

Evolutionary Theory and Normative Realism in Epistemology

Alessandro Di Nicola
Green Templeton College



A thesis submitted for the degree of Doctor of Philosophy
at the University of Oxford

Trinity Term 2014

Evolutionary Theory and Normative Realism in Epistemology

Alessandro Di Nicola, Green Templeton College
Oxford University, Trinity Term 2014

Abstract of Thesis Submitted for the Degree of Doctor of Philosophy

In this thesis I discuss one way in which evolutionary theory has been brought to bear on the evaluation of competing meta-normative views in epistemology. More specifically, I investigate whether normative realism in epistemology (epistemic realism) is compatible with the view that we are justified in holding many of the epistemic beliefs we hold, on the assumption that those beliefs can be explained in evolutionary terms.

In *Part I* I discuss normativity in epistemology and meta-epistemology. I begin by drawing attention to the fact that there are very different ways of understanding which concepts in epistemology are normative and what their normativity consists in. I focus on the concept of an *epistemic reason* to illustrate this point. I then discuss, in some detail, how different interpretations of epistemic-normativity will affect the form which normative-realist views in epistemology can take. I conclude by drawing a taxonomy of epistemic-realist views which is mindful of the different interpretations of epistemic normativity which I distinguished.

In *Part II* I turn to discuss the topic of Darwinian arguments against normative-realist views in epistemology. I begin by considering the form which Darwinian arguments have taken in recent *metaethical* literature. I argue that Darwinian arguments of a kind which is meant to raise a distinctively *epistemological* challenge for normative-realist views – I call those ‘Darwinian-epistemological’ arguments – represent a more interesting object of philosophical investigation than Darwinian arguments of a different ‘metaphysical’ kind. I then formulate a Darwinian-epistemological argument which targets normative-realist views in epistemology (DEA), explain how it works, and spend some time discussing its key premises.

In *Part III* I engage with the Darwinian-epistemological challenge against realism about epistemic normativity that Sharon Street presents in her paper ‘Evolution and the Normativity of Epistemic Reasons’ (2009). I argue that Street’s epistemological challenge is best viewed as an instance of a Darwinian-epistemological argument of the kind I formulated (DEA). I then go on to reply to Street’s argument on behalf of the epistemic realist. I conclude that arguments with the form of DEA do not represent a serious threat for normative realist views in epistemology.

Acknowledgements

During the time in which I have worked on this thesis – and in the years leading up to it – I have received support and encouragement, in various forms, from many individuals and more than one institution.

The Philosophy Faculty and Balliol College supported my research with a postgraduate studentship for the two years of my B.Phil. degree. The Philosophy Faculty and Brasenose College funded the subsequent three years of my doctoral research. Green Templeton College provided the ideal environment in which to bring this thesis to completion. I am very grateful to these institutions for their support.

On a more personal note, I would like to express my most sincere gratitude to my supervisors, Bill Child and Guy Kahane. Both Bill and Guy went through an inordinate number of drafts of the thesis at its various stages of development, they provided detailed comments on all of them, and were most generous with insightful criticism throughout. Their guidance, help, commitment and – when it was needed – encouragement were exemplary. I have not only learned a great amount of philosophy from them, but also about the personal and professional qualities of philosophers. It has been a great privilege to work under their supervision.

Thanks also to Scott Sturgeon, whose fearsome criticism of some earlier drafts of this thesis helped to make it a lot better than it would otherwise have been; and to John Broome, whose precious words of advice at a decisive moment were instrumental to my decision to remain at Oxford and begin work on this thesis.

My family in Italy have been extremely supportive throughout. They are too many to mention all but I would like to thank, in particular: my parents, my grandmothers, my brother (Birillo), my sister (Tara'), Giuseppe Bornago (Mendoza), Mauro Malvi, and Alessandro Trovini for looking after one another during difficult times and, by doing so, for giving me the opportunity to live abroad for so long in pursuit of my studies in philosophy.

