for the truth of the proposition is conclusive and believing the proposition. [...] The No Gap Principle states that one comes to believe that P upon realizing that there is conclusive evidence for it [...]. (Raz 2011, 38-39)

Similarly, Pamela Hieronymi writes:

Suppose you take certain considerations to bear on whether P, and, further, you take them to be sufficient to show that P [...] you yourself find the reasons convincing, you are convinced by them. Of course, if you take certain reasons to show that P, you therein believe P. [...] By finding such reasons convincing, you therein believe. (Hieronymi 2006, 51)

Sincerely judging that the case for the truth of P is conclusive plausibly implies or involves believing that P is true. If you do not yet believe that P is true, then it is hard to see how you can really think that the case for P is conclusive or sufficient.

But if this is right, then at least many cases of believing what you judge you epistemically should believe do not involve intending or trying to believe. At least in the case of conclusive epistemic judgments, sincerely making such epistemic judgments already means that you believe in accordance with that judgment. So once you make such an epistemic judgment, there is neither need nor room for intending or trying to believe what you think you epistemically should believe since you already believe it. No intending or trying needs to come in to bridge the gap between the (conclusive) epistemic judgment and the belief since, as Raz puts it, there is simply no gap to bridge. Put differently, the point is that one can only intend or try to ϕ if one has not yet ϕ -ed. But as soon as we make conclusive epistemic judgments, we believe that P. Therefore, if to be motivated to ϕ is to intend or be inclined to try to ϕ , then epistemic judgments do not typically motivate us to believe what we think we epistemically should believe, at least when it comes to conclusive epistemic judgments.

A third problem with doxastic EMT stems from the fact that trying and intending seem essentially tied to action. I cannot plausibly count as trying or intending to ϕ if I do not try or intend to perform some action that, I think, might lead to ϕ -ing. Of course, this is not to say that one can only try or intend to ϕ if ϕ -ing is *itself* an action. I can, for example, try to become healthier. However, it only makes sense to say that I am trying to become healthier if I am also trying to take actions that, I think, might make me healthier. Similarly, I can intend to have lower blood pressure by Christmas, but this does not make sense unless I also intend to do at least something that, I think might lower my blood pressure by Christmas. If it became clear to you that I do not intend to do anything to achieve these goals, then you would conclude that I am not really trying to become healthier or that I do not really intend to have

lower blood pressure by Christmas. So when ϕ -ing is not itself an action, trying or intending to ϕ requires trying or intending to perform actions that, one thinks, might bring it about that one ϕ -s.

But since beliefs are states and not actions, it follows that I cannot really count as trying or intending to *believe* that P unless I have the intention or the inclination to try to do something that, I think, might bring it about that I believe that P. Moreover, if being motivated to ϕ entails or amounts to intending or being inclined to try to ϕ , then I cannot really be motivated to believe that P unless I have the intention or the inclination to try to perform actions that, I think, might bring it about that I believe that P. This, in turn, means that according to doxastic EMT, if I judge that I epistemically should believe that P, then I have the intention or the inclination to try to perform actions that, I think, might bring it about that I believe that P.

The problem is that epistemic judgments do not entail having such intentions or inclinations. First, as I explained above, we often make *conclusive* epistemic judgments, which are such that we thereby believe in accordance with them. But as I have also pointed out, we can only intend to ϕ or try to bring it about that we ϕ if we have not yet ϕ -ed.

Second, even in cases where I do not yet believe what I think I epistemically should believe, it is implausible that I would necessarily intend to bring about that belief. Suppose I judge that I epistemically should believe that time travel is possible because, say, this is what best physicists claim. But suppose that for some reason, and despite making that epistemic judgment, I still do not believe in the possibility of time travel. Does that mean that necessarily, in such a case, I would have the intention or the inclination to try to change this and perform actions that might bring it about that I believe in time travel? Would I necessarily, for instance, intend or be inclined to try to reflect about the issue more carefully, read more about the topic, talk with physicists about it, and so on? Clearly not. It seems entirely possible that, in such a situation, I would remain completely unmoved by this issue and lack any intention to bring about the belief in question. I could very well fail to see any value in making the efforts required to change my belief about such an obscure and inconsequential matter. I could be perfectly content with having this particular tension in my beliefs and see no point in taking steps to eliminate it. There is, after all, a difference between judging 'I epistemically should believe that P' and judging 'I epistemically should bring

it about that I believe that P'. Even if I do not yet believe that P, I can very well make the former judgment without the latter.

One might reply by pointing out that if I do not intend or have the inclination to try to do things to change my belief in such a situation, then I am irrational. In particular, I am violating the enkratic requirement of rationality according to which, roughly, I am irrational if I do not intend to do what I think I ought to do.⁸⁹ So one could reply that this example is not relevant against doxastic EMT because EMT only applies to rational agents.⁹⁰ That is, one might claim that epistemic judgments (and normative judgments more generally) are necessarily motivating, but only for rational agents.

The problem is that I do not violate enkrasia in the time travel example. As I pointed out above, although I judge that I epistemically should believe in time travel, I do not judge that I should *bring it about* that I believe in time travel. Hence, I do not violate enkrasia by not intending to bring about that belief. Perhaps I still violate a distinctly *epistemic* enkratic requirement according to which, roughly, I am *epistemically* irrational if I do not *believe* what I think I epistemically should believe. Hence, one might reply that since EMT is a thesis specifically about epistemic judgments, the correct reply to my time travel case is that EMT only applies to epistemically rational agents. The idea would then be that epistemically rational agents necessarily intend or have the inclination to try to believe what they think they epistemically should believe.

However, this cannot work either because the two main components of this proposed revised version of EMT – i.e. (i) that I am epistemically rational and (ii) that I intend to believe what I think I epistemically should believe – cannot both be true at the same time. In particular, (ii) can only be true if (i) is false. On the one hand, if (i) is true (i.e. if I am epistemically rational), then I *do* believe what I think I epistemically should believe, in which case I won't intend or try to bring about that belief. On the other hand, if (ii) is true (i.e. if I intend or try to bring about the belief in question), then I do not yet believe what I think I epistemically should believe, in which case I am not epistemically rational.

⁸⁹ This is only a rough approximation, which overlooks the numerous complications and controversies that arise once we try to formulate the enkratic requirement more precisely. See e.g. Kolodny (2005), Brunero (2010), Broome (2013), and Way (Forthcoming b) for discussion.

⁹⁰ Many proponents of the Motivation Principle (MP) actually think it only holds for rational agents. See e.g. Korsgaard (1986), Smith (1994), and Wedgwood (2007).

At this point, proponents of EMT could concede that epistemic judgments do not necessarily entail having the intention or the inclination to try to believe that P, but add that we can be motivated to ϕ without intending or being inclined to try to ϕ . In particular, it might be said that epistemic judgments still necessarily *move* or *cause* epistemically rational agents to believe accordingly. After all, it does seem that whenever an epistemically rational agent judges 'I epistemically should believe that P', she also believes that P. Doesn't that mean that such judgments necessarily move such agents?

The problem, however, is that when agents believe what they think they epistemically should believe, it is typically not because they think that this is what they epistemically should believe. More often than not, in other words, it is not your judging that you epistemically should believe that P that, in turn, causes you or moves you to believe that P. Suppose I see a fox in my garden and judge that I epistemically should believe that there is a fox in my garden as a result. Since I am epistemically rational, I immediately form the belief that there is a fox in my garden. Plausibly however, what causes me or moves me to believe that there is a fox is not my thinking 'I epistemically should believe that P', but rather my seeing the fox. Similarly, suppose my biologist friend tells me that there was a fox in his garden and as a result I judge that I epistemically should believe that there was a fox in his garden. Since I am epistemically rational, I believe him and so I believe what I think I epistemically should believe. In such a situation, however, I would normally believe my friend not because I think that this is what I epistemically should believe, but rather because he told me.

Cases like these are familiar and they do not seem to be the exception. In fact, it is not even clear that there are cases where it is judging 'I epistemically should believe that P' that causes epistemically rational agents to believe that P. For one thing, as I pointed out above, if I am epistemically rational, then I already believe what I think I epistemically should believe. For another, it seems that if I am epistemically rational, then whatever causes me to judge 'I epistemically should believe that P' – e.g. seeing that P, being told that P, etc. – will also cause me to believe that P. It is hard to imagine a case, in other words, where what causes me to think 'I epistemically should believe that P' does not also cause me to believe that P, especially if I am epistemically rational. In all likelihood, if seeing that P or being told that P causes me to judge 'I epistemically should believe that P', it will be because I think that my seeing that P or my being told that P means that P is true. But if that is what I think, then I

am already convinced that P is true and so – insofar as I am epistemically rational – I will believe that P.

Therefore, even if we understand being motivated to believe that P simply as being moved or caused to believe that P, doxastic EMT is still not plausible because judging ‘I epistemically should believe that P’ is rarely – if ever – what moves us or causes us to believe that P. Doxastic EMT is thus implausible. Judging ‘I epistemically should believe that P’ does not necessarily – or even typically – motivate us to believe that P.

5.3.2 Practical EMT

The other obvious possible version of EMT is what I will call a *practical* version, according to which, even though judging ‘I epistemically should believe that P’ does not necessarily motivate you to believe that P (or to bring it about that one believes that P), it still necessarily motivates you to *do* something (other than bringing about the belief that P). The immediate problem with this idea however is that it is far from clear what acting in accordance with the judgment ‘I epistemically should believe that P’ could be besides trying to bring about the belief that P. Is there really something that agents necessarily intend or are inclined to try to do whenever they think they epistemically should believe something? One promising answer has recently been offered by Klemens Kappel and Emil Moeller (2013). According to them, epistemic judgments of the relevant form necessarily motivate us to *terminate inquiry with respect to P*.⁹¹ This suggestion is implausible for three reasons.

First, even if we admit that we are motivated to terminate inquiry whenever we make the relevant kind of epistemic judgment, it is far from clear our epistemic judgment itself necessarily has anything to do with that motivation. In fact, it is implausible that it is our making the relevant epistemic judgment that moves us or causes us to terminate inquiry. To see this, note first that the aim or goal of inquiring about whether P plausibly involves finding out or settling the question ‘is P is true?’ Consequently, if I take the question whether P to be settled, then I take the goal of inquiry about whether P to have been attained. As we have seen however, if you are epistemically rational and you judge that you epistemically should believe that P, then you thereby believe that P. But if you already believe that P, then in all likelihood, you

⁹¹ Kappel and Moeller only mention knowledge attributions, but their proposal plausibly extends to judgments of the form ‘I epistemically should believe that P’, or so I will assume.

also already take question 'is P true?' to be settled. This, in turn, means that at least in many cases where I judge 'I epistemically should believe that P', I also take the goal of inquiring about whether P to be attained. But if this is right, then in those cases, what causes me to stop inquiring is not my epistemic judgment, but rather my believing P. I stop trying to find out whether P not because I judge that I epistemically should believe that P, but rather because I already take P to be true. After all, solely taking P to be true – i.e. without also making an epistemic judgment – suffices to make me stop inquiring about whether P.

Second, although it does seem that we do not go on inquiring when we make the relevant kind of epistemic judgment, it is not clear this always mean that we are *motivated to stop inquiring* in those cases. If to be motivated to ϕ is to intend or be inclined to try to ϕ , then Kappel and Moeller's proposal would mean that we necessarily intend to terminate inquiry when we make the relevant kind of epistemic judgments. Now as Kappel and Moeller themselves point out, we are not always already inquiring when we make epistemic judgments. So what kind of intention are we supposed to have when we are already not inquiring? The natural answer is that we intend *not* to inquire further or to continue not inquiring in such cases. This, however, is only plausible if, in those cases, not inquiring either is or involves an act. As I pointed out above, intending to ϕ involves intending to act. But does such a 'not-inquiring' necessarily involves acting? There seems to be at least two senses in which I can 'not- ϕ '. Nicholas Rescher explains the distinction as follows:

When sitting at my desk writing I may refrain from scratching an itching mosquito bite—that is I “hold myself back” or “keep myself from” doing a certain action. This sort of keeping oneself from doing something that is at issue in refraining is importantly different from a second type of inaction which may be illustrated as follows. When sitting at my desk writing, there is an endless number of things I am not doing: reading the newspaper, chatting with a friend, driving a car, etc. But these nonactions are not doings of any sort. I am not somehow active in keeping myself from doing them. And thus, unlike refrainings, they are not actions at all. There is a critical difference between doing not-X, which is an action, and not-doing X, which need not be. (Rescher 1970, 248)

Applying this to inquiring, we can distinguish between not-inquiring as a mere 'nonaction' and not-inquiring as *actively refraining from inquiring*. Since only the latter

involves acting, it is only the latter that we can intend. Therefore, Kappel and Moeller must mean that whenever we make epistemic judgments, we intend to 'do' the latter, namely to actively refrain from inquiring.

The problem is that my not inquiring after making the relevant kind of epistemic judgment does not seem to typically involve actively refraining from inquiring. As I pointed out above, making such epistemic judgments at least often involves taking the question whether P to be settled. But because I take this question to be settled, I might very well simply not have any motivation or intention to try to find out whether P. Hence, I might very well simply continue *not* inquiring (in the sense of mere 'nonacting') without ever actively refraining from inquiring. Take the fox example I gave above. Because I saw the fox with my own eyes, I take the question whether there was a fox in my garden to be settled and so I am not trying to find out whether there was a fox in my garden. Does that mean that I now intend to actively refrain from inquiring about whether there was a fox in my garden? Do I thereby intend to actively refrain from e.g. asking neighbours if they saw it too, looking for fox footprints, and the like? It does not seem so. I could very well lack any such intention and spend the rest of my life merely not inquiring about that fox (in the sense of mere 'nonacting').

A potential reply is that epistemic judgments still necessarily entail being *disposed* to refrain from inquiring. As Kappel and Moeller argue, for example, if you make an epistemic judgment of the relevant kind, then you are such that if someone were to subsequently suggest that you should inquire, you would be disinclined to do so. *And perhaps being so disposed just is to intend to refrain from inquiring.* There are two problems with this response however.

First, even if we admit that we have such a disposition to refrain from inquiring whenever we make an epistemic judgment, what such a disposition entails is not that we *intend to refrain* from inquiring, but rather we are *disposed to intend* to refrain from inquiring. This is because there can be cases where you (i) make an epistemic judgment and take the question whether P to be settled, (ii) lack any intention to refrain from inquiring and instead simply go on not inquiring in the sense of mere nonacting, but only until (iii) you later find yourself in a position where you need to decide whether to inquire about whether P, in which case you then (and only then), form the intention to refrain from inquiring. So having a disposition of the kind Kappel and Moeller mention does not entail having the intention to refrain from inquiring.

Second, even if epistemic judgments do make us disposed to refrain from inquiring, their causing this kind of disposition cannot plausibly be sufficient for being *motivating*. This is because if it were, it would entail that an implausibly large class of judgments are necessarily motivating, including judgments that are obviously not normative. In particular, it would entail that simply judging that I have achieved some goal would be necessarily motivating as well. This is because just like epistemic judgments, these judgments also entail being disposed not to pursue the goal in question. For example, since I know that I have already received my Master's degree in Philosophy, I am now disposed not to try to get that degree anymore. Similarly, since I think that my mother has already received my birthday gift for her, I am disposed not to try to send her that gift. But if making us so disposed was all it took for a judgment to be motivating, then even judgments like these would be necessarily motivating. But this is implausible. If being necessarily motivating is a distinctive feature of normative judgments, then it must amount to more than a mundane feature that even clearly non-normative judgments like 'I received my Master's degree' have.

This proves too much especially by Kappel and Moeller's own lights. Following several philosophers before them, they argue that the Motivation Principle (MP) supports normative expressivism. This is because, they argue, ordinary representational beliefs alone cannot motivate. Only non-cognitive states can motivate. Kappel and Moeller's goal is to apply this kind of argument to epistemic judgments: since epistemic judgments are necessarily motivating and since ordinary beliefs alone cannot motivate, they conclude that we should be epistemic expressivists. This means that according to them, if a judgment is an ordinary belief, then it cannot be motivating. However, judgments like 'I was awarded my Master's degree' or 'my mother has received her gift from me' are clearly ordinary representational beliefs. But if this is correct, then they cannot say that such judgments are motivating because like most expressivists, they think that ordinary beliefs cannot be motivating.

For all those reasons, the practical version of EMT is not plausible either. Epistemic judgments do not necessarily motivate us to act in accordance with them.

5.3.3 Taking Stock

In this section, I examined and rejected the two main possible versions of the Epistemic Motivation Thesis (EMT), i.e. doxastic EMT and practical EMT. Given what

'being motivated to ϕ ' and acting in accordance with epistemic judgments might involve, there is no plausible sense in which epistemic judgments – even those the form 'I epistemically should believe that P' – necessarily motivate us to believe or to act in accordance with them.

5.4 Summary

A fourth commonly cited mark of normativity is a necessary connection between normative judgments and motivation. I took this fourth feature to be captured by the Motivation Principle (MP) according to which making a normative judgment of the form 'I should ϕ ' necessarily motivates one to behave in accordance with that judgment. Given MP, if Epistemic Normativity is true and epistemic judgments are normative judgments, then the Epistemic Motivation Thesis (EMT) must also be true: making an epistemic judgment of the form 'I epistemically ought to believe that P' must necessarily motivate one to behave in accordance with that judgment.

In this chapter, I offered a two-part case against EMT. First, I rejected expressivist accounts of epistemic judgments, which would plausibly entail EMT. Second, I turned to EMT directly and argued that, independently of the expressivism-cognitivism dispute, there is no plausible sense in which epistemic judgments – even those of the form 'I epistemically ought to believe that P' – necessarily motivate us to behave in accordance with them. I therefore conclude that the feature of normative judgments captured by MP – i.e. a necessary connection with motivation – is missing in epistemic judgments.

One final thing worth adding is that although MP is not universally accepted, many of those who deny it still think that cases where we are *not* motivated to do what we think we should do are very rare. Accordingly, many of them would still see it as a key feature of normative judgments that they *typically* or *ordinarily* motivate us behave in accordance with them. However, what I have shown in this chapter – especially in section 5.3 – is not only that epistemic judgments are not necessarily motivating, but that they are not even *typically* motivating. For one thing, we very often make what I called conclusive epistemic judgments, i.e. judgments where we think we epistemically should believe that P because we take the case for the truth P to be conclusive and thus take the question whether P to be settled. And as I have argued these typical epistemic judgments are not plausibly motivating since they already involve believing that P. So not only do epistemic judgments not have a

necessary connection with motivation, they do not even have a *close* connection with it.

Chapter 6: Control

6.1 Introduction: epistemology and 'ought' implies 'can'

Finally, a fifth feature that philosophers commonly attribute to normativity, but not with mere norm-relativity, is a necessary connection with *control*. In particular, many think that unlike merely norm-implying claims, normative claims according to which we should ϕ imply that ϕ -ing is, in some sense, under our control. 'Ought', as it is often put, implies 'can'. I will take this fifth commonly cited mark of normativity to be captured by the following principle:

Control principle (CP): *if S ought to ϕ then S can ϕ . And if S ought not to ϕ , then S can avoid ϕ -ing.*⁹²

Given CP, if Epistemic Normativity is true and epistemic claims are normative claims about what we ought or ought not to believe, then epistemic claims must also imply 'can'. That is, if CP and Epistemic Normativity are both true, then the following must also be true:

Epistemic Control Thesis (ECT): *if S is epistemically justified in believing that P, then S can believe that P. And if S is epistemically unjustified in believing that P, then S can avoid believing that P.*⁹³

The route from CP to ECT can be summarized as follows:

I. If S ought to ϕ , then S can ϕ . And if S ought not to ϕ , then S can avoid ϕ -ing. (CP)

⁹² See e.g. Zimmerman (1996), Haji (2002), Streumer (2003), Bykvist and Hattiangadi (2007), Vranas (2007), and Littlejohn (2012b).

⁹³ Or epistemically rational, reasonable, warranted, and the like. I only mention epistemic justification here for simplicity and because it is the most widely discussed of the allegedly normative epistemic notions. The arguments below apply equally to claims involving these other epistemic notions.