I am very grateful to the many friends – philosophers and non-philosophers – whose companionship and support I have benefited from in all those years. Thanks to Lorenzo Greco and Matteo Falomi (Il Reggimento), Ben Sorgiovanni and Jens Ziska (The Dumplings), Yair Levy, Giulia Orlando, Natalia Waights-Hickman, Carole Foucher, Marcus Nabielek, Miroslawa Alunowska Figueroa, Giang Nguyen, Vincent Keeling, Dominic Ryan, David Kilcline, David King, Lorenzo Migliorini, Karlis Zauers, Ann Geraghty, Paolo Di Marco, Alessio Andreoli, Cristina Ostorero, and Elisabetta Bassani.

Thanks also to Chris Timpson, Tom Sinclair, Peter Powell, Aase Villadsen, Chris Sauer, John and Sandra Feneley, Paul Monod, and Bernard Gowers for their advice, support, and kindness.

The time I spent working on this thesis was made all the more precious by the presence of some most dear friends. My very special thanks to Andrea Polo and Giovanni Merlo, who have been like brothers to me through all those years, and to Shivprasad Swaminathan for all that I have learned from him and for our countless discussions about Indian philosophy and mysticism.

Lastly, I would like to express my deepest gratitude to Betiel Wasihun, *la mia lucente stella Diana*. Without her love, patience, and unswerving support I cannot imagine what it would have been like to bring this thesis to an end.

To My Grandfather

Table of Contents

Introduction.....	7
Part I.....	13
Chapter 1: Coming to Terms with the Normative.....	14
1. Introduction	14
2. Normative Concepts and Normative Reasons	18
2.1. Normative Concepts.....	18
2.2. Normative Concepts through Normative Reasons.....	21
2.3. Normative Reasons.....	23
2.4. Normative Concepts and the question ‘Is X normative?’.....	28
Chapter 2: Locating Normativity in the Epistemic Domain	34
1. Introduction: Epistemic Normativity	34
2. Evidence	43
3. Three Concepts of Epistemic Reasons.....	47
3.1. Epistemic Reasons as Evidence-Based Reasons: Some Conceptual Distinctions	47
3.1.1. The Instrumentalist view of Epistemic Reasons (Epistemic Instrumentalism).....	49
3.1.2. The Evidentialist View of Epistemic Reasons	52
3.1.3. Evidence-as-Reasons Interpretations: Analytical Epistemic Naturalism and the Evidence as Irreducibly Normative Views.....	53
4. Importance of the Distinctions for the Project of Locating Normativity in Epistemology	55
Chapter 3: Varieties of Normative Realism in Epistemology.....	60
1. Introduction	60
2. Meta-Epistemology	61
3. A Note on the Definition of Normative Realism.....	65
3.1 Different Definitions of Normative Realism.....	65
3.2 Normative Realism and Normative Facts	67
4. Varieties of Epistemic Realism I: Irreducibility, Objectivity, Causality	68
4.1 Irreducibility	69
4.1.1. Challenging NECESSITY	72
4.1.2. Challenging NON-CIRCULARITY.....	73
4.1.3. Challenging ESSENCE.....	74
4.2 Irreducibility and Objectivity.....	79
4.3 Irreducibility and Causality.....	81
5. Varieties of Epistemic Realism II: Instrumentalism, Evidentialism, and Evidence-as-Reasons Views.....	85
5.1 Realism and Instrumentalism	86
5.2 Realism and Evidentialism	89
5.3 Realism and Evidence-as-Reasons views.....	94
6. A Spectrum of Epistemic Realist Views	96
Part II	98
Chapter 4: A Darwinian Epistemological Argument against Epistemic Realism.....	99
1. Introduction	99
2. Darwinian Arguments against Normative Realism.....	100
3. The Darwinian-Epistemological Challenge to Normative Realism in Epistemology	108
3.1. The Basic Form of the Argument.....	108
3.2. The Nature of the Argument: An Analogy	110
3.3. How the Argument Differs from Metaphysical-Darwinian Arguments and from other Epistemological Challenges.....	114

Chapter 5: The Evolutionary Hypothesis	116
1. Evolutionary Explanations of Beliefs	116
2. Evolutionary Explanations of Epistemic-Normative Beliefs	118
2.1. Their Explanandum	118
2.2. The Explanations at Work	121
3. <i>Strength</i> and <i>Pervasiveness</i> of Evolutionary Forces	125
3.1 Stronger and Weaker Evolutionary Hypotheses and Their Degrees of Confirmation.....	127
4. Evolutionary Forces, Epistemic-Normative Beliefs, and True Beliefs	130
4.1. <i>Desiderata</i> and <i>Constraints</i> on Explanations of Epistemic-Normative Beliefs	131
4.2. ‘Selectivity’ of Darwinian-Epistemological Arguments	133
4.3. How these <i>Desiderata</i> and <i>Constraints</i> Affect the Use of DEA as an Argument against Epistemic Realism.....	137
 Chapter 6: Evolution and the Lack of Correlation between Epistemic-Normative Beliefs and Epistemic-Normative Facts	 141
1. Introduction	141
2. Epistemic-Realism and the Darwinian-Epistemological Argument.....	141
2.1 The Target of the Darwinian-Epistemological Argument.....	145
3. Evolutionary Forces as ‘Off-Track’ Causal Influences	147
3.1. Off-track: ‘Chanciness’, and ‘Predictability’	147
3.2. ‘Off-track’ Practical-Normative Beliefs	149
3.3. ‘Off-track’ Epistemic-Normative Beliefs.....	153
3.4. Predictability, Epistemic-Normative beliefs and Two Senses of ‘Lack of Correlation/Off-track’	158
4. When Do Evolutionary Explanation of Beliefs Support the Claim that those beliefs are ‘Off-Track’?.....	161
4.1. Of Milvian Bridges	164
5. Conclusion.....	167
 Part III	 168
 Chapter 7: Street’s Darwinian-Epistemological Argument: An Epistemic-Realist Reply	 169
1. Introduction	169
2. Street’s on the Darwinian Epistemological Problems of Epistemic Realism	171
3. Street’s Epistemological Argument against Evidentialist Realism.....	173
3.1. A Better Argument.....	177
3.2. An Evidentialist-Realist Reply to Street’s Argument	179
3.2.1. Street on ‘Lack of Correlation’	180
3.2.2. The Evidentialist Reply: Different Senses of ‘Off-track’.....	184
4. Street’s Epistemological Argument against Evidence-as-Reasons Realism.	192
4.1. A Better Argument.....	199
4.2. An Evidence-As-Reasons Realist Reply to the Argument.....	205
4.2.1. The Point So Far.....	205
4.2.2. Evidence-as-Reasons Realism and Irreducibly Normative Facts	207
5. Conclusion.....	211
 Conclusion	 213
 Bibliography	 216

Introduction

One important strand of discussion in recent metaethical literature has focused upon the bearing of evolutionary considerations on the assessment of realist views of practical normativity. A number of philosophers have endorsed the view that realism is in tension with ‘the least controversial and most well-founded speculations of evolutionary psychology’ (Street: 2006). Different kinds of ‘Darwinian arguments’ have been appealed to in order to motivate this claim; each kind affording a different – and preferred – way for making the exact nature of the alleged tension explicit (Joyce: 2013). A sub-set of those Darwinian arguments, in particular, has taken the form of an *epistemological* challenge to realism. According to those arguments, if realism about practical normativity were true our practical-normative beliefs would be unjustified (Street: 2006, Gibbard: 1990 and 2003, Joyce: 2013). Some philosophers who have put forward arguments of this latter kind – such as Alan Gibbard and Sharon Street – have also defended the view that the scope of those arguments is not limited to the practical domain. They believe that realism about *epistemic* normativity would similarly be called into question by arguments of this form (Gibbard: 1990, Street: 2009).

What do Darwinian-epistemological arguments against realism about epistemic normativity amount to? And, are they successful? These are the questions which set the agenda for the investigations contained in this thesis and to which, in the end, this thesis will provide an answer. These questions are worth pursuing for various reasons; some of them are the following.