2. If S is epistemically justified in believing that P, then S (epistemically) ought to believe that P.⁹⁴ And if S is epistemically unjustified in believing that P, then S epistemically ought not to believe that P.⁹⁵ (*Epistemic Normativity*)
3. Therefore, if S is epistemically justified in believing that P, then S can believe that P and if S is epistemically unjustified in believing that P, then S can avoid believing that P. (*ECT*)

I argue for two main theses in this chapter.

First, ECT is false. Epistemic claims like attributions of epistemic justification do not imply ‘can’. I show this by (i) examining various commonly cited necessary conditions on ‘can’ in CP and (ii) offering examples where epistemic claims can still be true without these conditions being met. Therefore, no matter how exactly we construe the notion of control or ‘can’, in CP and ECT, it is not implied by epistemic claims.

Second, the falsity of ECT means that Epistemic Normativity is incompatible with CP. That is, given that epistemic claims do not imply ‘can’, if we accept that normative claims in general do imply ‘can’, then epistemic claims are not genuinely normative. I show this by ruling out a potential response according to which we can still keep Epistemic Normativity if ECT is false because (i) epistemic claims are not *deontic*, but rather *evaluative* norms, and (ii) only the former imply ‘can’. I argue that this avenue is not open to proponents of Epistemic Normativity. I concede that only deontic normative claims imply ‘can’, if CP is true. However, as I argued in chapter 2, epistemic norms do not necessarily imply genuine or normativity-grounding value. Therefore, if epistemic claims do turn out to be evaluative instead of deontic, then they are not genuinely normative since just like e.g. legal value, epistemic value does not necessarily constitute genuine or normativity-grounding value. Therefore, I conclude that the commonly cited mark of normativity that is captured by CP is missing in epistemic facts and claims.

⁹⁴ Some epistemologists argue that epistemic claims imply doxastic permissions and prohibitions, but not positive obligations or duties to believe. See for instance, Maitzen (1995), Nelson (2010), and Littlejohn (2012a). For these authors, epistemic justification, in other words, only has to do with what we *may* believe and what we *should not* believe. What I say in this chapter is entirely compatible with this thesis since (i) CP is commonly taken to apply also to permissions – i.e. many think that ‘may’ also implies ‘can’ – and (ii) even if CP turns out not to extend to permissions, it still applies to prohibitions and the arguments below show that epistemic prohibitions do not imply control.

⁹⁵ At this stage, I am assuming that epistemic norms are *deontic* or *prescriptive* rather than *evaluative*. I consider the suggestion that epistemic norms are evaluative in section 3.

6.2 Against the Epistemic Control Thesis (ECT)

6.2.1 'Can' explained

According to the Epistemic Control Thesis (ECT), epistemic claims like attributions of epistemic justification, rationality, and epistemic 'oughts' imply 'can' or control. But what does 'can' mean exactly in this context? What kind of control are normative claims supposed to imply according to the Control Principle (CP)? This is, of course, a deeply controversial question, which there is no space to adequately address here. Fortunately, as I will argue, it won't be necessary to answer it.

What I will do instead is consider four criteria that are invariably discussed in debates regarding necessary conditions for the truth of 'S can ϕ ' in 'ought' implies 'can'. I will refer to these four commonly discussed conditions as (i) Dual control, (ii) Know-how, (iii) Ability, and (iv) Opportunity. I understand each of these conditions as follows:

Dual control: *S can ϕ (avoid ϕ -ing) only if S can avoid ϕ -ing (can ϕ).*⁹⁶

Know-how: *S can ϕ (avoid ϕ -ing) only if S knows how to ϕ (to avoid ϕ -ing).*

Ability: *S can ϕ (avoid ϕ -ing) only if S is able to ϕ (to avoid ϕ -ing).*

Opportunity: *S can ϕ (avoid ϕ -ing) only if S has the opportunity to exercise her ability to ϕ (to avoid ϕ -ing).*

Of course, this list leaves several controversial questions unanswered. Let me mention two.

First, while Ability and Opportunity are widely accepted, there is extensive disagreement about whether to include the first two. On the one hand, some think that having only Ability and Opportunity would make CP too weak, thereby making 'oughts' too demanding. On the other, some reply that including the first two would

⁹⁶ A related condition is what philosophers such as Lavin (2004) call the *error constraint* on normativity. As Lavin puts it: "an agent is subject to a principle only if the agent can go wrong in respect of it." (2004, 425) If a norm requires you to ϕ , in other words, then you can violate that norm. As some of the examples below show, epistemic norms do not imply 'can violate' either.

make CP implausibly strong, thereby making 'oughts' too undemanding. Second, it is also controversial how exactly we should define each of these four forms of control. There is widespread disagreement, in other words, about what it takes to be able to do otherwise, to know how to ϕ , to have the ability to ϕ , and to have the opportunity to ϕ .

As I just mentioned however, there is neither space nor need for me to settle these issues here. This is because all I am going to argue is that whether or not 'can' implies Dual control and Know-how and no matter how we should understand each condition exactly, it is clear that epistemic claims do not imply *any* of them. That is, I will offer cases where S's belief that P is, on the face of it, epistemically justified or unjustified – i.e. where epistemic norms clearly apply to S and thus where S epistemically ought or ought not to believe that P – even though she (i) cannot do or believe otherwise, (ii) does not know how to believe that P, (iii) lacks the ability to believe that P, and (iv) has no opportunity to exercise her ability to believe that P. Therefore, I conclude that ECT is false. Epistemic claims do not imply 'can' or control.⁹⁷

6.2.2 ECT and Dual control

According to the Dual control condition, S 'can' ϕ only if S *can avoid* ϕ -ing. Conversely, S can avoid ϕ -ing only if S *can* ϕ . If this is right and if – as ECT maintains – epistemic claims imply 'can', then epistemic claims also imply Dual control. In other words, the following must also be true.

ECTdual-control: *if S is epistemically justified in believing that P, then S can avoid believing that P. And if S is epistemically unjustified in believing that P, then S can believe that P.*

One might think that ECTdual-control must be false because *doxastic voluntarism* is false, i.e. because we do not have direct voluntary control over our beliefs. After all, when I believe that it is raining, it is not as if I could just decide or will myself to believe that it is not raining. Yet I am epistemically justified in believing that it is raining.

⁹⁷ See also Mizrahi (2012) for the claim that the epistemic 'ought' does not imply 'can'.

This won't be my strategy however. To say that doxastic involuntarism suffices to rule out ECTdual-control is to claim that Dual Control – or 'can' more generally – requires *direct* control and that having *indirect* control is not sufficient for Dual Control. But this would be an implausibly demanding version of CP, which most of its contemporary proponents reject.⁹⁸ Even the strongest contemporary versions of CP allow that 'ought' only requires *indirect* control.⁹⁹

Instead, I will argue that even if it does not require direct control and doxastic voluntarism, ECTdual-control is vulnerable to counterexamples. Here is one example of an agent who clearly has an epistemically justified doxastic attitude, but for whom it is impossible to do otherwise.

Simon the neurosurgeon. Simon is a neurosurgeon who is obsessed epistemology and epistemic value. In particular, he is so deeply convinced of the badness of error that he invented a device that prevents his patients from believing what they do not know. Whenever there is any indication that a patient might be about to believe something she won't know, the device takes over her cognitive system and instantly causes her to suspend judgment about the matter at hand. Tanya is one of Simon's patients who has agreed to have the device implanted in her brain. After the procedure, Tanya goes out with a friend who, at some point in the night, asks her whether she thinks that the number of stars is even. Without any hesitation, Tanya replies that she has no way to know the answer and so suspends her judgment about that question. Because of her lack of hesitation, Simon's device did not activate and Tanya ended up suspending judgment by herself. However, if there had been the slightest indication that she might lean towards belief or disbelief, the device would have instantly taken over and caused her to suspend judgment.

In this example, Tanya could not have had the epistemically unjustified attitude of believing (or disbelieving) that the number of stars is even. That is, she could not have

⁹⁸ See e.g. Zimmerman (1996, 43-44) – especially his discussion of Ross (1939) and Pritchard (1949) – and Haji (2002, 23).

⁹⁹ This is why it is not plausible to reject, as e.g. Alston (1988) and Levy (2007) do, the possibility of doxastic 'oughts' on the basis of doxastic involuntarism and CP. Since doxastic involuntarism only entails the impossibility of direct doxastic control, their conclusion only follows if 'can' in CP implies direct control. See e.g. Chuard and Southwood (2009) and McHugh (2012b).

believed what she epistemically should not believe in that situation. Yet despite her inability to refrain from suspending judgment and her inability to violate the epistemic norm forbidding her to believe or disbelieve that the number of stars is even, Tanya's doxastic attitude is, by all accounts, epistemically justified. That is, she clearly believes what she epistemically ought to believe

So as this example shows, epistemic norms can apply to agents who cannot avoid doing as epistemic norms require. Therefore, ECT_{dual-control} is false. If 'can' requires dual control, then epistemic claims do not imply 'can' and ECT is false.

6.2.3 ECT and Know-how

According to the Know-how condition, S can ϕ only if S *knows how* to ϕ . Similarly, S can avoid ϕ -ing only if S knows how to avoid ϕ -ing. If this is right and if – as ECT maintains – epistemic claims imply 'can', then epistemic claims also imply Know-how. In other words, the following must also be true:

ECTknow-how: *if S is epistemically justified in believing that P, then S knows how to believe that P. And if S is epistemically unjustified in believing that P, then S knows how to avoid believing that P.*

ECTknow-how is also vulnerable to counterexamples however. Here are two examples where S is epistemically justified or unjustified in believing that P, but where S has no idea how to believe that P or to avoid believing that P.

Biased Bill. Bill regularly commits the gambler's fallacy. He is deeply convinced, for example, that if five coin tosses in a row land on heads, then the sixth one will almost certainly land on tails. However, Bill has no idea how probabilities work and what fallacies or biases are. Not only does he not know how to recognize fallacious reasoning, he does not even know that someone's beliefs can be fallacious or biased.

In this example, Bill does not know how to avoid the gambler's fallacy. Yet, that does not make his belief that the sixth coin toss will land on tails any less epistemically unjustified. It is still clearly something he epistemically should not believe.

Michael Jackson: Ever since Michael suffered a serious head injury in an accident, he has been convinced that he is the reincarnation of Michael Jackson. While he is aware that this is extremely implausible and improbable given everything he knows, he cannot help but believing it. He cannot even imagine himself as someone else.

In this example, Michael does not know how to believe that he is not Michael Jackson or to give up the belief that he is. Yet he is clearly epistemically unjustified in so believing. This is clearly still something that, from an epistemic point of view, he should not to believe.

As these examples show, there can be situations where S is epistemically unjustified in believing that P – i.e. where epistemic norms require S not to believe that P – even though S does not know how to avoid believing that P. Therefore, ECTknow-how is false. If ‘can’ requires Know-how, then epistemic claims do not imply ‘can’ and ECT is false.

6.2.4 ECT and ability

According to the Ability condition, S can ϕ only if S is *able* to ϕ . Similarly, S can avoid ϕ -ing only if S is able to avoid ϕ -ing. If this is right and if – as ECT maintains – epistemic claims imply ‘can’, then epistemic claims also imply ability. In other words, the following must also be true:

ECTability: *if S is epistemically justified in believing that P, then S is able to believe that P. And if S is epistemically unjustified in believing that P, then S is able to avoid believing that P.*

ECTability is also vulnerable to counterexamples however. Here are four cases where S is epistemically justified or unjustified in believing that P even though S is unable to believe or to avoid believing that P.

Modesty Blindspot. According to Maud, modesty is the most important of all the virtues. Unfortunately for her, she is also convinced that she is not a modest person at all and that she does not do enough to become more modest. Maud is wrong however. She is actually an extremely modest person.

This is supported by her evidence. All of her friends and relatives insist that she is extremely modest and she clearly thinks and behaves like a very modest person.

Maud's belief that she is not a modest person is, by all accounts, epistemically unjustified. Epistemically speaking, Maud should rather believe that she is modest or at least she epistemically should not believe that she is immodest.

Yet Maud cannot believe what her evidence tells her in this situation. As Sorensen (1988) argue, being modest is essentially a matter of underestimating one's self-worth. Hence Maud's going against her evidence is actually part of what makes her so modest. But if she followed her evidence and stopped underestimating her self-worth, she would automatically stop being a modest person and her believing that she is very modest would immediately provide her with conclusive evidence that she is *not* particularly modest. Her previously epistemically unjustified belief (that she is not modest) would therefore become justified and she would be immediately epistemically *unjustified* in believing that she is modest.

Here is a second example:

Evidential blindspot. Daniel is an experienced epistemology professor and one of his students, Mark, is the worst epistemology student he has ever seen. In class, Mark could never stay focused for more than a few seconds and could never understand what was going on. Even though he attended all the lectures, he still has literally no idea what epistemology is supposed to be about. Unsurprisingly, Mark failed the class miserably. His performance was so bad and his understanding so poor that Daniel concluded not only that Mark is incapable of passing an epistemology class, but also that *he cannot even form beliefs about evidence*. At least this is what the evidence indicates and what any reliable observer would conclude.

Daniel seems epistemically justified in believing that Mark cannot form beliefs about evidence. Interestingly, a lot of the evidence that support this conclusion is also possessed by Mark. He is well aware that he could never understand what was going on, that he could never focus more than one or two seconds in the class, that he failed miserably, and so on. Hence, just like Daniel, Mark is epistemically justified in

believing that he cannot form beliefs about evidence. This is, after all, what his evidence supports and what any reliable observer would conclude.

However, Mark cannot believe what his evidence supports in this situation. This is because if he were to form the belief *that he cannot form beliefs about evidence*, his so believing would immediately provide him with conclusive evidence that he *can* form beliefs about evidence. He would thereby become epistemically unjustified in believing that he cannot form beliefs about evidence. Mark is therefore unable to believe as he epistemically should.

Here is a third example:

Impostor syndrome. Anne suffers from severe impostor syndrome. She is deeply convinced that her intellectual aptitudes are well below average and that her achievements are all due to pure luck. This, however, clearly goes against what she knows. As she acknowledges herself, she always scores very high on standardized tests, everyone tells her how smart she is, and she has a long list of impressive academic and professional achievements. Yet she has this unshakable feeling that she is worthless. Anne knows this feeling is not supported by any evidence. Yet it is always there and it causes her to systematically disregard or explain away her actual evidence by citing luck and her ability to appear smart. It also prevents her from doing anything to get rid of that feeling and to do something about that impostor syndrome.

In this case, Anne lacks the ability to avoid believing that she is worthless and that her intellectual aptitudes are well below average. Nevertheless, that belief is, by all accounts, epistemically unjustified. It is clearly something that, epistemically speaking, she should not believe.

Finally, Mizrahi (2012) also provides an example that seems refute ECTability. It is a case that he borrows from the 2010 film *The Next Three Days*. He summarizes it as follows:

In this movie, the life of a family of three takes a turn for the worse when the wife is accused of murder. The forensic evidence against the wife, Lara, is compelling, and it includes fingerprints, an eyewitness account, ballistics reports, DNA from blood samples, and a clear motive. Lara is convicted of murder and sentenced to life in prison. But her husband, John, cannot believe

that she is a murderer. Now, in the film, there are hints pointing to Lara's innocence. For the sake of argument, however, suppose that she is in fact guilty of murder. [...] What we have here is a scenario in which John ought to believe that Lara is guilty of murder, but he simply cannot believe that his wife is a murderer. [...] John ought to believe that his wife is a murderer, and yet he cannot believe that his wife is a murderer. In other words, as an epistemic agent, John is required to believe that Lara is guilty of murder, but John lacks the specific ability and opportunity to believe that his wife is a murderer. (Mizrahi 2012, 832-833)

Although John lacks the ability to believe that P, he would be epistemically justified in believing that P. Moreover, even though John is unable to avoid believing that his wife is innocent, he is still epistemically unjustified in believing that. This is still something, in other words, that he epistemically should not believe.

As these examples show, there can be situations where S is epistemically justified or unjustified in believing that P even though S lacks the ability to believe that P or to avoid believing that P. Therefore, ECTability is false. If 'can' requires ability, then epistemic claims do not imply 'can' and ECT is false.

6.2.5 ECT and opportunity

Finally, according to the Opportunity condition, S can ϕ only if S has the *opportunity* to exercise her ability to ϕ . Similarly, S can avoid ϕ -ing only if S has the opportunity to exercise her ability to avoid ϕ -ing. If this is right and if – as ECT maintains – epistemic claims imply 'can', then epistemic claims also imply Opportunity. In other words, the following must also be true:

ECTopportunity: *if S is epistemically justified in believing that P, then S has the opportunity to exercise her ability to believe that P. And if S is epistemically unjustified in believing that P, then S has the opportunity to exercise her ability to avoid believing that P.*

ECTopportunity is also vulnerable to counterexamples however. Below are two cases where agents are clearly epistemically unjustified in believing that P despite being in situations that do *not* allow them to exercise their ability to avoid believing that P and to believe as epistemic norms require.

Logical Linda. Linda is a participant in a psychological experiment that requires her to solve a fairly easy logical puzzle. Linda is very intelligent, but she initially makes a silly reasoning mistake and concludes, incorrectly, that the answer is A. Since her evidence clearly indicates that the answer is B and since the answer is fairly obvious, she would realize her mistake and form the correct belief if she carefully thought about it for a few more minutes. Unbeknownst to Linda however, the neurosurgeon conducting the experiment does not want her to get it right because it would contradict his hypothesis. So once he sees that she gets it wrong, he immediately gives Linda an anaesthetic and installs a microscopic chip in her brain that will prevent her from thinking about the puzzle ever again. Every time she will start thinking about the puzzle, the chip will immediately distract her or make her think about something else. The experimenter therefore ensures that Linda won't ever have the opportunity to correct her initial belief. She thus goes on believing that the answer is A.

In this example, Linda lacks the opportunity to exercise her ability to figure out the right answer and to stop believing that the answer is A. Crucially however, her belief that the answer is A remains epistemically unjustified. This is still something she epistemically should not believe.

Fear of death. Because of his extreme fear of death Rob paid a hypnotist to 'program' him to instantly fall asleep whenever he starts thinking about his own death. Since before the treatment, Rob has had the belief that he will live at least until he is 85. Unfortunately, a few months after his treatment, Rob is diagnosed with a serious illness that will significantly reduce his life expectancy. After the diagnosis, Rob's evidence conclusively shows that he will *not* live in his eighties. He knows he has the illness and that no one who has had that illness went on to live older than 65. However, because of his hypnotism treatment, he never gets to revise his initial belief that he will live past 85 since he has to think about his own death in order to do so, which makes him fall asleep every time.

Even though Rob lacks the opportunity to change it, his initial belief that he will live past 85 is still epistemically unjustified. It is clearly false given what he knows and so this is still something he epistemically should not believe.

As these examples show, there can be situations where S is epistemically justified or unjustified in believing that P even though S lacks the opportunity to exercise her ability to believe that P or to avoid believing that P. Therefore, ECT_{opportunity} is false. If 'can' requires opportunity, then epistemic claims do not imply 'can' and ECT is false.

6.2.6 Taking stock

I identified four criteria that are often seen as necessary conditions for control or 'can' in 'ought' implies 'can'. These are Dual control, Know-how, Ability, and Opportunity. I have shown that given each of these four conditions, the Epistemic Control Thesis (ECT) is vulnerable to counterexamples. These are examples that show, in other words, that epistemic claims do not imply any of these four forms of control. Since there is no plausible sense of 'can' that is implied by epistemic claims, I conclude that ECT is false. Epistemic claims or 'oughts' do not imply 'can'.

6.3 *The incompatibility of CP and Epistemic Normativity.*

In this section, I argue that since ECT is false, Epistemic Normativity is incompatible with the Control Principle (CP). Since epistemic claims do not imply 'can', in other words, if normative claims imply 'can', then epistemic claims are not normative. This needs to be defended because even if ECT is false, there is a potential way for proponents of Epistemic Normativity to keep CP. This potential avenue for Epistemic Normativity is what I will call the *evaluative* option. According to the evaluative option, even if epistemic claims do not imply 'can', we can still keep both CP and Epistemic Normativity because (i) epistemic claims are not *deontic*, but rather *evaluative* normative claims, and (ii) only the former imply 'can'.

6.3.1 Epistemic norms: deontic or evaluative?

So far, I have assumed that epistemic norms or 'oughts' are best seen as *deontic* or *prescriptive* 'oughts'. Deontic 'oughts' are 'ought to do', i.e. 'oughts' that belong to, apply to, or place demands on agents. So on a deontic reading, epistemic norms make demands or prescriptions to agents about what to believe or not to believe. However,

norms can also be *evaluative* rather than deontic. Rather than implying prescriptions or demands on agents, evaluative ‘oughts’ only imply a claim about the value or goodness of something, i.e. that something ought to *be* the case.¹⁰⁰ This distinction matters for our purpose because if CP is true, then it only applies to deontic ‘oughts’. Even if deontic norms do imply control, in other words, evaluative norms do not.

6.3.2 The evaluative option

Given this distinction and its relation with CP, one potentially attractive option for proponents of Epistemic Normativity who want to keep CP is to argue that epistemic norms are evaluative instead of deontic. On this construal of epistemic norms, to say that S is epistemically justified or unjustified in believing that P is not to say that S is required, permitted, or forbidden to believe that P, but rather that S’s believing that P would be good or bad from an epistemic standpoint.¹⁰¹ If this is right, then the falsity of ECT – the fact that epistemic claims do not imply ‘can’ – is not an issue for Epistemic Normativity. It is simply due to the fact that epistemic norms are evaluative rather than prescriptive. This is what I call the evaluative option.

The problem with the evaluative option is that it only helps Epistemic Normativity if epistemic value necessarily constitutes genuine value. As I explained in chapter 2, we must distinguish genuine or normativity-grounding value from mere domain-relative value. Examples of the latter include fashion value, legal value, and the like. So although it might be right that epistemic claims are evaluative and that they necessarily imply epistemic value, it cannot be of any help to Epistemic Normativity unless epistemic value necessarily constitute genuine or normativity-grounding value and not merely domain-relative value. But as I argued in chapter 2, epistemic claims do *not* necessarily imply genuine or normativity-grounding value. Just like e.g. legal value or fashion value, epistemic value does not necessarily constitute genuine or normativity-grounding value. So if epistemic claims do turn out to be evaluative instead of deontic, then they are not normative since it is not necessarily genuinely good to do what is *epistemically* good. Therefore, the evaluative option is not open to proponents of Epistemic Normativity.

¹⁰⁰ For more on this distinction and its relation to CP, see e.g. Zimmerman (1996, chapter 3) and Haji (2002). See also e.g. Schroeder (2011) and Chrisman (2015) on the distinction itself.

¹⁰¹ See for instance Alston (1993), Chrisman (2008), Fassio (2011), McHugh (2012a), Jarvis (2012), and Engel (2013).

So in sum, epistemic norms or claims are either deontic or evaluative. If they are deontic, then given CP, Epistemic Normativity entails that epistemic claims imply 'can'. But since epistemic claims do not imply 'can', this first deontic option is not available to proponents of Epistemic Normativity who want to keep CP. If they are evaluative, then Epistemic Normativity does not entail that epistemic claims imply 'can' (even if CP is true), but it requires that epistemic value be necessarily genuine or normativity-grounding value. But since epistemic value is not necessarily normativity-grounding value, this evaluative option is not available to proponents of Epistemic Normativity either. Therefore, if CP is true, then Epistemic Normativity is false.

6.3.3 What about Epistemic Non-Normativity?

Is Epistemic Non-Normativity compatible with CP? Recall that according to Epistemic Non-Normativity, epistemic claims are norm-implying but not normative. Epistemic norms or 'oughts' do not necessarily have normative authority.

Unlike Epistemic Normativity, Epistemic Non-Normativity is compatible with the Control Principle. Even if CP is true, norms that lack necessary normative authority do not imply 'can'. Take the norms of fashion for example. Suppose you cannot help dressing in a way that is outrageously unfashionable and out of style. Perhaps you suffer from a psychological condition that compels you to continually wear a particular outfit. Or perhaps someone else physically forced you to dress in that way. But the fact that you cannot avoid wearing that outfit does not make it any less unfashionable and out of style. It is unfashionable and tacky regardless of whether you had any control over wearing it. So as this example illustrates, norms that lack necessary normative authority do not imply 'can'. Thus if, as Epistemic Non-Normativity claims, epistemic norms also lack necessary normative authority, then the falsity of ECT is precisely what we would expect, even if we accept CP in the first place i.e. even if we accept that deontic normative claims do imply 'can'.

One potential reply to this example is that fashion norms might very well be evaluative and not deontic. So then perhaps the reason why they do not imply 'can' is not their lack of normative authority, but rather their evaluative character. But what if epistemic norms are best seen as deontic and not evaluative? Wouldn't Epistemic Non-Normativity be equally incompatible with the Control Principle (CP)?

This worry is misguided however. First, norms can very well be deontic without having necessary normative authority. Norms or rules of etiquette, games, and the law, for example, clearly seem deontic and not evaluative. Hence, Epistemic Non-Normativity is compatible with both deontic and evaluative conceptions of epistemic norms. Second, these deontic norms that lack necessary normativity do not imply 'can' either. For example, even if you suffer from extreme, treatment-resistant kleptomania (and so cannot refrain from stealing) it remains true that the law forbids you to steal. Whether or not you have control over your stealing, in other words, your stealing remains illegal and something that you legally ought not to do.

Therefore, whether epistemic norms are best seen as deontic or evaluative, Epistemic Non-Normativity is compatible with CP because whether or not CP is true, norms that lack necessary normative authority do not imply 'can'.

6.4 Summary

A fifth commonly cited mark of normativity is a necessary connection with control or 'can'. I took this feature to be captured by the Control Principle (CP) according to which if S ought to ϕ then S can ϕ , and if S ought not to ϕ , then S can avoid ϕ -ing. Given CP, if Epistemic Normativity is true and epistemic claims are normative claims, then the Epistemic Control Thesis (ECT) must also be true: if S is epistemically justified in believing that P, then S can believe that P. And if S is epistemically unjustified in believing that P, then S can avoid believing that P.

I argued for two main theses in this chapter. First, ECT is false. Epistemic claims like attributions of epistemic justification do not imply 'can'. I showed this by offering counter-examples to each of the most commonly cited necessary conditions on 'can' in CP. Second, the falsity of ECT means that Epistemic Normativity is incompatible with CP. That is, given that epistemic claims do not imply 'can', if we accept that normative claims do imply 'can', then epistemic claims are not normative. I showed this by ruling out what I called the evaluative option for proponents of Epistemic Normativity. Hence, I conclude that epistemic claims do not bear the mark of normativity that is captured by the Control Principle (CP).

PART II: EXPLAINING THE MISSING MARKS OF NORMALITY

Chapter 7: Two rival explanations

In part I, I argued that epistemic facts and claims lacks five features that are commonly attributed to normative facts and claims, but not to merely norm-implying ones. Of course, this is not sufficient to establish Epistemic Non-Normativity. After all, I have not shown that any of these features actually are necessary features of normative facts and claims, i.e. that any of the five principles from chapters 2-6 are actually true. There is no space to show that here and hence it won't be my strategy for arguing in favour of Epistemic Non-Normativity.

Instead, what I would like to show in this second part is that the conclusions of part I are *best explained* by Epistemic Non-Normativity. In this chapter, I clarify the explanations that Epistemic Normativity and Epistemic Non-Normativity can each provide of the conclusions of chapters 2-6. I also argue that Epistemic Non-Normativity can accommodate additional features of epistemic facts and claims which they seem to have in common with normative facts and claims. In the next chapter, I will argue that given commonly invoked explanatory virtues or criteria for theory choice, the explanation provided by Epistemic Non-Normativity is superior to the one that Epistemic Normativity must give.

7.1 Epistemic Normativity's explanation of the missing marks: Normative Pluralism

If epistemic facts and claims are normative facts and claims – as Epistemic Normativity maintains – then why do they lack the five commonly cited features of normativity from part I? How can proponents of Epistemic Normativity answer this question?

A first thing to note is that part of their answer must be a rejection of the five principles I outlined in part I. Given the conclusions of chapters 2-6, Epistemic Normativity is incompatible with the Value Principle (VP), the Instrumental Principle (IP), the Autonomy Principle (AP), the Motivation Principle (MP), and the Control Principle (CP). If epistemic facts and claims are genuinely normative, in other words, then it cannot be the case that all normative facts and claims imply (i) value, (ii) the promotion of our desires, (iii) autonomy from the non-normative, (iv) motivation, and (v) control. Therefore, an Epistemic Normativity-friendly explanation of part I must involve the claim that none of the five commonly cited features are actually essential marks of normativity *in general*. It must claim that facts and claims can lack all five features and still be genuinely normative.

A second thing to note, however, is that some normative facts and claims plausibly do imply at least some of these features. That is, even if the features from part I are not marks of normativity *in general*, it is plausible that at least some of them are marks of *moral* and *prudential* normativity. This is something that Epistemic Normativity must accommodate. Even if it rejects all five principles from chapters 2-6 (since these principles purport to be about normativity in general), Epistemic Normativity must still allow the possibility of moral, prudential, or more generally *practical* versions of these principles. This is because most accounts of practical normativity are committed to at least some of these principles.

First, most – if not all – versions of practical expressivism as well as many versions of cognitivism imply a practical version of the Motivation Principle (MP) from chapter 5, i.e. the thesis that practical normative judgments of the relevant form are necessarily motivating. Second, *non-naturalist* accounts of practical normativity are typically committed to a practical version of the Autonomy Principle from chapter 4. This is because non-naturalists typically invoke this principle to motivate their view and to reject normative naturalism.¹⁰² We are therefore left with accounts of practical normativity that view practical facts as (i) natural facts that (ii) are not necessarily motivating. But the problem is that it is hard to find accounts that meet those two conditions, but that do not also accept at least a moral or prudential version of the Value Principle from chapter 2, i.e. accounts that do not also claim that it is necessarily at least *pro tanto* good to conform to moral norms or to prudential norms.

Therefore, if it is to have any plausibility, Epistemic Normativity's explanation must allow that at least some of the five features might still be essential marks of *other* kinds of normativity such as moral and prudential normativity. Consequently, a second part of Epistemic Normativity's explanation must be that there is not only one kind of normativity, but rather a plurality of them. That is, it must claim that there are different kinds of normative facts, which can have different essential characteristics. Some of these kinds of normativity – e.g. moral and prudential normativity – have at least some of the features from part I as essential marks, while others – including at least epistemic normativity – do not. The resulting explanation can be summarized as follows:

¹⁰² See for instance Moore (1903), Johnston (1989), Scanlon (1998), (2014), Dancy (2006), Fitzpatrick (2008), (2010), Enoch (2011a), and Parfit (2011).

Normative Pluralism: *Epistemic facts and claims do not bear the five commonly cited marks of normativity from part I because (i) there is a plurality of kinds of normativity and (ii) the kind of normativity implied by epistemic facts and claims – i.e. epistemic normativity – lacks these five features.*¹⁰³

I will take this to be the most promising explanation that proponents of Epistemic Normativity can give of the conclusions of part I.

7.2 *Epistemic Non-Normativity's explanation of the missing marks*

How does Epistemic Non-Normativity explain the fact that epistemic facts and claims lack the features from part I? Just like Normative Pluralism, the Epistemic Non-Normativity explanation has two key components. The first and most straightforward stems from Epistemic Non-Normativity's thesis that epistemic facts and claims are not normative. If this is right, then it is not surprising that alleged essential features of normative facts and claims are missing in epistemic facts and claims. To the contrary, it is exactly what we would expect.

The second stems from Epistemic Non-Normativity's more positive claim, i.e. that epistemic facts and claims are still *norm-implying* even though they are not normative. Recall that according to Epistemic Non-Normativity, epistemic claims still have to do with what we should or should not believe according to epistemic norms. It is just that these epistemic norms lack necessary normative authority. The second main component of its explanation is therefore the following. Whether or not the five features from part I are essential features of normative facts and claims, they are *not* essential features of merely norm-implying facts and claims. Even if the five characteristics from chapters 2-6 are marks of normativity, in other words, they are not marks of mere norm-relativity. The resulting explanation can therefore be summarized as follows:

¹⁰³ Several prominent normativity theorists can be interpreted as defending something like Normative Pluralism. One example is Raz (1999), (2011) who distinguishes between value-based *practical* normativity on the one hand, and *adaptive* normativity, which does not imply value, on the other. See also e.g. Scanlon (1998), (2014), Skorupski (2010), and Parfit (2011). See Reisner (2015) for relevant discussion. Note that what I call Normative Pluralism should not be confused with Tiffany's (2007) 'normative pluralism'.

Epistemic Non-Normativity explanation: *the five marks of normativity from part I are missing in epistemic facts and claims because (i) epistemic facts and claims are not normative, but rather (ii) merely norm-implying, and the five features from part I are not essential features of merely norm-implying facts and claims.*

While the first part of this explanation is straightforward, the second could conceivably be questioned. That is, someone could conceivably question the claim that the features from chapters 2-6 are not essential features of norm-implying facts and claims.

In the rest of this section, I rule out this possibility. More precisely, I give examples to show that merely norm-implying facts and claims do not necessarily imply (i) value, (ii) the promotion of one's desires, (iii) autonomy from the non-normative, (iv) motivation, and (v) control. Even if these are marks of normativity, in other words, they are not marks of mere norm-relativity.

7.2.1 Norm-relativity and value

The first commonly cited mark of normativity I examined is a necessary connection with value. I took this first characteristic to be captured by the Value Principle (VP):

Value principle (VP): *there is a normative reason to ϕ if and only if it is good to ϕ .*

Given VP, if Epistemic Normativity is true, then the following must also be true:

Epistemic value thesis (EVT): *it is necessarily good to conform to epistemic norms.*

However, as I argued in chapter 2, EVT is false. It is not necessarily good to conform to epistemic norms. Hence, epistemic facts and claims do not have the necessary connection with value that is captured by VP. According to Epistemic Non-Normativity, EVT is false because epistemic facts and claims are merely norm-implying and not normative. That is, the norms that they imply lack necessary normative authority and it is not necessarily good to conform to norms that lack such authority.

Even if we accept that normative claims are necessarily connected to value, in other words, merely norm-implying ones are not.

To see this, consider other examples of norms that lack necessary normative authority. Plausibly, there is not necessarily something good in conforming to e.g. etiquette, fashion, the law, tradition, and so on. To use the distinction I drew in chapter 2 between mere domain-relative goodness and genuine or normativity-grounding goodness, it is not necessarily genuinely good to do what is legal, fashionable, decorous, traditional, and so on (i.e. what is good or correct from a legal, fashion, etiquette, or traditional point of view). Unlike, say, moral value or prudential value, legal value, fashion value, etiquette value, traditional value, and the like are merely domain-relative. They do not necessarily constitute normativity-grounding value. There is not necessarily anything genuinely good, for example, in conforming to immoral or absurd laws, traditions, or etiquette rules, especially if you have no chance of getting caught or punished. More generally, although it is very often – perhaps almost always – genuinely good to do what is legal, fashionable, decorous, traditional, and so on, it is not necessarily so.

The conclusion of chapter 2 is that epistemic norms are just like that. Although it is very often – perhaps almost always – genuinely good to conform to them, it is not necessarily so. Unlike perhaps moral and prudential value, epistemic value is merely domain-relative and does not necessarily constitute genuine or normativity-grounding value. According to Epistemic Non-Normativity, this is exactly what we should expect since epistemic facts and claims are merely norm-implying and this is precisely the kind of connection that merely norm-implying facts and claims have with genuine or normativity-grounding value.

7.2.2 Norm-relativity and desire

The second commonly cited mark of normativity I examined is a necessary connection with the satisfaction of one's desires. I took this second feature to be captured by the following principle:

Instrumental Principle (IP): *There is a normative reason for one to ϕ if and only if ϕ -ing promotes the satisfaction of one's desires.*

Given IP, if Epistemic Normativity is true, then the following must also be true:

Epistemic Instrumental Thesis (EIT): *conforming to epistemic norms necessarily promotes the satisfaction of one's desires.*

However, as I argued in chapter 3, EIT is false. Conforming to epistemic norms does not necessarily promote the satisfaction of one's desires. Hence, epistemic facts and claims do not have the necessary connection with desires that is captured by IP.

According to Epistemic Non-Normativity, EIT is false because epistemic facts and claims are merely norm-implying and not normative. That is, the (epistemic) norms that they imply lack necessary normative authority and conforming to norms that lack such authority does not necessarily promote the satisfaction of one's desires. Even if we accept that normative claims are necessarily connected to the satisfaction of one's desires, in other words, merely norm-implying ones are not. To see this, consider other examples of norms that lack necessary normative authority. Plausibly, conforming to e.g. norms of law, fashion, etiquette, tradition, religious sects, and the like does not necessarily promote the satisfaction of one's desires. Fashion could require you to wear outfits that you hate and to avoid wearing outfits that you love. Similarly, tradition could require you to marry someone you hate and forbid you to have a relationship with someone you love. You could also get caught in a sect that requires you to commit suicide. More generally, there can obviously be cases where being lawful, fashionable, decorous, traditional, and so on does *not* promote any of your desires.

The conclusion of chapter 3 is that epistemic norms are just like that. Although conforming to them very often – perhaps almost always – does promote the satisfaction of one's desires, it does not necessarily do so. According to Epistemic Non-Normativity, this is exactly what we should expect since epistemic facts and claims are merely norm-implying and this is precisely the kind of connection that merely norm-implying facts and claims have with desire.

7.2.3 Norm-relativity and autonomy

The third commonly cited mark of normativity I examined is a necessary autonomy from the non-normative. I took this third feature to be captured by what is often known as Hume's law or as I put it:

Autonomy Principle (AP): *normative claims cannot be derived from non-normative claims alone.*

Given AP, if Epistemic Normativity is true, then the following must also be true:

Epistemic Autonomy Thesis (EAT): *epistemic claims cannot be derived from non-normative claims alone.*

However, as I argued in chapter 4, EAT is false. Epistemic claims *can* be derived from non-normative claims alone. Hence, epistemic facts and claims are not autonomous from the non-normative.

According to Epistemic Non-Normativity, EAT is false because although they are norm-implying, epistemic claims are not normative and even if normative claims are necessarily autonomous from non-normative ones, merely norm-implying claims are not. Even if we accept AP, in other words, that principle does not extend to merely norm-implying claims like epistemic claims.

Consider legal claims for instance. The fact that ϕ -ing is legally required, permitted, or forbidden can very well be derived from non-normative premises alone. In particular, from the (non-normative) fact that a country's code of law says that it is forbidden to ϕ – e.g. if it is written in its constitution or was asserted by judges on its supreme court – we can conclude that in that country, it is illegal to ϕ , that we *legally should not* ϕ , and the like.

In chapter 4, I rejected the autonomy of epistemic claims from the non-normative by examining two phenomena, namely (i) first-person epistemic deliberation and (ii) what I called descriptive-epistemic questions. I argued that non-normative facts alone *can* settle epistemic deliberation from the first-person perspective, and that descriptive-epistemic questions can be closed or trivial questions and are not necessarily open or substantial questions. Note that the same is true of merely norm-implying claims like legal claims. First, non-normative facts alone can settle for oneself the legal deliberative question 'legally speaking, should I ϕ ?' In particular, it can be settled by non-normative facts about what the law says. Once I come to the conclusion that the law in my country says it is forbidden to ϕ , then the question whether I should ϕ from a legal point of view is settled for me.

Second, descriptive-legal questions can very well be closed or trivial questions. For example, the question ‘ ϕ -ing is disallowed according to the US constitution, but is it legal to ϕ in the US?’ is not an open question.