I have just mentioned the arguments of philosophers who have defended the view that Darwinian considerations can be drawn upon in order to *raise* a distinctively epistemological challenge to normative realism. It is a fact worth noticing that some realist philosophers have defended the exact opposite claim. That is to say, they have argued that evolutionary

considerations can be appealed to in order to *respond to a* general epistemological challenge that is often levelled against the form of realism they favour. To this extent, Enoch writes:

Street's way of thinking about the epistemological challenge is also a good way to think about how to solve it. (2010:426)

Parfit agrees, and he thinks this is especially so in the case of *epistemic* reasons (2011: ch.7). This is a somewhat striking disagreement;¹ and it draws attention to an interesting and substantive philosophical question.

Secondly, our epistemic evaluative practices rely crucially – at least *prima facie*² – on the notion of an epistemic reason (White; 2010: 573). In particular, epistemic reasons are the kinds of reasons we appeal to in the justification of our beliefs; they are what we consider in assessing whether A knows that P, or whether A is rational or irrational in believing that P, and so forth. A meta-normative view that had, as a consequence, that our beliefs about epistemic reasons are unjustified would not only entail that we do not have knowledge about epistemic reasons. It would also, for instance, entail that we do not know *whether we are ever justified in believing anything* (Parfit; 2011:522). This ought to be considered a rather worrying result. Because of this, assessing the epistemological challenge that Darwinian arguments raise against epistemic-normative realism would seem all the more urgent. If those arguments were sound, there would be serious pressure to abandon realist views of epistemic normativity. And if turned out that we cannot do so, we will be left to face a troublesome form of scepticism.

Thirdly, specific discussion of the impact of Darwinian considerations on the assessment of competing meta-normative views in *epistemology* has so far received next to no attention, as opposed to the attention which considerations of a similar kind have attracted in the *moral* and

¹ The disagreement over the relevance of evolutionary considerations for the epistemology of normative belief is all the more surprising if one adds to the views I have just mentioned that of those who believe that evolutionary considerations *have no obvious or specific bearing* on this issue. (See, for instance, Nagel: 1997, Dworkin: 1996, White: 2010.)

² One might think that being justified (only) requires having (sufficient) *evidence*. The relation between epistemic reasons and evidence is far from being as straightforward and uncontroversial as it might seem (see Chapter 2 of this thesis). I believe we can appreciate the intuitive force of the point I am making here while keeping aside the complexities of that discussion.

practical normative domains. The debate is thus remarkable for its absence in the philosophical literature. This reflects – and perhaps is to be explained in terms of – the little attention which certain meta-normative issues in epistemology have, in general, received so far.

But this oversight, one might think, can also be otherwise explained. Perhaps it rests on the expectation that whatever the outcome of the debates in *metaethics* turns out to be, it will carry over unchanged to affect meta-normative views in the *epistemic* domain. Often the above expectation seems to be the result of the following assumption. In investigating the relevance of evolutionary theory to adjudicating between realist and anti-realist views of normativity, the discussion can be carried out at a level of generality that can safely ignore the distinction between normativity of *different kinds*. This latter assumption seems indeed to be in place in some influential treatments of those matters, such as David Enoch's (2010). Gibbard's anti-realist and broadly Darwinian account of the origin and function of normative practices and discourse is also meant to cover, in the same breath, the normativity of actions, emotions, and beliefs or, as he puts it, 'what to do or think or feel' (Gibbard; 1990: 72).

Are the *expectation* and the *assumption* which I have mentioned above justified? It seems to me that one can hardly decide on whether this is the case, without first paying close – and separate – attention to the impact that Darwinian considerations have on the assessment of meta-normative views in epistemology. Prior to doing so, we may think that the *specific differences* between the practical and the epistemic domains ought rather to make us wary of discussing 'under the same heading' the fate of normative realism in each of those domains.

How evolutionary theory bears on the assessment of competing views in what I will refer to as 'meta-epistemology' is thus not only a much under-addressed topic. It is a question the answer to which bears on the broader issues which I have just drawn attention to.

The first problem one encounters when attempting to deal with this topic is that locating normativity in the epistemic domain is far from a straightforward task. Despite the fact that – as most would agree – normative language figures prominently in epistemology, it is far from

obvious how one is to go about determining which concepts and notions are *normative* in this domain. A distinct, but related, problem is an even more conspicuous one. It amounts to the fact that specific discussions of normative realism in epistemology are hard to come by. This is all the more surprising, given the current degree of sophistication of discussions of *metaethical* realism in recent and contemporary philosophical literature.