The conclusion of chapter 4 is that epistemic claims are just like that. Just like legal conclusions, epistemic conclusions can be derived from non-normative conclusions alone. According to Epistemic Non-Normativity, this is exactly what we should expect since epistemic facts and claims are merely norm-implying and this is precisely the kind of connection that merely norm-implying claims have with non-normative claims.

7.2.4 Norm-relativity and motivation

The fourth commonly cited mark of normativity I examined is a necessary connection with motivation. I took this fourth feature to be captured by the Motivation Principle (MP):

Motivation principle (MP): *making a normative judgment of the form ‘I should ϕ ’ necessarily motivates one to behave in accordance with that judgment.*

Given MP, if Epistemic Normativity is true, then the following must also be true:

Epistemic Motivation Thesis (EMT): *making an epistemic judgment of the form ‘I epistemically should ϕ ’ necessarily motivates one to behave in accordance with that judgment.*

However, as I argued in chapter 4, EMT is false. Epistemic judgments are not necessarily motivating.

According to Epistemic Non-Normativity, EMT is false because although they are norm-implying, epistemic judgments are not normative and even if normative judgments are indeed motivating, merely norm-implying ones are not. Even if we accept MP, in other words, that principle does not extend to merely norm-implying judgments like epistemic judgments. Consider judgments about what one should or should not do relative to certain gender stereotypes. I can very well judge, for instance, that I should not cry in public given standards of masculinity or machismo

and not be motivated at all to refrain from crying in public. Similarly, I can very well judge sincerely that I should ϕ according to e.g. fashion, the law, or etiquette without being motivated to ϕ to any extent.

The conclusion of chapter 5 is that epistemic judgments are just like that. Just like judgments about what I should or should not do according to masculinity standards, the law, fashion, etiquette, and the like, I can very well make judgments about what I *epistemically* should or should not believe without being motivated to behave accordingly to any extent. According to Epistemic Non-Normativity, this is exactly what we should expect since epistemic judgments are merely norm-implying and this is precisely the kind of connection that merely norm-implying claims have with motivation.

7.2.5 Norm-relativity and control

The final commonly cited mark of normativity I examined is a necessary connection with control. I took this fifth feature to be captured by the well-known thesis that ‘ought’ implies ‘can’ or as I put it:

Control Principle (CP): *if S ought to ϕ then S can ϕ . And if S ought not to ϕ , then S can avoid ϕ -ing.*

Given CP, if Epistemic Normativity is true, then the following must also be true:

Epistemic control thesis (ECT): *if S is epistemically justified in believing that P, then S can believe that P. And if S is epistemically unjustified in believing that P, then S can avoid believing that P.*

However, as I argued in chapter 6, ECT is false. Epistemic claims, norms, and ‘oughts’ do not imply ‘can’. Hence, epistemic facts and claims do not have the necessary connection with control captured by CP.

According to Epistemic Non-Normativity, ECT is false (epistemic claims do not imply ‘can’) because although they are norm-implying, epistemic claims are not normative. So even if normative claims imply control, merely norm-implying ones do not. That is, norms and ‘oughts’ that lack necessary normative authority do not imply

'can'. Once again consider legal claims. As I mentioned in chapter 6, even if you suffer from extreme, treatment-resistant kleptomania (and so cannot refrain from stealing) it remains *illegal* for you to steal. Whether or not you have control over your stealing, in other words, it remains true that you legally ought not to steal.

The conclusion of chapter 6 is that epistemic claims, norms, and 'oughts' are just like that. Just like in the case of legal 'oughts', the fact that you *epistemically* ought (ought not) to believe that P does not entail that you can – in the relevant sense – believe (not believe) that P. According to Epistemic Non-Normativity, this is exactly what we should expect since epistemic claims are merely norm-implying and this is precisely the kind of connection that merely norm-implying claims have with control.

7.3 Interlude: explaining apparent remaining marks of normativity

So far, I have focused on the need to explain why the five commonly cited marks of normativity from part I are missing in epistemic facts and claims. As we have just seen, Epistemic Non-Normativity's explanation is simply that epistemic facts and claims are not genuinely normative, but rather merely norm-implying.

An important thing to note, however, is that although they do not have the features from part I, epistemic facts and claims still have some features that one might claim to be additional marks of normativity. In particular, epistemic facts still seem to be (i) *important*, (ii) *inescapable*, (iii) *sources of epistemic reasons*, and (iv) *universal or absolute*. But one might doubt that epistemic facts can still have these features if they are merely norm-implying. So this raises a second explanatory question: if Epistemic Non-Normativity is true and epistemic facts and claims are not normative, then why do they have these features that also seem to be hallmarks of normative facts and claims?

In the rest of this chapter, I argue that Epistemic Non-Normativity can provide a satisfactory answer to this second explanatory questions as well. That is, it can convincingly accommodate these four additional features of epistemic facts. Its answer is, roughly, that these characteristics are not exclusive to normative facts. Even if epistemic facts are merely norm-implying, they can still be important, inescapable, a source of epistemic reasons, and universal or absolute. I defend each of these claims in the rest of this chapter.

7.3.1 Norm-relativity and importance

It seems that just like e.g. moral norms and unlike e.g. norms of etiquette or fashion, epistemic norms *matter*. However, if epistemic norms have the same kind of authority as the latter, doesn't it follow, implausibly, that epistemic norms are similarly trivial and unimportant?

Note first that the claim cannot be that it *always* or *necessarily* matters whether you conform to epistemic norms. This is because this would amount to the Epistemic Value Thesis (EVT), which I rejected in chapter 2. The point made in chapter 2, in other words, was precisely that it does not necessarily matter whether you conform to epistemic norms.

Instead, this first feature is most plausibly interpreted as the claim that epistemic norms are *an important kind* of norm or that they matter *in general*. This, however, is something that Epistemic Non-Normativity can easily accommodate. All that is needed for a norm N to lack necessary normative authority, is that there *can* be cases where there is no normative reason to conform to N. Crucially however, this is compatible with such cases being very rare. Perhaps there is *almost always* a normative reason to conform to N. But this, in turn, leaves the possibility that it *almost always* matters whether you are conforming to N. But if this is so, then N is clearly an important kind of norm. Even if there is not necessarily a normative reason to conform to the law, for instance, there is still almost always such a reason. After all, it is moral and prudent to conform to most laws most of the time. Similarly, it is very often a good idea to conform to norms of fashion and etiquette. Very often, in other words, there is a normative reason to do as these norms say. So even norms that lack necessary normative authority can have tremendous importance.

Epistemic Non-Normativity is perfectly compatible with epistemic norms being just like that, namely norms to which there is almost always, but not necessarily a normative reason to conform. This, in turn, is compatible with the possibility that it almost always matters whether you conform to epistemic norms. There are, after all, countless subjects and situations about which it is important to know the truth and not be mistaken. So as Epistemic Non-Normativity can maintain, the situations where conforming to epistemic norms does not matter are plausibly very rare, which means that epistemic norms are still a very important kind of norms.

7.3.2 Norm-relativity and inescapability

Just like moral norms, epistemic norms seem *inescapable* or *categorical*. Agents' acts are morally right or wrong regardless of what they care about. Intuitively, even if you do not care at all about being moral, it remains true that you morally should not e.g., humiliate others. In the same way, epistemic facts apply to us regardless of what we want or care about. Whether our beliefs are epistemically justified, rational, correct, and the like does not depend on our desires or goals. Even if a wishful thinker does not want to believe the truth, avoid error, follow her evidence, and the like, her wishful belief is still epistemically unjustified. How can Epistemic Non-Normativity explain that?

The fact that a norm is inescapable does not entail that it has necessary normative authority, i.e. that it is *inescapably normative*.¹⁰⁴ Norms that lack necessary normative authority can still apply to agents independently of what they care about. It is true that you cannot simply escape or opt out of morality by not caring about morality or by not playing the 'morality game'. But this is also true of norms like etiquette, fashion and the law. You cannot simply opt out of these norms by not caring about them or by not playing the 'etiquette game', the 'fashion game', or the 'legal game'. Your unfashionable outfits are unfashionable whether or not you care about being fashionable, your illegal acts are illegal whether or not you care about the law, and your indecorous acts are indecorous whether or not you care about etiquette.

According to Epistemic Non-Normativity, epistemic norms are just like that. Even though they lack necessary normative authority – i.e. even though they are not inescapably normative – epistemic norms still inescapably apply to us. Your epistemically unjustified beliefs are epistemically unjustified whether or not you care about avoiding error, believing the truth, following the evidence, and the like. But this is so even if you are in a situation where there is no normative reason for you to conform to epistemic norms and believe what is epistemically justified. Just like you cannot simply opt out of the law by simply not caring about it or by not playing the 'legal game', you cannot simply escape or opt out of epistemology by not playing the 'epistemology game' or by not caring about things like truth, knowledge, and evidence. So proponents of Epistemic Non-Normativity can very well accommodate the inescapability of epistemic norms.

¹⁰⁴ This is one of the main lessons from Foot (1972). See also Joyce (2001) and Olson (2014).

One might object that the datum to explain is not only that epistemic norms are inescapable, but that they are inescapably normative like, say, moral norms. This reply is question-begging however. To assert that epistemic norms are inescapably normative just is to assert that epistemic norms have necessary normative authority, which is precisely what is at issue. A better reply however would be that epistemic norms do seem inescapably normative because they plausibly imply *epistemic reasons*. This leads me to a third feature that Epistemic Non-Normativity needs to explain.

7.3.3 Norm-relativity and epistemic reasons

Just like moral facts imply moral reasons, epistemic facts plausibly imply epistemic reasons. If ϕ -ing is morally right or wrong, then there is a moral reason to ϕ or not to ϕ . In the same way, if believing that P is epistemically justified or unjustified, then there is an epistemic reason to believe that P or not to believe that P. So just like moral norms seem to be a source of moral reasons, epistemic norms seem to be a source of epistemic reasons. Moreover, just like moral reasons, epistemic reasons seem to be inescapable or utterly independent of what we care about. For example, consider the question whether you epistemically should believe that the number of stars is even. Whether or not you care about believing the truth, avoiding error, and so on, there is an epistemic reason for you *not* to believe that the number of stars is even and instead suspend judgment about that matter. How can Epistemic Non-Normativity explain this second feature of epistemic facts?

It is true that just like moral facts imply moral reasons, epistemic facts imply epistemic reasons. But similarly, as I already mentioned in chapter 4, legal facts imply legal reasons, etiquette facts imply etiquette reasons, tradition facts imply traditional reasons, and so on. As we have already seen, terms like ‘should’, ‘ought’, ‘good’, ‘correct’, and so on can be used to convey mere norm-relativity and not genuine normativity. For example, even if there is not necessarily a normative reason to conform to the law, illegal acts are still necessarily *legally* forbidden, bad, incorrect, and so on. The point is that these terms can be used to merely convey that ϕ -ing is required, good, correct, and so on, *relative to* or *according to* a particular norm N. For all that these mere norm-implying or norm-relative should-claims, goodness-claims, and correctness-claims imply, there might be no normative reason to ϕ . In the same way, although reasons-claims are sometimes used – especially by philosophers – to

convey genuine normativity, they too can very well be read in a merely norm-relative or norm-implying way. That is, they can very well be used to merely convey that there is a reason to ϕ *relative to* or *according to* N. For all these mere norm-implying reasons-claims imply, there might be only norm-relative reasons to ϕ and no *normative* reason to ϕ . As Jonas Olson explains:

‘[O]ught’ and ‘reason’ are both polysemous terms. [...] For example, it seems perfectly fine to say that according to Catholicism there is *reason* to abstain from eating meat on Fridays and to mean by this simply that it is incorrect according to the rules of Catholicism to eat meat on Fridays. Other examples of the same kind are not difficult to find. Consider grammar or etiquette. The claim that there are reasons for writers in English not to split the infinitive might simply mean that splitting the infinitive is inappropriate according to (some) rules of grammar; the claim that there are reasons for male guests to wear a tie at formal dinners might simply mean that this is required by the rules of etiquette. Or consider chess or football. The claim that chess players have reasons not to move the rook diagonally might simply mean that this is incorrect according to the rules of chess; the claim that football players have reasons not to play the ball to their own goalkeeper while under pressure might simply mean that doing so is likely to provide the opponent team with opportunities to score. To give a final example, the claim that soldiers have reasons to comply with the orders of a general might simply mean that doing so is to comply with the orders of someone superior in military rank, which is part of the role of being a soldier. (Olson 2014, 119-120)

Even when there is no *normative* reason to do as e.g. the law, fashion, and etiquette says, it remains true, trivially, that there are *legal reasons*, *fashion reasons*, and *etiquette reasons* to do so. It is just that in these situation such reasons lack genuine normative force or authority. These are situations, in other words, where legal reasons, fashion reasons, and etiquette reasons do not constitute genuinely normative reasons, but are only norm-relative reasons. There can very well be cases where there is no normative reason to do what there is e.g. legal reasons, etiquette reasons, and traditional reasons to do.

According to Epistemic Non-Normativity, epistemic reasons are just like that. While epistemic norms do imply epistemic reasons, they do so in the same trivial or norm-relative sense that e.g. legal norms imply legal reasons. Just like legal reasons, epistemic reasons do not necessarily constitute genuinely normative reasons or have normative force or authority. They are sometimes just mere norm-relative reasons.

There might be, in other words, no normative reason to do what there is epistemic reasons to do. Therefore, the sole claim that there is an epistemic reason to believe that P – just like the sole claim that there is a legal reason to ϕ – leaves open the further, normative question whether there is a genuinely normative reason to believe what there is epistemic reason to believe. Hence, the merely norm-implying character of epistemic facts and claims is perfectly compatible with their implying epistemic reasons.

This explanation also accommodates the *inescapability* of epistemic reasons. That is, it also entails, correctly, that epistemic reasons are independent from our desires or goals. When there is an epistemic reason for you to believe that P, this is so whether or not you care about believing the truth, avoiding error, and the like. Epistemic Non-Normativity delivers that result because mere norm-relative reasons can very well be inescapable. As Olson writes:

A soldier might not *desire* to comply with the general's order and he might have no other desires whose satisfaction would be promoted by his compliance, but he can still be said to have reasons to comply with the general's order since complying with the orders of those of superior military rank is part of the role of being a soldier. The same goes for chess players and football players; they might not desire to play by the rules and they need not even desire to win. Agents can occupy roles they have no desire to fulfil and engage in activities they have no desire to succeed in. (Olson 2014, 121)

Similarly, whether or not you care about conforming to e.g. the law, fashion, or tradition, there are legal reasons for you not to violate the law, fashion reasons not to violate fashion standards, and traditional reasons not to go against tradition. In the same way, there are epistemic reasons for you to conform to epistemic norms regardless of what you care about. It is just that just like legal reasons, fashion reasons, and traditional reasons, epistemic reasons do not necessarily have genuine normative force.

7.3.4 Norm-relativity and universality

I have been comparing epistemic norms with norms – such as the law, etiquette, and fashion – which are plausibly not *absolute* or *universal*, but rather *relative* and

conventional. When it comes to the latter kinds of norms, it seems that there is not one *code* or *system*, or *framework* that is absolutely ‘correct’ or ‘valid’. Rather there is a plurality of legal, etiquette, and fashion codes, which are all equally ‘valid’ or ‘correct’. It makes no sense, after all, to claim that something is illegal, unfashionable, or indecorous *period* or *absolutely*. Rather, any such claim must be relativized to a particular legal, etiquette, or fashion code or framework. This is because, plausibly, the law, etiquette, and fashion are not necessary facts that we discovered, but rather contingent – if not arbitrary – conventions that we came up with. Facts about what is e.g. legal, decorous, and fashionable, in other words, are – or are grounded in – conventional facts about cultural or social artefacts.

The problem is that for most philosophers, epistemic facts are nothing like this. Epistemic norms do not seem to be relative and conventional in that way. It rather seems that, much like e.g. moral claims, epistemic claims can be true absolutely and even perhaps necessarily. For most epistemologists, there is not a plurality of equally valid epistemic ‘codes’, but rather one absolute and universal system of epistemic norms. It does make sense, after all, to claim that a belief is epistemically justified, rational, or correct *period* or *absolutely*. Unlike e.g. legal claims, epistemic claims do not need to be relativized to a particular epistemic system to make sense. This is because – again, according to most philosophers – these epistemic norms are plausibly not just a matter of contingent or arbitrary conventions that we came up with. Epistemic facts are not mere conventional facts about cultural or social artefacts. Rather, they are discoverable facts that exist independently of such conventions.¹⁰⁵ How can Epistemic Non-Normativity explain this if, as it maintains, epistemic norms are akin to paradigmatically relative and conventional norms like etiquette, fashion, and the law?

Epistemic Non-Normativity’s answer is that norms can be absolute, nonconventional, and universal without having necessary normative authority. Merely norm-implicating claims, in other words, do not necessarily imply relative and conventional norms. They can very well imply absolute and universal norms. One way in which this can be the case is if the norms in question are grounded in or derived from some fundamental *aim* or *goal* that is itself absolute and universal. But as we

¹⁰⁵ This is, of course, controversial. Several epistemologists defend forms of epistemic relativism or conventionalism. See for instance, Rorty (1981), Kusch (2002), Neta (2007), Hazlett (2013), Kerr and Carter (2015), and Carter (2016). See e.g. Boghossian (2006) and White (2007) for criticism.

have already seen, this is precisely how many epistemologists construe epistemic norms. As I explained in chapter 4, one popular conception of the architecture of epistemic norms is a *veritist* conception according to which, roughly, epistemic norms are the norms that derive from a fundamental goal of *truth*, i.e. the goal or standard of believing what is true and not what is false. If this is right and if, as most epistemologists maintain, truth is universal and not just a matter of convention, then there might very well be only one true system of epistemic norms, i.e. the one that derives from the goal of truth. Crucially however, none of this entails that this one epistemic system or framework has necessary normative authority. In particular, it does not mean that there is necessarily a normative reason to achieve the universal and absolute goal of truth. So even proponents of Epistemic Non-Normativity can maintain that epistemic norms are those that derive from the absolute and universal goal of truth and thus that there is not a plurality of equally valid epistemic systems.¹⁰⁶

This kind of story can also accommodate the non-conventional character of epistemic facts. Suppose epistemic facts are essentially facts about those norms that derive from the fundamental goal of truth. Whether some belief-forming process is conducive to the goal of believing the truth and avoiding error is plausibly not just a matter of contingent or arbitrary conventions that we came up with, but rather a matter of discoverable facts that exist independently of such conventions.

To be clear, my goal here is not to establish epistemic absolutism and reject epistemic relativism or conventionalism. Instead, I only aim to show that Epistemic Non-Normativity is not committed to epistemic relativism. It can accommodate absolutism by invoking stories – like veritism – that are already popular among epistemologists. Nor am I trying to defend or formulate veritism. For one thing, the sketch I offered above overlooks the disagreements regarding how to formulate the truth goal exactly. For another, the fundamental epistemic goal might not be truth, but rather e.g. knowledge.¹⁰⁷

¹⁰⁶ Although most veritists accept Epistemic Normativity, some epistemologists appear to be sympathetic to the story sketched here. This is roughly how Grimm (2009) reads Sosa (2007). Some proponents of naturalized epistemology such as Quine (1969), (1992), Maffie (1990), and Laudan (1990a), (1990b), also seem to invoke versions of this approach in response to the kind of objection from epistemic normativity raised by Sellars (1956) and Kim (1988). See e.g. Wrenn (2006) for discussion.

¹⁰⁷ See e.g. Williamson (2000), (Forthcoming), Bird (2007), Hattiangadi (2010), McHugh (2011), Littlejohn (2013), (Forthcoming b).

At this point, one might object that there is still something arbitrary or even conventional with this kind of story. Why are epistemic norms those that derive from the goal of truth? Why do we evaluate beliefs according to this fundamental epistemic goal in the first place? Isn't *that* conventional or arbitrary?