Without a clear idea of what normative realism in epistemology amounts to, one can hardly go about assessing arguments – Darwinian or otherwise – which allegedly call a view of this kind into question. And without a clear idea of how to go about identifying the normative concepts and the normative notions which are at work in epistemic discourse, one can hardly set out on the task of formulating epistemic realism in epistemology. For one would not know – for instance – what the realists about normativity in epistemology are being realists *about*.

Thus addressing questions regarding the nature and strength of Darwinian-epistemological arguments against normative realism in epistemology can be expected to require some preliminary work on the issues I have just drawn attention to. The structure of the present thesis reflects – and is determined by – this fact.

The thesis is divided into three parts. In *Part I* (Chapters 1, 2, and 3) I discuss normativity in epistemology and meta-epistemology. I begin by drawing attention to the fact that there are very different ways of understanding which concepts in epistemology are normative and what their normativity consists in. I focus on the concept of an *epistemic reason* to illustrate this point. I then discuss – in some detail – how different interpretations of epistemic-normativity will affect the form which normative-realist views in epistemology can take. I conclude by drawing a taxonomy of epistemic-realist views which is mindful of the different interpretations of epistemic normativity which I distinguished.

In *Part II* (Chapters 4, 5, and 6), I turn to discuss the topic of Darwinian arguments against normative-realist views in epistemology. I begin by considering the form which Darwinian arguments have taken in recent *metaethical* literature. I argue that Darwinian arguments of a kind

which are meant to raise a distinctively *epistemological* challenge for normative-realist views – I call those ‘Darwinian-epistemological’ arguments – represent a more interesting object of philosophical investigation than Darwinian arguments of a different ‘metaphysical’ kind. I then formulate a Darwinian-epistemological argument which targets normative-realist views in epistemology (DEA), explain how it works and spend some time discussing its premises.

In *Part III* (Chapter 7) I engage with the Darwinian-epistemological argument against realism about epistemic normativity that Sharon Street presents in her paper ‘Evolution and the Normativity of Epistemic Reasons’ (2009). I argue that the argument which Street presents there is best viewed as an instance of a Darwinian-epistemological argument of the kind I formulated (DEA). I then go on to reply to Street’s argument on behalf of the epistemic realist. I conclude that arguments with the form of DEA do not represent a serious threat for normative realist views in epistemology.

It seems to me that the themes and investigations which comprise this thesis should be of interests to metaethicists (and philosophers of normativity more generally) and epistemologists alike. They should be of interest to the former because the study of normativity has been traditionally dominated by discussions of *practical* and *ethical* normativity. To the extent to which the results of these discussions have then been taken to apply to normativity *in general*, their outcomes have – arguably – suffered from an ‘insufficient diet of examples’. A case in point is the lack of debate surrounding the applicability to other normative domains of Darwinian arguments originally formulated to challenge *metaethical* realism. Reflection on normative and meta-normative questions in epistemology will also be relevant – and indirectly contribute to – the assessment of important debates in the philosophy of normativity, such as those regarding the nature and value of so called ‘companion-in-guilt’ defences of moral realism (Cuneo: 2007), or the prospects for a theory of normativity which is ‘fully unified across the practical and epistemic domains’ (Street; 2009: 213).

But the investigations of this thesis should also be of interest to epistemologists. This is because, as I have already pointed out, the use of normative language is ubiquitous in this discipline. Yet it is often unclear what the normativity of central normative concepts in epistemology is taken to amount to. In turn, it is unclear which of the central concepts in epistemology are normative, and what makes them so. Is evidence normative? Is justification understood in terms of the reliability of belief-forming methods or processes normative? Moreover, questions regarding the *ontology* of normativity in epistemology are rarely asked. Is there – for instance – a need for the appeal to *irreducibly normative facts* in epistemology? If so, what must these facts be like? If not, why not? Addressing these and similar questions would, arguably, lead to a better understanding of epistemic discourse. And an examination of these and similar questions would clearly benefit from employing some of the tools and terminology which are at work in traditional debates in ethics and metaethics.