Not necessarily. Another popular thesis that even proponents of Epistemic Non-Normativity can invoke is what I called in previous chapters the *aim of belief* thesis. Recall that according to that thesis, part of what it is for something to be a belief is for it to be, in some sense, directed or regulated towards being true and not false. It is partly constitutive of what beliefs are, in other words, that they aim at the truth. Crucially, this purports to be a necessary truth. Necessarily, if a state or propositional attitude is a belief, then it is the kind of state that aims at truth in some sense. Therefore, if the aim of belief thesis is true, then epistemic norms and the fundamental epistemic goal may have a clear non-arbitrary, non-conventional, and even necessary source or ground, namely the essential nature of belief, i.e. the necessary fact that belief constitutively aims at truth.¹⁰⁸ The epistemic norms of belief are just those, in other words, that derive from the constitutive aim of belief.

One might object that this kind of *constitivist* story is not available to proponents of Epistemic Non-Normativity or that it is self-defeating because it seems to establish Epistemic Normativity.¹⁰⁹ According to some philosophers, it would be sufficient to establish the necessary normative authority of certain norms if we could show that these norms derive from an aim that agents cannot possibly escape or avoid. This is what we might call *normative constitutivism*.¹¹⁰ But if epistemic norms derive from the constitutive aim of belief, then it seems that normative constitutivism could be applied to Epistemic Normativity since beliefs are plausibly inescapable for agents. Being an agent, in other words, seems to constitutively involve forming and revising beliefs. But if belief constitutively aims at truth and if epistemic norms derive

¹⁰⁸ An alternative possible story is that epistemic norms are just those that derive from the constitutive aim of *inquiry* or of *epistemic agency*.

¹⁰⁹ The rest of this section draws in part from Côté-Bouchard (Forthcoming).

¹¹⁰ Normative constitutivism is especially popular in the practical domain. Many claim, for instance, that we can ground practical or even moral normativity in what is constitutive of action or agency. See for instance, Velleman (2000), Wallace (2001), Rosati (2003), O'Hagan (2004), (2005), Ferrero (2009), Korsgaard (1996), (2008), (2009), Tubert (2010), (2011), Berdea (2013), and Katsafanas (2013), (Forthcoming). See Enoch (2006) for criticism.

from that aim, doesn't it follow that we cannot possibly escape the aim of belief and therefore that they necessarily have normative force for us?¹¹¹

As I have argued elsewhere, such an epistemic or doxastic version of normative constitutivism is a non-starter.¹¹² No plausible version of the aim of belief thesis can ground the necessary normative authority of epistemic norms, i.e. Epistemic Normativity. First, as I already argued in chapter 3, the aim of belief cannot be interpreted as an aim that agents necessarily *have* or as something that agents necessarily want. A more plausible option is to construe it as a metaphor for the claim that a belief is good, well-functioning, correct, or fitting qua belief if and only if it is true. But as I argued in chapter 2, this thesis is not sufficient to ground genuine normativity. The sole fact that something like an activity, state, or role X has an essential function or correctness condition does not entail that there is any normative reason to meet that condition. This claim alone, as we have seen, is silent on whether there is any normative reason to be a good, well-function, fitting, or correct X, and to avoid being a bad, defective, or incorrect X.

However, as the story from the previous paragraph indicates, even if they do not entail genuine normativity, constitutive aims, functions, or correctness conditions can still ground *norms*. That is, even if they cannot ground genuine normative reasons, they still provide or set standards relative to which we can evaluate people and their behaviour. So even if *normative* constitutivism is implausible, it does not follow that *norms* constitutivism – i.e. constitutivist explanations of *norms* but not of normativity – is also implausible. Hence, proponents of Epistemic Non-Normativity can very well adopt some form of norms constitutivism in the case of epistemic norms. That is, one can very well claim that epistemic norms are those that derive from the constitutive aim of belief and at the same time maintain that such norms lack necessary normative authority since there is not necessarily a normative reason to achieve the aim of belief.

7.4 Summary

In this chapter, I asked how Epistemic Normativity and Epistemic Non-Normativity can answer two explanatory questions or challenges. The first, which is the one I am

¹¹¹ See for instance Velleman (2000), Wedgwood (2002), Shah and Velleman (2005), O'Hagan (2005), Steglich-Petersen (2009), Tubert (2010), Wiland (2012, chapter 6), and Nolfi (2013).

¹¹² Côté-Bouchard (Forthcoming). See also Railton (1997) and Papineau (2013).

mainly concerned with in this thesis, is to explain the conclusions of part I, i.e. that the five commonly cited marks of normativity from chapters 2-6 are missing in epistemic facts and claims.

On the one hand, I have argued that the most promising explanation for Epistemic Normativity is what I called Normative Pluralism. According to the pluralist explanation, epistemic facts and claims do not bear the five commonly cited marks of normativity from part I because (i) there is a plurality of kinds of normativity and (ii) the kind of normativity implied by epistemic facts and claims – i.e. epistemic normativity – lacks these five features. According to Epistemic Non-Normativity, on the other hand, the five marks of normativity from part I are missing in epistemic facts and claims because (i) epistemic facts and claims are not normative, but rather (ii) merely norm-implying and the five features from part I are not essential features of merely norm-implying facts and claims.

The second explanatory challenge was to explain why epistemic facts and claims also have other features which one might see as additional hallmarks of normativity. I have argued that Epistemic Non-Normativity can meet that challenge because these features are not exclusive to normative facts and claims. Even if they are norm-implying, epistemic facts can still be important, inescapable, a source of epistemic reasons, and universal or absolute.

In the next chapter, I return to the first explanatory question and evaluate the two rival explanations on the basis of commonly invoked explanatory virtues or criteria. I will argue that given these criteria, Epistemic Non-Normativity is the best explanation of the conclusions of chapters 2-6.

Chapter 8: Epistemic non-normativity as the best explanation

In the previous chapter, I formulated the explanations that Epistemic Normativity and Epistemic Non-Normativity can each give of the conclusions of part I. In particular, I argued that the two rival explanations are the following:

Normative pluralism: *Epistemic facts and claims do not bear the five commonly cited marks of normativity from part I because (i) there is a plurality of kinds of normativity and (ii) the kind of normativity implied by epistemic facts and claims – i.e. epistemic normativity – lacks these five features.*

Epistemic Non-Normativity explanation: *the five marks of normativity from part I are missing in epistemic facts and claims because (i) epistemic facts and claims are not normative, but rather (ii) merely norm-implying and the five features from part I are not essential features of merely norm-implying facts and claims.*

In this final chapter, I argue that Epistemic Non-Normativity's explanation is the best explanation of the fact that the features from chapters 2-6 are missing in epistemic facts and claims. More precisely, I argue that given the most commonly invoked explanatory virtues or criteria for theory choice, Epistemic Non-Normativity clearly comes out as best. In particular, in comparison with Epistemic Non-Normativity, Normative Pluralism has either less or no more *ontological simplicity, explanatory simplicity, breadth, depth, coherence with background knowledge, intrinsic plausibility, and avoidance of ad hoc elements.*¹¹³ I therefore conclude that Epistemic Non-Normativity is, on the whole, the best explanation for the conclusions of part I. I consider the simplicity criteria in 8.1 and the other criteria in 8.2.

8.1 Simplicity

The most commonly invoked explanatory criterion or virtue is *simplicity*. It is standardly claimed that simpler theories, hypotheses, or explanations are preferable other things being equal. In this first section, I distinguish various ways in which

¹¹³ The list of explanatory criteria or virtues I use below is borrowed in large part from Beebe (2009).

theories can be more or less simple and argue that Epistemic Non-Normativity is, on the whole, simpler than Normative Pluralism. In particular, I argue that the former is both *ontologically* (or *semantically*) and *explanatorily* (or *syntactically*) simpler than the latter.¹¹⁴

8.1.1 Ontological simplicity

The first facet of the simplicity criterion has to do with what theories or explanations postulate as part of our ontology. Common labels for this first explanatory virtue include *ontological* simplicity, *semantic* simplicity, and *parsimony*. Beebe (2009, 609) distinguishes between two ways in which a theory can be more or less parsimonious or ontologically simple. The first has to do with the *kinds* of entities that theories postulate:

Ontological Simplicity I: *Other things being equal, a theory that postulates the existence of fewer kinds of entities should be preferred to a theory that postulates more.*

Normative Pluralism postulates the existence of more kinds of entities than Epistemic Non-Normativity.

Not only does Pluralism postulate the existence of normative facts, it also postulates the existence of at least two kinds of normative facts. In comparison, Epistemic Non-Normativity does not postulate the existence of normative facts at all. To the contrary, it is perfectly compatible with there being no such things. It is also compatible with there being only one kind of normative facts or normativity. The only kind of facts it postulates are merely norm-implying facts. Crucially however, Normative Pluralism is also committed to the existence of merely norm-implying facts such as legal facts, etiquette facts, and fashion facts. Hence, Pluralism is less parsimonious than Epistemic Non-Normativity when it comes to how many kinds or types of entities it postulates.

¹¹⁴ Beebe (2009) includes a third, *psychological*, facet of the simplicity criterion according to which, roughly, theories that are easier to understand are preferable other things being equal. I do not include this third criterion since it is less commonly invoked and because Epistemic Non-Normativity and Normative Pluralism are, on the face of it, equally easy to understand.

The second way in which an explanation can be more or less parsimonious has to do not with kinds of entities, but rather with the quantity of *individual* entities postulated:

Ontological Simplicity II: *Other things being equal, a theory that postulates the existence of fewer entities should be preferred to a theory that postulates more.*

Epistemic Non-Normativity postulates the existence of fewer individual entities than Normative Pluralism.

The ontological commitments of both theories are plausibly the same when it comes to individual non-normative facts. Normative Pluralism does not postulate fewer non-normative entities than Epistemic Non-Normativity. As we have just seen however, Pluralism also postulates normative facts. Epistemic Non-Normativity, on the other hand, is only committed to the existence of non-normative facts. Therefore, the only ontological difference between the two explanations is that the former postulates normative facts in addition to the same quantity of non-normative facts. Of course, it might turn out that there really are normative facts, and Epistemic Non-Normativity is entirely compatible with that possibility. But even if we are already committed to the existence of normative facts, Epistemic Non-Normativity still means fewer individual *normative* entities than Normative Pluralism since unlike the latter, it does not postulate *epistemic* normative facts. Hence, Epistemic Non-Normativity postulates fewer individual entities than Normative Pluralism.

In response, one might suggest that Normative Pluralism does not *add* normative facts to the same quantity of non-normative facts, but rather *replaces* some non-normative facts with normative facts. One might reply, in other words, that Normative pluralism and Epistemic Non-Normativity both postulate epistemic facts, but that the former views them as normative facts and the latter views them as non-normative. Hence the total quantity of individual entities remains the same.

But this is mistaken. Theories that postulate normative facts still postulate norm-implying facts as well. Even if moral facts are normative facts, there are still also merely norm-implying facts about what we should or should not do according to moral norms. It is just that the latter are not genuine moral facts. There is more to moral facts, in other words, than mere norm-implying facts about what is required, permitted, or forbidden according to moral norms. Similarly, even if epistemic facts

are normative facts as Normative Pluralism maintains, there are still norm-implicating facts about what we should or should not do *relative* to epistemic norms. Normative Pluralism's claim is just that the latter are not the real epistemic facts. There is more to epistemic facts, according to Normative Pluralism (and Epistemic Normativity more generally), than mere norm-implicating facts about what is required, permitted, or forbidden according to epistemic norms. So the Pluralist point is not that the latter non-normative facts do not exist, but rather that they are not epistemic facts.

A second potential reply is that Normative Pluralism is not committed to more individual entities because although they are normative, epistemic facts are identical or analysable purely in terms of non-normative facts. That is, some form of *analytic naturalism* might be true in the case of epistemic facts. However, the main problem with this suggestion is that even if we admit that such an analytic naturalistic account of epistemic normativity is plausible, analytic naturalism is much less plausible in the case of other kinds of normativity such as *moral* normativity. If, as most metaethicists maintain, *moral* facts are not identifiable or analysable in terms of purely non-normative facts, then even if we accept a form of analytic naturalism in the case of epistemic normativity, Normative Pluralism will still entail more entities, namely other normative facts such as moral fact, which are not analysable in non-normative terms.

Perhaps there are better replies to my defence of Epistemic Non-Normativity's postulation of fewer individual entities. However, the best that my opponents can hope to show is that Normative Pluralism postulates *no more* individual entities than Epistemic Non-Normativity. But even if this were the case, the latter would still come out as the ontologically simpler theory on the whole since it would still postulate fewer *kinds* of entities. Hence, I conclude that Epistemic Non-Normativity is, on the whole, ontologically simpler or more parsimonious than Normative Pluralism.

8.1.2 Explanatory simplicity

The second main facet of the simplicity criterion concerns the theories or hypotheses or explanations themselves rather than the entities they postulate. This is often referred to as *explanatory* or *syntactic* simplicity. Once again, I follow Beebe (2009) and distinguish three ways in which explanations can be more or less explanatorily

simple. I argue given these versions of the criterion, Epistemic Non-Normativity is explanatorily simpler, on the whole, than Normative Pluralism.

The first way in which a theory or explanation can be more or less explanatorily simple has to do with how *elegant* or *straightforward* it is:

Explanatory Simplicity I: *Other things being equal, a theory whose structure is more elegant or straightforward should be preferred to a theory that is less elegant or straightforward.*

Epistemic Non-Normativity is at least as elegant as Normative Pluralism. Its explanation of the conclusions of part I could hardly be more straightforward since it simply takes these conclusions at face value. The reason why epistemic facts and claims do not bear these common marks of normative facts and claims is, at bottom, that they are simply not normative.

A second way in which explanations can be more or less syntactically simple concerns the *quantity of additional explanatory questions* that they raise:

Explanatory Simplicity II: *Other things being equal, a theory that raises fewer further explanatory questions should be preferred to a theory that raises more.*

Normative Pluralism clearly raises more additional explanatory questions than Epistemic Non-Normativity. Here are some of them.

Recall that according to Normative Pluralism, there is a plurality of kinds of normative facts, and at least one of them – i.e. the epistemic kind of normative facts – lacks the features from part I. But if they lack these features and if, as I argued in the previous chapter, merely norm-implying facts lack these features as well, then (i) why do epistemic facts still count as normative facts? What other feature do they have that other normative facts have and that norm-implying facts lack? And (ii) why does this distinctly epistemic kind of normativity lack these features that other kinds of normativity seem to have? Moreover, (iii) if the possession or absence of the five features from part I is not sufficient to distinguish normative facts from merely norm-implying facts, then what is? What is it that distinguishes normative facts from merely norm-implying ones? (iv) What is it that *grounds* or *explains* this distinctly epistemic kind of normativity? Why or in virtue of what, in other words, do epistemic norms

necessarily have normative authority? Where does their necessary normative force come from if, as I argued in chapters 2 and 3, it cannot come from facts about value or about our desires? (v) Since epistemic facts are not autonomous from non-normative ones – as shown in chapter 4 – what is the relation between those two kinds of facts? In particular, does it mean that the latter is reducible or analysable in terms of the latter? (vi) If so, why are they still genuinely normative? (vii) If they are still genuinely normative despite being analysable or reducible in terms of non-normative facts, does it mean that such reductive or analytic naturalism is also acceptable in the case of other kinds of normative facts such as moral facts? (viii) If it is, then how should we answer the standard objections to, say moral naturalism? If it is not, then why don't the standard objections to normative naturalism apply to epistemic naturalism as well? Finally, (ix) if epistemic facts are normative facts, then how can we respond to the metanormative objections like those outlined in chapter 1 and defend the existence of normative epistemic facts?

We do not need to answer these questions if we adopt Epistemic Non-Normativity however. Of course, Epistemic Non-Normativity also raises further explanatory questions, which I examined in chapter 7. As I explained however, these questions are not especially problematic. Epistemic Non-Normativity can easily accommodate the fact that epistemic facts are important, inescapable, sources of epistemic reasons, and absolute.

In contrast, the additional explanatory questions raised by Normative Pluralism are not only more numerous, they are also more problematic or difficult. For instance, if Normative Pluralism is true and epistemic facts have their own distinctive kind of normativity, then what are such normative epistemic facts like? They cannot plausibly be *non-natural* facts since the Autonomy Principle (AP) from chapter 4 is a central commitment or motivation of non-naturalist accounts of normative facts. As I explained above, it is a central tenet of normative non-naturalism that there is a fundamental gap between normative and natural non-normative facts. Even if the autonomy of normative facts from non-normative natural ones does not automatically entail the truth of normative non-naturalism, the *non*-autonomy of normative facts from non-normative facts does render non-naturalism implausible or at least unmotivated. If we can infer normative conclusions from natural non-normative facts alone, then it is unclear why such normative truths could not also be natural truths. Consequently, *epistemic* non-naturalism would plausibly require that

the Autonomy Principle be true in the case of epistemic facts and claims – i.e. that what I called the Epistemic Autonomy Thesis (EAT) be true. But since, as I argued in chapter 4, we *can* infer epistemic conclusions from purely natural non-normative premises – i.e. since EAT is false – epistemic facts are not plausibly natural facts and not non-natural facts. So if Normative Pluralism is correct, then epistemic facts are most plausibly *natural* normative facts.

However, this means that Normative Pluralism faces what I called the *Losing normativity* problem in chapter 1. The problem, in short, was that purely natural facts do not seem to be the kind of thing that can have normative authority. After all, in the case of *moral* normativity for instance, it seems that purely natural facts cannot settle first person moral deliberation and that natural-moral questions are always *open* or *substantial* questions. Why is it any different with the epistemic kind of normativity? Why can this kind of normativity be generated by natural facts while others cannot?

The question of the *source* or *grounds* of this epistemic kind of (natural) normativity is also highly problematic. Where does this distinctive kind of normative authority that epistemic facts have come from? *Why* or *in virtue of what* do epistemic norms have necessary normative authority? As we have already seen in chapters 2 and 3, the answer cannot be that epistemic normativity comes from facts about *value* or *desires*. As I have also explained in the previous chapter, it is not plausibly grounded in facts about the constitutive aim of belief. Hence, it is far from clear what kind of natural facts are supposed to ground or explain epistemic normativity.

One alternative option is that the normativity of epistemic facts is not grounded or explained by anything else. That is, perhaps there is nothing else in virtue of which epistemic facts have necessary normative authority. Perhaps it is just a *brute* or *primitive* fact that there is necessarily a normative reason to conform to epistemic norms. However, this would make Normative Pluralism much more ad hoc and much less illuminating than Epistemic Non-Normativity. It seems that such a view of epistemic normativity would do little more than simply positing a kind of normativity that simply fits the explanandum.¹¹⁵

Finally, why does the distinctly epistemic kind of normativity count as a kind of normativity and not mere norm-relativity? What does it have in common with

¹¹⁵ I return to the issue of ad hoc explanations and explanatory depth below.

other kinds of normativity that makes it count as a kind of normativity and not an instance of mere norm-relativity? The most natural answer is that just like, say, moral or prudential reasons, the epistemic reasons entailed by epistemic facts are a kind of normative reasons.¹¹⁶ However, given that epistemic non-naturalism is ruled out, Normative Pluralism must view these normative epistemic reasons as natural entities or relations. This raises at least three issues.

First, this seems to go against the dominant approach in contemporary normativity theory according to which the reasons relation or ‘counting in favour’ relation or favouring relation is the single fundamental or primitive normative unit.¹¹⁷ For one thing, many proponents of this approach embrace non-naturalism and do not view this relation as a natural phenomenon.¹¹⁸ For another, Normative Pluralism means that there is not one single and fundamental kind of favouring relation, but rather a plurality of kinds of ‘counting in favour’. In particular, we have the epistemic kind of favouring on the one hand, which is a natural phenomenon, and on the other we have non-epistemic kinds of favouring, which may or may not be natural.

Second, and relatedly, suppose naturalism turns out to be implausible in the case of non-epistemic favouring. Suppose, in other words, that e.g. practical ‘counting in favour’ is best seen as a non-natural relation. In virtue of what would they both count as normative relations? What would they have in common that would make them both normative? Relatedly, if the practical reasons relation cannot be natural, then why can it be natural in the epistemic case? Conversely, if that relation can be natural in the epistemic case, why can’t it be natural in the practical case?¹¹⁹

Finally, if we allow that epistemic reasons involve just another kind of ‘counting in favour’, why should we not also see e.g. legal reasons as involving yet a further, legal, kind of ‘counting in favour’? In virtue of what are certain kinds of favouring normative and others non-normative or merely norm-relative?

¹¹⁶ See e.g. Scanlon (1998), (2014), Raz (1999), (2014), Cuneo (2007), Schroeder (2007), Alvarez (2010), Skorupski (2010), Enoch (2011a), Parfit (2011), and Rowland (2013).

¹¹⁷ See e.g. Scanlon (1998), (2014), Raz (1999), (2014), Skorupski (2010), and Parfit (2011).

¹¹⁸ E.g. Scanlon (1998), (2014), Cuneo (2007), Enoch (2011a), and Parfit (2011).

¹¹⁹ An additional point worth mentioning is that if it turned out that (i) Normative Pluralism is true, (ii) that epistemic normativity is a natural phenomenon, but (iii) that moral normativity is a *non*-natural phenomenon, then the argument in favour of the Unitary stance – which I presented in chapter 1 – would also lose force. Since normative epistemic facts would be natural facts, they would escape many of the metaethical objections summarized in the first chapter. The metaethical objections would not constitute metanormative objections and their soundness would not entail that there are no epistemic facts after all.

So in sum, Normative Pluralism raises more additional explanatory questions and many of these additional questions seem more difficult and problematic than those raised by Epistemic Non-Normativity. Therefore, Pluralism is less explanatorily simple in this second respect.

The final way in which theories can be more or less syntactically simple has to do with the quantity of primitive *explanatory notions* that they posit:

Explanatory Simplicity III: *Other things being equal, a theory that posits fewer primitive explanatory notions should be preferred to one that posits more.*

Epistemic Non-Normativity does not posit any explanatory notion that Normative Pluralism does not posit. The main notion at work in the former explanation is that of a merely norm-implying fact or of the property of norm-relativity, to which Normative Pluralism is also committed as we have already seen. Therefore, Normative Pluralism posits at least as many primitive explanatory notions than Epistemic Non-Normativity.

8.1.3 Taking stock

In this first section, I compared the two rival explanations of the conclusion of part I relative to the criteria of ontological and explanatory simplicity. I argued that Epistemic Non-Normativity postulates fewer kinds of entities, postulates fewer individual entities, is no less straightforward, raises fewer additional explanatory questions, and posits no more primitive explanatory notions than Normative Pluralism. Hence, I conclude that the former is the simplest explanation, on the whole, of the fact that the commonly invoked marks of normativity from chapters 2-6 are missing in epistemic facts and claims.

8.2 Additional explanatory criteria

In this second section, I evaluate Epistemic Non-Normativity and Normative Pluralism in light of five additional explanatory criteria or virtues. These are (i) *explanatory breadth*, (ii) *explanatory depth*, (iii) *coherence with background knowledge*, (iv) *intrinsic plausibility*, and (v) *avoidance of ad hoc elements*. I argue that Epistemic Non-Normativity performs either better or no worse than Normative Pluralism relative

to these criteria. Hence, I conclude that just like simplicity, these additional criteria show, on the whole, that the former is superior to the latter.

The first additional explanatory criterion or virtue concerns how much a theory can explain. Beebe (2009, 609) summarizes this criterion as follows:

Explanatory Breadth: *Other things being equal, a theory that explains a wider range of phenomena should be preferred to a theory that explains a narrower range.*

Epistemic Non-Normativity explains just as many phenomena as Normative Pluralism. As I argued in chapter 7 it is not only able to explain the conclusions of part I, but also the fact that epistemic facts are important, inescapable, a source of epistemic reasons, and absolute or universal.

The second additional explanatory criterion has to do with how *deep* or *illuminating* an explanation is:

Explanatory Depth: *Other things being equal, a theory that provides a more illuminating explanation of the relevant data should be preferred to a theory that provides a less illuminating explanation*

Epistemic Non-Normativity arguably provides a more illuminating explanation of the conclusions of part I than Pluralism. Normative Pluralism's answer to part I consists essentially in positing a kind of normativity that fits the data, i.e. that lacks the features from chapters 2-6. But this does little to illuminate or further our understanding of epistemic and normative facts. For one thing, it does not tell us why epistemic facts and claims lack the features from part I. For another, it leaves us in the dark regarding what it takes for facts to be normative and not merely norm-implying.

Perhaps Pluralism could be developed further to provide a deeper explanation, but Epistemic Non-Normativity seems in a better position to illuminate the data from part I. It helps us understand why epistemic facts are as they are described in part I since it shows, among other things, that such facts and claims are part of a larger class of facts and claims (merely norm-implying facts and claims), which also lack the features from chapters 2-6.

The third additional explanatory virtue concerns how well an explanation coheres with background knowledge or widely accepted claims. As Beebe (2009, 610) puts it:

Coherence with Background Knowledge: *Other things being equal, a theory that fits better with other widely accepted theories and background knowledge should be preferred to a theory that fits less well.*

This criterion is trickier since it is hard to think of any widely accepted theories concerning the issues relevant for the debate between Epistemic Non-Normativity and Epistemic Normativity. Instead, most major questions in metaethics and normativity theory have several popular competing answers that each have prominent defenders. However, one crucial difference between Epistemic Non-Normativity and Normative Pluralism is that the former is compatible with a far greater number of prominent theories than the latter. Hence, Epistemic Non-Normativity might be said to be more coherent with background knowledge since for all we know, several of these popular positions could turn out to be true and unlike Pluralism, it is compatible with most of them.¹²⁰

For one thing, whereas Epistemic Non-Normativity is compatible with any account of the nature or source of normativity, Normative Pluralism rules out several such accounts. First, Pluralism is incompatible with value-based theories according to which all normativity is ultimately grounded in facts about value or goodness. This is ruled out by the fact that Pluralism must accommodate the conclusion of chapter 2, namely that it is not necessarily good to conform to epistemic norms. Second, it is incompatible with desire-based theories according to which all normativity is ultimately grounded in facts about our desires. This is ruled out by the fact that Pluralism must accommodate the conclusion of chapter 3, namely that conforming to epistemic norms does not necessarily promote one's desires. Third, it is incompatible with *non-naturalist* accounts of normativity according to which all normative facts are irreducible and unanalysable in terms of natural facts. This is because Pluralism must

¹²⁰ Understood in that sense, the coherence criterion overlaps with the criterion that is commonly called *conservatism*. According to Beebe (2009, 610), a theory is more conservative if it "results in a smaller change in one's overall view." Since, as I will argue, Epistemic Non-Normativity is compatible with far more prominent accounts of normativity and metaethical theories than Normative Pluralism, the former will plausibly result in smaller changes in people's overall view.

accommodate the non-autonomous character of epistemic conclusions from non-normative premises – i.e. the conclusion of chapter 4 – and as we have already seen, normative non-naturalism is not plausibly compatible with some normative facts being *non*-autonomous from non-normative natural claims. Fourth, Normative Pluralism is incompatible with *expressivist* accounts of normativity according to which all normative claims express non-cognitive motivating states. This is because Pluralism must accommodate the rejection – argued for in chapter 5 – of epistemic expressivism. Finally, Pluralism is also incompatible with the widely held view that normative ‘oughts’ imply ‘can’. This is because it needs to accommodate the conclusions of chapter 6, namely that epistemic claims do not imply ‘can’.

A second respect in which Epistemic Non-Normativity is more coherent with popular theories is that unlike Normative Pluralism, it is compatible with both realism and anti-realism about normative facts. Unlike Pluralism, Epistemic Non-Normativity does not postulate the existence of any normative facts since it simply claims that epistemic facts are merely norm-implying and not normative. So on the whole, Epistemic Non-Normativity appears to be more coherent with widely accepted theories than Normative Pluralism.

The fourth additional explanatory criterion is what Beebe (2009, 610) labels intrinsic plausibility:

Intrinsic Plausibility: *Other things being equal, a theory that is more intrinsically plausible should be preferred to a theory that less intrinsically plausible.*

Epistemic Non-Normativity is no less intrinsically plausible than Normative Pluralism. As I have shown in the previous chapter, it can not only account for the conclusions of chapters 2-6, but also for additional features of epistemic facts that one might associate with normative facts. Moreover, what I have said so far shows that there is nothing intrinsically implausible with the idea that epistemic facts are merely norm-implying and not normative. That there are merely norm-implying facts which lack the features from part I is something we already find intrinsically plausible. Therefore, it is far from implausible to conclude that epistemic facts are also merely norm-implying.

Finally, a fifth additional explanatory criterion has to do with how *ad hoc* theories or explanations are. Beebe (2009, 610) summarizes it as follows:

Avoidance of Ad Hoc Elements: *Other things being equal, a theory that has fewer ad hoc elements should be preferred to a theory that has more.*

As I already mentioned when discussing explanatory simplicity, Normative Pluralism is, on the face of it, a more ad hoc explanation of the conclusions of part I than Epistemic Non-Normativity. Normative Pluralism does little more than simply postulating a kind of normativity or normative facts that happen to fit the explanandum. Meanwhile, Epistemic Non-Normativity makes sense of the explanandum simply by invoking the already widely accepted fact that some facts are merely norm-implying and not genuinely normative. It does not look, in other words, like an explanation that was put together out of the blue just to fit the data.

8.3 Summary

In this chapter, I asked which of the two rival explanations outlined in chapter 7 – i.e. Epistemic Non-Normativity and Normative Pluralism – is the best explanation of the conclusions of part I. In order to answer this question, I evaluated the two explanations relative to various commonly invoked explanatory criteria. I argued that Epistemic Non-Normativity clearly comes out as the best explanation according to these criteria. More precisely, I argued that compared to Normative Pluralism, Epistemic Non-Normativity (i) postulates the existence of fewer kinds of entities, (ii) postulates fewer individual entities, (iii) is no less elegant or straightforward, (iv) raises fewer additional (difficult) explanatory questions, (v) posits no more primitive explanatory notions, (vi) explains just as many phenomena, (vii) is more illuminating, (viii) is more coherent with widely accepted theories, (ix) is no less intrinsically plausible, and (x) has fewer ad hoc elements. Since it performs better, on the whole, than Normative Pluralism according to these explanatory criteria, I conclude that Epistemic Non-Normativity is the best explanation of the fact that the commonly invoked marks of normativity chapters 2-6 are missing in epistemic facts and claims. Moreover, since it best explains the relevant data, I also conclude that in all likelihood, Epistemic Non-Normativity is true and Epistemic Normativity is false. Epistemic facts and claims are best seen not as genuinely normative, but rather as merely norm-implying.

Conclusion

Epistemology is widely seen as a normative discipline like ethics. Just like moral facts, in other words, epistemic facts – i.e. facts about our beliefs' epistemic justification, rationality, reasonableness, correctness, warrant, and the like – are *normative* facts. Yet, whereas countless metaethicists have rejected the existence of moral facts, few philosophers have raised similar objections to the existence of epistemic facts.

According to what I called the *Unitary stance*, this *Dual stance* regarding the vulnerability of moral and epistemic facts to metaethical objections is mistaken. If moral facts really are objectionable, queer, mysterious, and so on, then so are epistemic facts. This is because, as recent developments in metaethics and normativity theory made clear, those metaethical objections really are *metanormative* objections. That is, arguments against the existence of moral facts really are arguments against normative facts more generally. It is only because of their necessary or inherent normative character that moral facts are seen as objectionable in the first place. But since epistemic facts are normative like moral facts, their existence is no less threatened by metaethical objections. This is what I called the *argument from Epistemic Normativity* in favour of the Unitary stance and against the Dual stance:

The argument from Epistemic Normativity

1. If any of the metaethical objections is sound, then there are no normative facts; metaethical objections really are metanormative objections. (*Metanormative thesis*)
2. Epistemic facts and claims are normative facts and claims. (*Epistemic Normativity*)
3. Therefore, if any of the metaethical objections is sound, then there are no epistemic facts. (*Unitary stance*)

In this thesis, I offered a response to this argument. I argued that we should reject its second premise and adopt instead what I called Epistemic Non-Normativity. Contrary to what is commonly claimed or assumed, epistemic facts and claims are best seen not as normative like e.g. moral facts, but rather as merely norm-implying like legal facts, etiquette facts, and the like. Epistemic facts imply norms – epistemic norms – that lack necessary normative authority. Unlike e.g. moral norms, epistemic norms

do not necessarily entail or provide normative reasons. Instead, just like e.g. legal norms, etiquette norms and the like, there is not necessarily a normative reason to conform to them.

I offered the following two-part abductive argument in favour of Epistemic Non-Normativity and against Epistemic Normativity:

The abductive argument from the marks of normativity

1. Epistemic facts and claims lack five commonly cited marks of normativity, i.e. five features that merely norm-implying facts and claims lack, but which are commonly attributed to normative facts and claims. (PART I)
2. The best explanation of premise 1 is that Epistemic Non-Normativity is true, i.e. that epistemic facts and claims are merely norm-implying and not normative. (PART II)
3. Therefore, in all likelihood, Epistemic Non-Normativity is true and Epistemic Normativity is false.

I defended the first premise in part I. The five commonly cited marks of normativity that are missing in merely norm-implying facts are (i) a necessary connection with *value*, (ii) a necessary connection with *desire*, (iii) a necessary autonomy from the non-normativity, (iv) a necessary connection with *motivation*, and (v) a necessary connection with *control*.

In chapter 2, I argued that epistemic facts and claims do not necessarily imply value since it is not necessarily genuinely good to conform to epistemic norms. In chapter 3, I argued that epistemic facts and claims do not necessarily imply the promotion of one's desire satisfaction since conforming to epistemic norms does not necessarily promote one's desires. In chapter 4, I argued that epistemic facts and claims are not necessarily autonomous from the non-normative since we can derive epistemic conclusions from non-normative premises alone. In chapter 5, I argued that epistemic facts and claims are not necessarily connected with motivation since epistemic expressivism is implausible and there is no plausible sense of being motivated in which epistemic judgments are necessarily motivating. Finally, in chapter 6, I argued that epistemic facts and claims do not imply control or 'can'. Hence, just

like merely norm-implicating facts and claims, epistemic facts and claims do not bear any of these five commonly cited features of normativity.

I defended the second premise of the abductive argument in the two chapters of part II. First, in chapter 7, I presented how Epistemic Non-Normativity and Epistemic Normativity can each explain the fact that epistemic facts and claims lack the five commonly cited marks of normativity from part I. On the one hand, I argued that Epistemic Normativity must give what I called a *Pluralist* explanation of part I. According to such a Normative Pluralism, epistemic facts and claims do not bear the five commonly cited marks of normativity from part I because (i) there is a plurality of kinds of normativity and (ii) the kind of normativity implied by epistemic facts and claims – i.e. epistemic normativity – lacks these five features.

On the other hand, Epistemic Non-Normativity's explanation of the results of part I is that epistemic facts and claims are merely norm-implicating (and not normative) and that the five features from part I are not features of merely norm-implicating facts. That is, even if they are marks of normativity, these features are not marks of mere norm-relativity. I ended chapter 7 by explaining how Epistemic Non-Normativity's explanation can also accommodate additional characteristics of epistemic facts and claims that might be seen as further features of normative facts and claims. I argued that these additional features are not exclusive to normative facts and claims. Merely norm-implicating facts can also be (i) important, (ii) inescapable, (iii) sources of epistemic reasons, and (iv) universal or absolute.

Finally, in chapter 8, I argued that given a variety of commonly invoked explanatory virtues or criteria, Epistemic Non-Normativity clearly comes out as the best explanation of the conclusions of part I. In particular, I argued that Epistemic Non-Normativity fares better than Normative pluralism relative to criteria of ontological simplicity, explanatorily simplicity, depth, coherence with background knowledge, and avoidance of ad hoc elements. I also argued that Epistemic Non-Normativity has no less breadth and intrinsic plausibility than Normative pluralism. Hence, the best explanation of the relevant data is, on the whole, that epistemic facts and claims are not genuinely normative, but rather merely norm-implicating. I therefore conclude that Epistemic Non-Normativity is most probably true. Unlike e.g. moral facts and claims, epistemic facts and claims are most plausibly seen not as genuinely normative, but rather as merely norm-implicating.

This means, in turn, that we should also reject the above argument in favour of the Unitary stance and against the Dual stance. Even if metaethical objections really are metanormative objections, it does not follow that epistemic facts are vulnerable to them since such facts are *not* best seen as normative facts.

This conclusion is likely to have important consequences for epistemology and metaethics. Epistemologists and metaethicists often invoke the alleged normative character of epistemology in contemporary debates. But if Epistemic Non-Normativity is true and epistemic facts are not genuinely normative, then arguments that rely on Epistemic Normativity lose force. I will end by giving three examples of debates where my conclusion is likely to make a difference.

The first is the dispute concerning the existence of moral facts. As I mentioned in chapter 1, metaethicists like Terence Cuneo have recently defended moral realism by invoking epistemic facts as moral facts' 'companions in guilt'.¹²¹ According to him, given Epistemic Normativity and the Unitary stance, the inexistence of moral facts would mean that there are no epistemic facts either. But since, he argues, epistemic facts exist, moral facts must exist as well.

Such a defence of moral realism loses force, however, if epistemic facts are merely norm-implicating and not genuinely normative. As I explained in chapter 1, it is precisely because epistemic facts are normative that, according to authors like Cuneo, such facts are targeted by metaethical objections like moral facts. But if Epistemic Non-Normativity is true, then epistemology won't help moral realism since the existence of epistemic facts might be due to their escaping metaethical objections by being merely norm-implicating and not normative.

Second, my conclusion also matters for the question whether epistemology can be 'naturalized'. In "Epistemology Naturalized" Quine argues that we should treat epistemic phenomena like knowledge as *empirical* phenomena to be studied by science. Epistemology, according to this approach, "[...] falls into place as a chapter of psychology." (1969, 82) One of the most prominent objections against this idea is that it cannot accommodate the normative character of epistemology. Quine's conception of epistemology as a naturalized scientific discipline seems to make it devoid of normativity. After all, paradigmatically normative disciplines like ethics are

¹²¹ Cuneo (2007). See also Stratton-Lake (2002) and Rowland (2013).

clearly *not* chapters of psychology or any other science. Science alone plausibly cannot tell us what is morally right or wrong. But as Jaegwon Kim writes:

[...] it is difficult to see how an “epistemology” that has been purged of normativity, one that lacks an appropriate normative concept of justification or evidence, can have anything to do with the concerns of traditional epistemology. And unless naturalized epistemology and classical epistemology share some of their central concerns, it's difficult to see how one could replace the other, or be a way (a better way) of doing the other [...]. For epistemology to go out of the business of justification is for it to go out of business. (Kim 1988, 391)

Similarly, according to Wilfrid Sellars:

The essential point is that in characterizing an episode or state as that of *knowing*, we are not giving an empirical description of that episode or state; we are placing it in the logical space of reasons, of justifying and being able to justify what one says. (Sellars 1956, Sect. 36)

Their thought is that since epistemic notions like epistemic justification are normative, adopting Quine's normativity-free approach would mean abandoning epistemic notions and epistemology altogether.¹²²

However, this argument loses force if Epistemic Non-Normativity is true. Making epistemology normativity-free only entails giving up epistemology if Epistemic Normativity is true in the first place. If epistemic facts and claims are not normative, then the sole fact that Quine views epistemology as normativity-free does not entail that his naturalism is also epistemology-free. So although my conclusion does not mean that Quine's approach is correct, it casts serious doubt on Sellars and Kim's objection to it.

Finally, a third issue where the rejection of Epistemic Normativity is likely to matter is the debate between internalists and externalists about epistemic justification. A common way to draw the distinction between internalism and externalism is that the former, but not the latter, puts an awareness or accessibility condition on epistemic justification. For internalists, S can only have an epistemically

¹²² See e.g. Wrenn (2006) for discussion.

justified belief if she is aware (or could become aware) of what justifies that belief. Externalists deny that.