Lastly, a quick word on how this thesis can be read. One obvious way to read the thesis is in the way I have presented it. That is to say, as raising a question – that regarding the force of a challenge of the Darwinian-epistemological kind against epistemic-normative realism – and as forging the tools which are required for answering it. But the thesis can also be read, first and foremost, as an exploration of some neglected issues regarding the nature of normative notions at work in epistemology; and as an illustration by means of a ‘case study’ – i.e. the discussion of Darwinian-epistemological arguments – of the value of work of this kind for the framing and, hopefully, the solution of interesting philosophical problems.

Part I

Chapter 1

Coming to Terms with the Normative

1. Introduction

The aim of this and of the next two chapters is two-fold. Firstly, it is to examine some central aspects of the use and function of normative concepts in epistemology, and to do so in a way which is mindful of discussions of normativity in the practical and moral sphere. Secondly, it is to explore the extent to which the familiar taxonomy of positions developed in metaethics can be used as a blueprint for providing a map of the hitherto largely unexplored field of meta-epistemology.

There are many reasons for wanting to engage in a pursuit of this kind. First-level and conceptual discussions of ‘the normative’ in the practical and moral sphere have reached a degree of complexity which, arguably, is not matched by parallel discussions of normativity in epistemology. And while vocabulary which is recognised and referred to as ‘normative’ is ubiquitous in both the epistemic and the practical domain, what the normativity in question is taken to consist in often differs significantly in each case: it is not uncommon to have the sense that what is being talked about are slightly different notions.

A comprehensive study of normativity would benefit from a discussion that attempted to bridge the gap between the use of normative concepts in epistemology and that in the practical sphere respectively. This enquiry would be preliminary to raising – and perhaps settling – questions regarding the extent to which the two domains are similar, or the respects in which they differ. Bringing the more developed conceptual framework in use in discussion of practical normativity to bear on epistemological matters can also help to disambiguate and clarify central

concepts in epistemology and be thus taken as an independent contribution to progress in this field.

A quick illustration of the benefits of conceptual work of the kind at issue is the following. It is not uncommon to come across philosophical works in which epistemic justification is referred to as a paradigmatic normative notion. But we might wonder whether this judgment would stand scrutiny if one were to apply to it an account of normativity which is often at work in discussions of normativity in the practical domain. The idea that it is a requirement for concepts to count as normative that they be *reason-implying* is often taken for granted in those discussions. But if this criterion is anything to go by, then on some account of what justification consists in, justification would not automatically qualify as a normative concept.

For instance, on some externalist accounts of justification, the justification of a belief is dependent upon the reliability of the method or process by which it is acquired – where the reliability at issue is a matter of the probability that the method or process in question will result in the formation of true beliefs. And one might think³ that the fact that a belief that P is justified in the above sense does not yet provide one with *a reason* to hold the belief in question. That is to say, one might think that there being a reason for believing the justified belief that P for agent A is contingent upon the presence in A of a cognitive goal of the appropriate kind (for instance the goal to form a true belief as to whether P). If this is the case, and if we adopt a ‘reason-implying’ criterion for the normativity of concepts, then it is not at all clear that the notion of epistemic justification can be taken as a paradigmatically normative one. On the externalist account I have mentioned, the issue as to whether epistemic justification is normative – far from being settled from the outset – would rather seem to be a substantive one. Being clear as to what conception of normativity is at work in epistemology and about *where* and *how* exactly normative concepts apply, can not only help disambiguate claims such as those regarding the normativity of justification, or provide the tools to show how substantive questions can hide behind apparent

³ As some externalists indeed do; see, for instance, Papineau (2006).

‘platitudes’. It can also help one raise questions and frame – and perhaps settle – disagreements regarding the alleged normativity of other crucial epistemic notions: such as, for instance, those of warrant, entitlement, evidence, and – in the epistemology of science – confirmation.