Importantly however, that debate is often thought to go hand in hand with the choice between theories that construe epistemic justification as a deontic notion having to do with 'ought' and those that do not.¹²³ This is because many think that there is an internal condition on the truth of deontic 'ought' claims. It can only be the case that you ought to ϕ , the thought goes, if you are aware (or could be aware) that you ought to ϕ . If this is right, then the truth of the deontic construal would be problematic for externalism since it puts no such awareness or accessibility condition on epistemic justification. Similarly, the falsity of the deontic construal would be problematic for internalism since it would rob it of one of its main motivations and would open the door for construals that favour externalism.

However, even if we accept that there is an awareness or accessibility requirement on deontic norms or 'oughts', such a condition does not apply or extend to merely norm-implying norms. Deontic 'oughts' that lack necessary normative authority can very well apply to agents in cases where they cannot be aware of them. For example, if ϕ -ing is illegal, then I legally ought not to ϕ whether or not I can become aware of that law. My lacking cognitive access to that norm does not make my ϕ -ing any less illegal or legally forbidden. Similarly, if ϕ -ing is forbidden by a particular code of etiquette, then my ϕ -ing is a violation of that code even if there is no way for me to become aware of that norm. Even if I lack such access, ϕ -ing will still count as something I ought not to do according to that code of etiquette. So when it comes to merely norm-implying 'oughts' like legal or etiquette 'oughts', cognitive access to them is not necessary for being subject to them.

But if Epistemic Non-Normativity is true, then even if epistemic norms or 'oughts' are deontic, they are merely norm-implying just like etiquette and the law. But since merely norm-implying deontic 'oughts' do not necessarily have an awareness or accessibility condition, it follows that the issue whether epistemic justification is deontic won't settle the dispute between epistemic internalism and externalism.

¹²³ See e.g. Bonjour (1985), Steup (1988), Feldman (1988), Plantinga (1993), Goldman (1999b), and Bergmann (2006).

As these examples show, the alleged normativity of epistemology plays an important role in contemporary epistemological and metaethical disputes. But since most contemporary philosophers have defended or assumed Epistemic Normativity, the conclusion reached in this thesis is likely to open up new avenues in these debates.

Bibliography

- Adler, J. 2002. *Belief's Own Ethics* MIT Press.
- Ahlstrom-Vij, Kristoffer. 2013a. "In Defense of Veritistic Value Monism." *Pacific Philosophical Quarterly* 94 (1): 19-40.
- . 2013b. "Moderate Epistemic Expressivism." *Philosophical Studies* 163 (2): 337-357.
- Almotahari, Mahrad and Adam Hosein. 2015. "Is Anything just Plain Good?" *Philosophical Studies* 172 (6): 1485-1508.
- Alston, William. P. 1988. "The Deontological Conception of Epistemic Justification." *Philosophical Perspectives* 2: 257-299.
- . 1989. *Epistemic Justification* Cornell University Press.
- . 1993. "Epistemic Desiderata." *Philosophy and Phenomenological Research* 53 (3): 527-551.
- . 2005. *Beyond Justification: Dimensions of Epistemic Evaluation* Ithaca: Cornell University Press.
- Alvarez, Maria. 2010. *Kinds of Reasons: An Essay in the Philosophy of Action* Oxford University Press.
- Anscombe, G. E. M. 1957. *Intention* Harvard University Press.
- Ayer, A. J. 1936. *Language, Truth and Logic* London, V. Gollancz, Ltd.
- Bandura, Albert. 1989. "Human Agency in Social Cognitive Theory." *American Psychologist* 44 (9): 1175.
- Bedke, Matthew. 2009. "Intuitive Non-Naturalism Meets Cosmic Coincidence." *Pacific Philosophical Quarterly* 90 (2): 188-209.
- . 2010. "Might all Normativity be Queer?" *Australasian Journal of Philosophy* 88 (1): 41-58.
- . 2014. "A Menagerie of Duties? Normative Judgments are Not Beliefs about Non-Natural Properties." *American Philosophical Quarterly* 51 (3): 189-201.
- Beebe, James R. 2009. "The Abductivist Reply to Skepticism." *Philosophy and Phenomenological Research* 79 (3): 605-636.
- . 2007. "BonJour's Abductivist Reply to Skepticism." *Philosophia* 35 (2): 181-196.

- Bergmann, Michael. 2006. *Justification without Awareness: A Defense of Epistemic Externalism* Oxford University Press.
- Berteau, Stefano. 2013. "Constitutivism and Normativity: A Qualified Defence." *Philosophical Explorations* 16 (1): 81-95.
- Bird, Alexander. 2007. "Justified Judging." *Philosophy and Phenomenological Research* 74 (1): 81-110.
- Blackburn, Simon. 1971. "Moral Realism." In *Morality and Moral Reasoning*, edited by John Casey: Methuen.
- . 1984. *Spreading the Word* Clarendon Press.
- . 1993. *Essays in Quasi-Realism* Oxford University Press.
- . 1996. "Securing the Nots: Moral Epistemology for the Quasi-Realist." In *Moral Knowledge? New Readings in Moral Epistemology*, edited by Walter Sinnott-Armstrong and Mark Timmons, 82-100: Oxford University Press.
- . 1998. *Ruling Passions* Oxford University Press.
- Bloomfield, Paul. 2001. *Moral Reality* Oxford University Press.
- Boghossian, Paul. 2006. *Fear of Knowledge: Against Relativism and Constructivism* Oxford University Press UK.
- . 2008. *Content and Justification: Philosophical Papers* OUP Oxford.
- BonJour, Laurence. 1985. *The Structure of Empirical Knowledge* Harvard University Press.
- Bortolotti, Lisa. 2014. *Irrationality* Polity Press.
- Boyd, Richard. 1988. "How to be a Moral Realist." In *Essays on Moral Realism*, edited by G. Sayre-McCord, 181-228: Cornell University Press.
- Brink, David Owen. 1989. *Moral Realism and the Foundations of Ethics* Cambridge University Press.
- Broome, John. 2007. "Is Rationality Normative?" *Disputatio* 2 (23): 161-178.
- . 2013. *Rationality through Reasoning* Wiley-Blackwell.
- Brunero, John. 2010. "The Scope of Rational Requirements." *The Philosophical Quarterly* 60 (238): 28-49.
- Burge, Tyler. 2003. "Perceptual Entitlement." *Philosophy and Phenomenological Research* 67 (3): 503-548.

- . 2010. *Origins of Objectivity* OUP Oxford.
- Bykvist, Krister and Anandi Hattiangadi. 2007. "Does Thought Imply Ought?" *Analysis* 67 (296): 277-285.
- Côté-Bouchard, Charles. 2015. "Epistemic Instrumentalism and the Too Few Reasons Objection." *International Journal of Philosophical Studies* 23 (3): 337-355.
- . Forthcoming. "Can the Aim of Belief Ground Epistemic Normativity?" *Philosophical Studies*.
- Carter, J. Adam. 2011. "Kvanvig on Pointless Truths and the Cognitive Ideal." *Acta Analytica* 26 (3): 285-293.
- . 2016. *Metaepistemology and Relativism* Springer.
- Carter, J. Adam and Matthew Chrisman. 2012. "Is Epistemic Expressivism Incompatible with Inquiry?" *Philosophical Studies* 159 (3): 323-339.
- Chrisman, Matthew. 2007. "From Epistemic Contextualism to Epistemic Expressivism." *Philosophical Studies* 135 (2): 225-254.
- . 2008. "Ought to Believe." *Journal of Philosophy* 105 (7): 346-370.
- . 2012. "Epistemic Expressivism." *Philosophy Compass* 7 (2): 118-126.
- . 2015. *The Meaning of 'Ought': Beyond Descriptivism and Expressivism in Metaethics* Oxford University Press Usa.
- Chuard, Philippe and Nicholas Southwood. 2009. "Epistemic Norms without Voluntary Control." *Noûs* 43 (4): 599-632.
- Clifford, William Kingdon. 1999. *The Ethics of Belief and Other Essays* Prometheus Books.
- Conee, Earl and Richard Feldman. 2004. *Evidentialism* Oxford University Press.
- Cowie, C. 2014. "Why Companions in Guilt Arguments Won't Work." *Philosophical Quarterly* 64 (256): 407-422.
- . 2015. "Good News for Moral Error Theorists: A Master Argument Against Companions in Guilt Strategies." *Australasian Journal of Philosophy* 94 (1): 115-130.
- Cuneo, Terence. 2007. *The Normative Web: An Argument for Moral Realism* Oxford University Press.
- Cuneo, Terence and Russ Shafer-Landau. 2014. "The Moral Fixed Points: New Directions for Moral Nonnaturalism." *Philosophical Studies* 171 (3): 399-443.

- Dancy, Jonathan. 2006. "Nonnaturalism." In *The Oxford Handbook of Ethical Theory*, edited by David Copp. Oxford: Oxford University Press.
- Darwall, Stephen, Allan Gibbard, and Peter Railton. 1992. "Toward Fin De Siècle Ethics: Some Trends." *Philosophical Review* 101 (1): 115-189.
- David, Marian. 2001. "Truth as the Epistemic Goal." In *Knowledge, Truth, and Duty*, edited by M. Steup, 151-169: New York: Oxford University Press.
- . 2005a. "Truth as the Primary Epistemic Goal: A Working Hypothesis." In *Contemporary Debates in Epistemology*, edited by M. Steup and E. Sosa : Malden, MA: Blackwell Publishing.
- . 2005b. "On 'Truth is Good'." *Philosophical Books* 46 (4): 292-301.
- Dreier, James. 1992. "The Supervenience Argument Against Moral Realism." *Southern Journal of Philosophy* 30 (3): 13-38.
- Dworkin, Ronald. 1996. "Objectivity and Truth: You'd Better Believe it." *Philosophy and Public Affairs* 25 (2): 87-139.
- Engel, Pascal. 2013. "Doxastic Correctness." *Aristotelian Society Supplementary Volume* 87 (1): 199-216.
- Enoch, David. 2006. "Agency, Shmagency: Why Normativity Won't Come from what is Constitutive of Action." *Philosophical Review* 115 (2): 169-198.
- . 2011a. *Taking Morality Seriously: A Defense of Robust Realism* Oxford University Press.
- . 2011b. "Reason-Giving and the Law." In *Oxford Studies in Philosophy of Law*, edited by Leslie Green and Brian Leiter: Oxford University Press.
- . 2011c. "Shmagency Revisited." In *New Waves in Metaethics*, edited by Michael Brady: Palgrave Macmillan.
- Fassio, Davide. 2015. "Belief, Aim of." In *Internet Encyclopedia of Philosophy*.
- . 2011. "Belief, Correctness and Normativity." *Logique Et Analyse* 54 (216): 471.
- Feldman, Richard. 1988. "Epistemic Obligations." *Philosophical Perspectives* 2: 235-256.
- . 2000. "The Ethics of Belief." *Philosophy and Phenomenological Research* 60 (3): 667-695.
- Ferrero, Luca. 2009. "Constitutivism and the Inescapability of Agency." *Oxford Studies in Metaethics* 4: 303-333.

- Field, Hartry. 1998. "Epistemological Nonfactualism and the A Prioricity of Logic." *Philosophical Studies* 92 (1/2): 1-24.
- . 2009. "Epistemology without Metaphysics." *Philosophical Studies* 143 (2): 249-290.
- Fine, Kit. 2002. "Varieties of Necessity." In *Conceivability and Possibility*, edited by Tamar Szabo Gendler and John Hawthorne, 253-281: Oxford Up.
- Finlay, Stephen. 2008. "The Error in the Error Theory." *Australasian Journal of Philosophy* 86 (3): 347-369.
- . 2010. "Recent Work on Normativity." *Analysis* 70 (2): 331-346.
- . 2014. *Confusion of Tongues: A Theory of Normative Language* Oxford University Press.
- Fitzpatrick, William. 2008. "Robust Ethical Realism, Non-Naturalism, and Normativity." *Oxford Studies in Metaethics* 3: 159-205.
- . 2011. "Ethical Non-Naturalism and Normative Properties." In *New Waves in Metaethics*, edited by Michael Brady: Palgrave Macmillan.
- Foley, Richard. 1987. *The Theory of Epistemic Rationality* Harvard University Press.
- . 1991. "Evidence and Reasons for Belief." *Analysis* 51 (2): 98-102.
- . 1992. *Working without a Net: A Study of Egocentric Epistemology* Oxford University Press.
- Foot, Philippa. 1972. "Morality as a System of Hypothetical Imperatives." *Philosophical Review* 81 (3): 305-316.
- . 2001. *Natural Goodness* Oxford University Press.
- Garner, Richard T. 1990. "On the Genuine Queerness of Moral Properties and Facts." *Australasian Journal of Philosophy* 68 (2): 137-146.
- Geach, P. T. 1956. "Good and Evil." *Analysis* 17 (2): 33-42.
- Gibbard, Allan. 1990. *Wise Choices, Apt Feelings: A Theory of Normative Judgment* Harvard University Press.
- . 2003. *Thinking how to Live* Harvard University Press.
- . 2005. "Truth and Correct Belief." *Philosophical Issues* 15 (1): 338-350.
- Gibbons, John. 2013. *The Norm of Belief* OUP Oxford.

- Goldman, Alan H. 2009. *Reasons from within: Desires and Values* Oxford University Press.
- Goldman, Alvin I. 1999a. *Knowledge in a Social World* Oxford University Press.
- . 1999b. "Internalism Exposed." *Journal of Philosophy* 96 (6): 271-293.
- Greco, Daniel. 2015. "Epistemological Open Questions." *Australasian Journal of Philosophy* 93 (3): 509-523.
- Grimm, Stephen R. 2009. "Epistemic Normativity." In *Epistemic Value*, edited by Adrian Haddock, Alan Millar and Duncan Pritchard, 243-264: Oxford: Oxford University Press.
- Haji, Ishtiyaque. 2002. *Deontic Morality and Control* Cambridge University Press.
- Hampton, Jean. 1998. *The Authority of Reason* Cambridge University Press.
- Hare, R. M. 1952. *The Language of Morals* Oxford Clarendon Press.
- Harman, Gilbert. 1975. "Moral Relativism Defended." *Philosophical Review* 84 (1): 3-22.
- . 1977. *The Nature of Morality: An Introduction to Ethics* Oxford University Press.
- Hattiangadi, Anandi. 2007. *Oughts and Thoughts: Rule-Following and the Normativity of Content* Oxford University Press.
- . 2010. "The Love of Truth." *Studies in History and Philosophy of Science Part A* 41 (4): 422-432.
- Hazlett, Allan. 2013. *A Luxury of the Understanding: On the Value of True Belief* OUP Oxford.
- Heathwood, Chris. 2009. "Moral and Epistemic Open-Question Arguments." *Philosophical Books* 50 (2): 83-98.
- Hieronymi, Pamela. 2006. "Controlling Attitudes." *Pacific Philosophical Quarterly* 87 (1): 45-74.
- Hood, Bruce. 2012. *The Self Illusion: Why there is no 'You' Inside Your Head* Hachette UK.
- Horgan, Terence and Mark Timmons. 1992. "Troubles on Moral Twin Earth: Moral Queerness Revived." *Synthese* 92 (2): 221-260.
- . 2006. "Cognitivist Expressivism." In *Metaethics After Moore*, edited by Terry Horgan and Mark Timmons, 255-298: Oxford University Press.

- Horwich, Paul. 2006. "The Value of Truth." *Noûs* 40 (2): 347-360.
- Hume, David. 1978. *A Treatise of Human Nature, 1739, 2d Edition of 1888, Edited by LA Selby Bigge, Revised by PH Nidditch* Oxford: Clarendon.
- Jackson, Frank. 1998. *From Metaphysics to Ethics: A Defence of Conceptual Analysis* Oxford University Press.
- James, William. 1995. *The Will to Believe: And Other Writings from William James* Image Books.
- Jarvis, Benjamin. 2012. "Norms of Intentionality: Norms that Don't Guide." *Philosophical Studies* 157 (1): 1-25.
- Joyce, Richard. 2001. *The Myth of Morality* Cambridge University Press.
- . 2006. *The Evolution of Morality* MIT Press.
- Joyce, Richard and Simon Kirchin. 2010. *A World without Values* Springer.
- Kahane, Guy. 2011. "Evolutionary Debunking Arguments." *Noûs* 45 (1): 103-125.
- Kappel, Klemens. 2010. "Expressivism about Knowledge and the Value of Knowledge." *Acta Analytica* 25 (2): 175-194.
- Kappel, Klemens and Emil F. L. Moeller. 2013. "Epistemic Expressivism and the Argument from Motivation." *Synthese* 191 (7): 1-19.
- Karmo, Toomas. 1988. "Some Valid (but no Sound) Arguments Trivially Span the 'Is'-'Ought' Gap." *Mind* 97 (386): 252-257.
- Katsafanas, Paul. 2013. *Agency and the Foundations of Ethics: Nietzschean Constitutivism* Oxford University Press.
- . Forthcoming. "Constitutivism about Practical Reasons." In *Oxford Handbook of Reasons and Normativity*, edited by Daniel Star: Oxford.
- Kelly, Thomas. 2002. "The Rationality of Belief and Other Propositional Attitudes." *Philosophical Studies* 110 (2): 163-196.
- . 2003. "Epistemic Rationality as Instrumental Rationality: A Critique." *Philosophy and Phenomenological Research* 66 (3): 612-640.
- . 2007. "Evidence and Normativity: Reply to Leite." *Philosophy and Phenomenological Research* 75 (2): 465-474.
- Kerr, Eric T. and J. Adam Carter. 2015. "Richard Rorty and Epistemic Normativity." *Social Epistemology* 30 (1): 3-24.

- Kim, Jaegwon. 1988. "What is Naturalized Epistemology?" *Philosophical Perspectives* 2: 381-405.
- Kitcher, Philip. 2005. "Biology and Ethics." In *The Oxford Handbook of Ethical Theory*, edited by David Copp: Oxford University Press.
- . 2011. *The Ethical Project* Harvard University Press.
- Kolodny, Niko. 2005. "Why be Rational?" *Mind* 114 (455): 509-563.
- Kornblith, Hilary. 1993. "Epistemic Normativity." *Synthese* 94 (3): 357-376.
- . 2002. *Knowledge and its Place in Nature* Oxford University Press.
- Korsgaard, Christine M. 1986. "Skepticism about Practical Reason." *Journal of Philosophy* 83 (1): 5-25.
- . 1996. *The Sources of Normativity* Cambridge University Press.
- . 2008. *The Constitution of Agency: Essays on Practical Reason and Moral Psychology* Oxford University Press.
- . 2009. *Self-Constitution: Agency, Identity, and Integrity* Oxford University Press.
- Kusch, Martin. 2002. *Knowledge by Agreement: The Programme of Communitarian Epistemology* Oxford University Press.
- Kvanvig, Jonathan L. 2003. *The Value of Knowledge and the Pursuit of Understanding* Cambridge University Press.
- . 2008. "Pointless Truth." *Midwest Studies in Philosophy* 32 (1): 199-212.
- Kyriacou, Christos. 2012. "Habits-Expressivism about Epistemic Justification." *Philosophical Papers* 41 (2): 209-237.
- Laudan, Larry. 1990a. "Aim-Less Epistemology?" *Studies in History and Philosophy of Science Part A* 21 (2): 315-322.
- . 1990b. "Normative Naturalism." *Philosophy of Science* 57 (1): 44-59.
- Lavin, Douglas. 2004. "Practical Reason and the Possibility of Error." *Ethics* 114 (3): 424-457.
- Leite, Adam. 2007. "Epistemic Instrumentalism and Reasons for Belief: A Reply to Tom Kelly's Epistemic Rationality as Instrumental Rationality: A Critique." *Philosophy and Phenomenological Research* 75 (2): 456-464.
- Leiter, Brian. 2001. "Moral Facts and Best Explanations." *Social Philosophy and Policy* 18 (2): 79.

- . 2006. "The Hermeneutics of Suspicion: Recovering Marx, Nietzsche, and Freud." In *The Future for Philosophy*, edited by Brian Leiter: Clarendon Press.
- Lillehammer, Hallvard. 2004. "Moral Error Theory." *Proceedings of the Aristotelian Society* 104 (2): 93-109.
- . 2007. *Companions in Guilt: Arguments for Ethical Objectivity* Palgrave Macmillan.
- . 2013. "The Argument from Queerness." In *The International Encyclopedia of Ethics*, edited by Hugh LaFollette.
- . Forthcoming. "Ethics, Evolution and the a Priori: Ross on Spencer and the French Sociologists." In *The Cambridge Handbook of Evolutionary Ethics*, edited by Robert Richards Michael Ruse.
- Littlejohn, Clayton. 2012a. *Justification and the Truth-Connection* Cambridge University Press.
- . 2012b. "Does 'Ought' Still Imply 'Can'?" *Philosophia* 40 (4): 821-828.
- . 2013. "The Russellian Retreat." *Proceedings of the Aristotelian Society* 113 (3pt3): 293-320.
- . Forthcoming a. "A Plea for Epistemic Excuses." In *The New Evil Demon Problem*, edited by Fabian Dorsch and Julien Dutant: Oxford University Press.
- . Forthcoming b. "The Truth in Gnosticism." *Analisis Revista De Investigacion Filosofica*.
- . Forthcoming c. "Objectivism and Subjectivism in Epistemology." In *The Factive Turn*, edited by Veli Mitova: Cambridge University Press.
- . Forthcoming d. "The Right in the Good: A Defense of Teleological Non-Consequentialism in Epistemology." In *Epistemic Consequentialism*, edited by Kristoffer Ahlstrom-Vij and Jeff Dunn: Oxford University Press.
- Littlejohn, Clayton and John Turri. 2013. *Epistemic Norms: New Essays on Action, Belief and Assertion* Oxford: Oxford University Press.
- Lockard, Matthew. 2013. "Epistemic Instrumentalism." *Synthese* 190 (9): 1701-1718.
- Loeb, D. 1998. "Moral Realism and the Argument from Disagreement." *Philosophical Studies* 90 (3): 281-303.
- Lynch, Michael. 2004. *True to Life: Why Truth Matters* Cambridge, Mass.: MIT Press.
- . 2009. "The Truth of Values and the Values of Truth'." In *Epistemic Value*, edited by Pritchard, Haddock and Millar: Oxford: Oxford University Press.

- Mackie, J. L. 1977. *Ethics: Inventing Right and Wrong* Penguin.
- Maffie, James. 1990. "Naturalism and the Normativity of Epistemology." *Philosophical Studies* 59 (3): 333-349.
- Maitzen, Stephen. 1995. "Our Errant Epistemic Aim." *Philosophy and Phenomenological Research* 55 (4): 869-876.
- Maguire, Barry. 2015. "Grounding the Autonomy of Ethics." *Oxford Studies in Metaethics* (Vol. 10).
- . 2016. "The Value-Based Theory of Reasons." *Ergo, an Open Access Journal of Philosophy* 3.
- Maguire, Barry and Jack Woods. Ms. "Explaining Epistemic Normativity."
- Manne, Kate. 2014. "Internalism about Reasons: Sad but True?" *Philosophical Studies* 167 (1): 89-117.
- Markovits, Julia. 2014. *Moral Reason* OUP Oxford.
- McHugh, Conor. 2011. "What do we Aim at when we Believe?" *Dialectica* 65 (3): 369-392.
- . 2012a. "The Truth Norm of Belief." *Pacific Philosophical Quarterly* 93 (1): 8-30.
- . 2012b. "Epistemic Deontology and Voluntariness." *Erkenntnis* 77 (1): 65-94.
- . 2014. "Fitting Belief." *Proceedings of the Aristotelian Society* 114: 167-187.
- McHugh, Conor and Daniel Whiting. 2014. "The Normativity of Belief." *Analysis* 74 (4): 698-713.
- McHugh, Conor and Jonathan Way. Forthcoming. "Fittingness First." *Ethics*.
- McKay, Ryan T. and Daniel C. Dennett. 2009. "The Evolution of Misbelief." *Behavioral and Brain Sciences* 32 (6): 493.
- McPherson, Tristram. 2011. "Against Quietist Normative Realism." *Philosophical Studies* 154 (2): 223-240.
- . 2012. "Mark Schroeder's Hypotheticalism: Agent-Neutrality, Moral Epistemology, and Methodology." *Philosophical Studies* 157 (3): 445-453.
- Millar, Alan. 2004. *Understanding People: Normativity and Rationalizing Explanation* Oxford University Press.
- Moore, G. E. 1903. *Principia Ethica* Dover Publications.

- Nagel, Thomas. 1986. *The View from Nowhere* Oxford University Press.
- Nelson, Mark T. 2010. "We have no Positive Epistemic Duties." *Mind* 119 (473): 83-102.
- Neta, Ram. 2007. "In Defense of Epistemic Relativism." *Episteme* 4 (1): 30-48.
- Nietzsche, Friedrich. 2008. *On the Genealogy of Morals: A Polemic by Way of Clarification and Supplement to My Last Book 'Beyond Good and Evil'* Oxford University Press.
- Nolfi, Kate. 2013. "Why is Epistemic Evaluation Prescriptive?" *Inquiry* 57 (1): 97-121.
- . 2015. "How to be a Normativist about the Nature of Belief." *Pacific Philosophical Quarterly* 96 (2): 181-204.
- Nozick, Robert. 1993. *The Nature of Rationality* Princeton University Press.
- Oddie, Graham. 2005. *Value, Reality, and Desire* Clarendon Press.
- O'Hagan, Emer. 2004. "Practical Identity and the Constitution of Agency." *Journal of Value Inquiry* 38 (1): 49-59.
- . 2005. "Belief, Normativity and the Constitution of Agency." *Philosophical Explorations* 8 (1): 39-52.
- Olson, Jonas. 2014. *Moral Error Theory: History, Critique, Defence* OUP Oxford.
- Owens, David. 2000. *Reason without Freedom: The Problem of Epistemic Normativity* Routledge.
- . 2003. "Does Belief have an Aim?" *Philosophical Studies* 115 (3): 283-305.
- Papineau, David. 1999. "Normativity and Judgment." *Proceedings of the Aristotelian Society* 73 (73): 16-43.
- . 2013. "There are no Norms of Belief." In *The Aim of Belief*, edited by T. Chan.
- Parfit, Derek. 2011. *On what Matters* Oxford University Press.
- Pettit, Philip and Michael Smith. 1996. "Freedom in Belief and Desire." *Journal of Philosophy* 93 (9): 429-449.
- Pigden, Charles R. 1989. "Logic and the Autonomy of Ethics." *Australasian Journal of Philosophy* 67 (2): 127-151.
- . 2010. *Hume on is and Ought*. Palgrave Macmillan.
- Plantinga, Alvin. 1993. *Warrant and Proper Function* Oxford University Press.

- Prior, A. N. 1960. "The Autonomy of Ethics." *Australasian Journal of Philosophy* 38 (3): 199-206.
- Prichard, Harold Arthur. 1949. *Moral Obligation* Oxford, Clarendon Press.
- Quine, W. V. 1969. "Epistemology Naturalized." In *Ontological Relativity and Other Essays*: New York: Columbia University Press.
- . 1992. *Pursuit of Truth* Harvard University Press.
- Railton, Peter. 1986. "Moral Realism." *Philosophical Review* 95 (2): 163-207.
- . 1994. "Truth, Reason, and the Regulation of Belief." *Philosophical Issues* 5: 71-93.
- . 1997. "On the Hypothetical and Non-Hypothetical in Reasoning about Belief and Action." In *Ethics and Practical Reason*, edited by Garrett Cullity and Berys Nigel Gaut, 53-79: Oxford University Press.
- Raz, Joseph. 2011. *From Normativity to Responsibility* OUP Oxford.
- Reisner, Andrew. 2008. "Weighing Pragmatic and Evidential Reasons for Belief." *Philosophical Studies* 138 (1): 17-27.
- . 2009a. "Abandoning the Buck Passing Analysis of Final Value." *Ethical Theory and Moral Practice* 12 (4): 379-395.
- . 2009b. "The Possibility of Pragmatic Reasons for Belief and the Wrong Kind of Reasons Problem." *Philosophical Studies* 145 (2): 257-272.
- . 2013. "Leaps of Knowledge." In *The Aim of Belief*, edited by Timothy Chan, 167-183: OUP.
- . 2014. "A Short Refutation of Strict Normative Evidentialism." *Inquiry* (5): 1-9.
- . 2015. "Normative Conflicts and the Structure of Normativity." In *Weighing and Reasoning: Themes from the Work of John Broome*, edited by Iwao Hirose and Andrew Reisner: OUP.
- . Forthcoming. "Pragmatic Reasons for Belief." In *The Oxford Handbook of Reasons and Normativity*, edited by Daniel Star: Oxford University Press.
- Rescher, Nicholas. 1970. "On the Characterization of Actions." In *The Nature of Human Action*, edited by M. Brand: Glenview, Ill: Scott, Foresman and Co.
- Ridge, Michael. 2007. "Expressivism and Epistemology: Epistemology for Ecumenical Expressivists." *Aristotelian Society Supplementary Volume* 81 (1): 83-108.
- . 2014. *Impassioned Belief* OUP Oxford.

- Roojen, Mark van. 2015. *Metaethics: A Contemporary Introduction* Routledge.
- Rorty, R. 1981. *Philosophy and the Mirror of Nature* Princeton University Press.
- Rosati, Connie S. 1995. "Naturalism, Normativity, and the Open Question Argument." *Noûs* 29 (1): 46-70.
- Ross, W. D. 1939. *Foundations of Ethics*. Oxford: Clarendon Press.
- Rowland, Richard. 2013. "Moral Error Theory and the Argument from Epistemic Reasons." *Journal of Ethics and Social Philosophy* 7 (1): 1-24.
- Ruse, Michael. 1986. *Taking Darwin Seriously: A Naturalistic Approach to Philosophy* Prometheus Books.
- Scanlon, T. M. 1998. *What we Owe to each Other* Harvard University Press.
- . 2014. *Being Realistic about Reasons* Oxford University Press.
- Schroeder, Mark. 2007. *Slaves of the Passions* Oxford University Press.
- . 2008. *Being for: Evaluating the Semantic Program of Expressivism* Oxford University Press.
- . 2011. "Ought, Agents, and Actions." *Philosophical Review* 120 (1): 1-41.
- Sellars, Wilfrid S. 1956. "Empiricism and the Philosophy of Mind." *Minnesota Studies in the Philosophy of Science* 1: 253-329.
- Sepielli, Andrew. forthcoming. "Subjective and Objective Reasons." In *Oxford Handbook of Reasons and Normativity*, edited by Daniel Star: Oxford.
- Setiya, Kieran. 2003. "Explaining Action." *Philosophical Review* 112 (3): 339-393.
- Shafer-Landau, Russ. 2003. *Moral Realism: A Defence* Oxford University Press.
- . 2006. "Ethics as Philosophy : A Defense of Ethical Nonnaturalism." In *Metaethics After Moore*, edited by Terry Horgan and Mark Timmons: Oxford University Press.
- Shah, Nishi. 2006. "A New Argument for Evidentialism." *Philosophical Quarterly* 56 (225): 481-498.
- Shah, Nishi and J. David Velleman. 2005. "Doxastic Deliberation." *Philosophical Review* 114 (4): 497-534.
- Sinclair, Neil. 2009. "Recent Work in Expressivism." *Analysis* 69 (1): 136-147.
- Skorupski, John. 2010. *The Domain of Reasons* Oxford University Press.

- Smith, Michael. 1994. *The Moral Problem* Blackwell.
- Smithies, Declan. 2012. "The Normative Role of Knowledge." *Noûs* 46 (2): 265-288.
- Sorensen, Roy A. 1988. *Blindspots* Oxford University Press.
- Sosa, Ernest. 2007. *A Virtue Epistemology* Oxford University Press.
- Srinivasan, Amia. 2015. "Normativity without Cartesian Privilege." *Philosophical Issues* 25 (1): 273-299.
- Steglich-Petersen, Asbjorn. 2008. "Does Doxastic Transparency Support Evidentialism?" *Dialectica* 62 (4): 541-547.
- . 2009. "Weighing the Aim of Belief." *Philosophical Studies* 145 (3): 395-405.
- . 2011. "How to be a Teleologist about Epistemic Reasons." In *Reasons for Belief*, edited by Asbjorn Steglich-Petersen and Andrew Reisner, 13-33: Cambridge University Press.
- . 2013. "Truth as the Aim of Epistemic Justification." In *The Aim of Belief*, edited by Timothy Chan: Oxford University Press.
- Steup, Matthias. 1988. "The Deontic Conception of Epistemic Justification." *Philosophical Studies* 53 (1): 65-84.
- Stevenson, Charles L. 1937. "The Emotive Meaning of Ethical Terms." *Mind* 46 (181): 14-31.
- . 1963. *Facts and Values* Yale University Press.
- Strandberg, Caj. 2004. *Moral Reality. A Defence of Moral Realism* Lund University.
- Stratton-Lake, Philip. 2002. "Introduction" In W.D. Ross *The Right and the Good*: Oxford University Press.
- Street, Sharon. 2006. "A Darwinian Dilemma for Realist Theories of Value." *Philosophical Studies* 127 (1): 109-166.
- . 2008. "Constructivism about Reasons." *Oxford Studies in Metaethics* 3: 207-245. Oxford University Press.
- . 2009. "Evolution and the Normativity of Epistemic Reasons." *Canadian Journal of Philosophy* 39 (supl): 213-248.
- Streumer, Bart. 2003. "Does 'ought' conversationally Implicate 'can'?" *European Journal of Philosophy* 11 (2): 219-228.
- Sturgeon, Nicholas. 1998. "Moral Explanations." In *Ethical Theory I: The Question of Objectivity*, edited by James Rachels: OUP Oxford.

- Sutton, Jonathan. 2005. "Stick to what You Know." *Nous* 39 (3): 359-396.
- . 2007. *Without Justification* MIT Press.
- Sylvan, Kurt. 2012. "Truth Monism without Teleology." *Thought: A Journal of Philosophy* 1 (3): 161-163.
- Taylor, Shelley E. 1989. *Positive Illusions: Creative Self-Deception and the Healthy Mind*. Basic Books.
- Taylor, Shelley E. and Jonathon D. Brown. 1988. "Illusion and Well-being: A Social Psychological Perspective on Mental Health." *Psychological Bulletin* 103 (2): 193.
- . 1994. "Positive Illusions and Well-being Revisited: Separating Fact from Fiction." .
- Taylor, Shelley E., Margaret E. Kemeny, Lisa G. Aspinwall, Stephen G. Schneider, Richard Rodriguez, and Mark Herbert. 1992. "Optimism, Coping, Psychological Distress, and High-Risk Sexual Behavior among Men at Risk for Acquired Immunodeficiency Syndrome (AIDS)." *Journal of Personality and Social Psychology* 63 (3): 460.
- Taylor, Shelley E., Jennifer S. Lerner, David K. Sherman, Rebecca M. Sage, and Nina K. McDowell. 2003. "Are Self-Enhancing Cognitions Associated with Healthy Or Unhealthy Biological Profiles?" *Journal of Personality and Social Psychology* 85 (4): 605.
- Thomson, Judith Jarvis. 2008. *Normativity* Open Court.
- Tiberius, Valerie. 2000. "Humean Heroism: Value Commitments and the Source of Normativity." *Pacific Philosophical Quarterly* 81 (4): 426-446.
- Tiffany, Evan. 2007. "Deflationary Normative Pluralism." *Canadian Journal of Philosophy* 37 (5): 231-262.
- Toppinen, Teemu. 2015. "Expressivism and the Normativity of Attitudes." *Pacific Philosophical Quarterly* 96 (2): 233-255.
- Tubert, Ariela. 2010. "Constitutive Arguments." *Philosophy Compass* 5 (8): 656-666.
- . 2011. "Korsgaard's Constitutive Arguments and the Principles of Practical Reason." *Philosophical Quarterly* 61 (243): 343-362.
- Urmson, J. O. 1968. *The Emotive Theory of Ethics* London, Hutchinson.
- Vahid, Hamid. 2006. "Aiming at Truth: Doxastic Vs. Epistemic Goals." *Philosophical Studies* 131 (2): 303-335.
- Van Roojen, Mark. 2015. *Metaethics: A Contemporary Introduction* Routledge.

- Vayrynen, Pekka. 2006. "Resisting the Buck-Passing Account of Value." In *Oxford Studies in Metaethics, Volume 1*, edited by Russ Shafer-Landau, 295-324: Oxford University Press.
- Velleman, David. 2000. "On the Aim of Belief." In *The Possibility of Practical Reason*, edited by David Velleman, 244-281: Oxford University Press.
- Vranas, Peter. 2007. "I Ought, therefore I can." *Philosophical Studies* 136 (2): 167-216.
- Wallace, Jay. 2001. "Normativity, Commitment, and Instrumental Reason." *Philosophers' Imprint* 1.
- Wallace, R. Jay. 2006. *Normativity and the Will: Selected Papers on Moral Psychology and Practical Reason* Oxford University Press.
- Way, Jonathan. 2013. "Value and Reasons to Favour." *Oxford Studies in Metaethics* 8.
- . Forthcoming a. "Two Arguments for Evidentialism." *Philosophical Quarterly*.
- . Forthcoming b. "Reasons and Rationality." In *Oxford Handbook of Reasons and Normativity*, edited by Daniel Star: Oxford University Press.
- Wedgwood, Ralph. 2002. "The Aim of Belief." *Philosophical Perspectives* 36 (s16): 267-297.
- . 2007. *The Nature of Normativity* Oxford University Press.
- White, Roger. 2007. "Epistemic Subjectivism." *Episteme* 4 (1): 115-129.
- Whiting, Daniel. 2010. "Should I Believe the Truth?" *Dialectica* 64 (2): 213-224.
- . 2012. "Does Belief Aim at the Truth?" *Pacific Philosophical Quarterly* 93 (2): 279-300.
- . 2013a. "Truth: The Aim and Norm of Belief." *Teorema: International Journal of Philosophy* 32 (3): 121-136.
- . 2013b. "Nothing but the Truth: On the Norms and Aims of Belief." In *The Aim of Belief*, edited by Timothy Chan: Oxford University Press.
- . 2013c. "The Good and the True (Or the Bad and the False)." *Philosophy* 8 (2): 219-242.
- Wiland, Eric. 2012. *Reasons Continuum*.
- Williams, Bernard. 1973. "Deciding to Believe." In *Problems of the Self*, edited by Bernard Williams, 136-151: Cambridge University Press.
- . 1979. "Internal and External Reasons." In *Rational Action*, edited by Ross Harrison, 101-113: Cambridge University Press.

- . 1985. *Ethics and the Limits of Philosophy* Harvard University Press.
- Williamson, Timothy. 2000. *Knowledge and its Limits* Oxford University Press.
- . Forthcoming. "Justifications, Excuses, and Sceptical Scenarios." In *The New Evil Demon Problem*, edited by Fabian Dorsch and Julien Dutant: Oxford University Press.
- Wong, David B. 1984. *Moral Relativity* University of California Press.
- Wrenn, Chase. 2006. "Epistemology as Engineering?" *Theoria* 72 (1): 60-79.
- . 2010. "True Belief is Not Instrumentally Valuable." In *New Waves in Truth*, edited by Cory D. Wright and Nikolaj J. L. L. Pedersen: Palgrave Macmillan.
- Wright, Crispin. 1992. *Truth and Objectivity* Harvard University Press.
- . 1995. "Truth in Ethics." *Ratio* 8 (3): 209-226.
- Wright, G. H. von. 1963. *Norm and Action* New York, Humanities.
- Zagzebski, Linda. 2003. "The Search for the Source of Epistemic Good." *Metaphilosophy* 34 (1-2): 12-28.
- Zimmerman, Michael J. 1996. *The Concept of Moral Obligation* Cambridge University Press.
- . 2002. "Taking Luck Seriously." *Journal of Philosophy* 99 (11): 553-576.
- . 2008. *Living with Uncertainty: The Moral Significance of Ignorance* Cambridge University Press.