An enquiry into meta-normative issues arising in epistemology would likewise seem of great interest. *Given the analogies* between the practical and the epistemic domains – via the notion of *norms* and *value* and the central use of concepts such as *rational* and *reasons* which are employed in both – meta-epistemology is a surprisingly much less developed field than metaethics. It is a field – one might say – still ‘remarkable for its absence’. Witness to this is the fact that debates regarding the ontology and epistemology of normative properties have traditionally, if not exclusively, taken place in the context of discussions of the nature of *practical* normativity.

The extent to which traditional taxonomies in metaethics carry over to epistemology is also worth exploring for other, more specific, reasons. One among them is the following. Some realist positions in metaethics are often argued against by claiming that the kind of normative properties and facts they are committed to are ‘too queer’ for them to be taken seriously. Realists often reply to this charge by claiming that normative concepts also figure centrally in epistemology, and that providing an anti-realist understanding of these concepts in epistemic contexts is not a viable option. But, the realist continues, the normative properties and facts an ‘epistemic realist’ must be committed to are significantly similar to the ‘queer’ moral properties and facts the existence of which the anti-realist believed to be problematic. The intended upshot of these considerations is that moral properties are either not-so-queer, after all, or that their queerness should not put us off believing in them.⁴

This reply to the ‘queerness’ charge relies crucially on the analogies between epistemology and ethics. It similarly relies on the claim that anti-realism about normative discourse in epistemology is not a viable option, and that the realism one must adopt in epistemology will end up sharing

⁴ See, for instance, Cuneo (2008), Putnam (1990), Jackson (1999), and Parfit (2011). For discussion of related issues, see Railton (1997) and Velleman (2000).

the commitments of the versions of realism which anti-realists in ethics found ‘queer’. This and similar ‘companions-in-guilt’ defences of moral realism are often deployed in the philosophical literature, but rarely articulated in detail.⁵ Yet they clearly depend upon establishing in some detail whether the analogies between moral and epistemic normativity which they rely on really hold. What does being an epistemic realist about normative discourse in epistemology entail? And is it true that epistemic realism must share the features of moral realism that many anti-realists find unpalatable? Arguably, addressing these and related questions would require knowing more about how the field of metaethics and that of meta-epistemology are related. And – arguably – this task calls for a clearer map of the similarities and differences between the two than the one we currently have.

One specific version of the ‘companion-in-guilt’ defence of moral realism is particularly worth mentioning here. I have already (in the introduction) remarked upon how in recent years philosophers have appealed to evolutionary theory in order to formulate arguments against realist positions in metaethics. These arguments have rested upon the premise that moral realism is in tension with plausible empirical claims from evolutionary psychology (Joyce: 2006, Street: 2006, Gibbard: 1990 and 2003). One way of responding to this kind of challenge consists in holding that epistemic realism would similarly be called into question by this argument and that abandoning realism in epistemology is not a viable option.⁶ Once again, assessing what is at stake here will depend on investigating the strength and the specific nature of the analogies between normative discourse in ethics and epistemology and on their meta-normative underpinnings.

In this and the next two chapters I will be occupied with this endeavour. A comparative study of normativity in the epistemic and practical domain raises many and deep questions, and it would be unrealistic to attempt to provide a comprehensive treatment of all. My purpose, then, will have to be more modest. On the one hand, I will seek to provide the sketch for a ‘translation

⁵ With Cuneo (2007) among the notable exceptions.

⁶ Some of the considerations in Korsgaard (1996), Nagel (1997), and Putnam (1983) – for instance – can be read to provide support for this claim.

manual' that would allow epistemologists and philosophers working in the practical and the moral sphere to speak the same language, so to say, when it comes to the use of normative vocabulary. Ideally, a translation manual of this kind would have to be as neutral as it is allowed with respect to open and debated controversies in each field, so as to be of as large a use as possible. On the other hand, I will be occupied with sketching a 'rough guide' to meta-epistemology on the blueprint of the distinctions which are customarily drawn and observed in metaethics. All along, I will be concerned with bringing to the fore some issues which have been – to my knowledge – largely overlooked in discussions of normativity in the epistemic domain. Addressing those issues will turn out to be of crucial importance to the project that lies at the core of this thesis.

2. Normative Concepts and Normative Reasons

2.1. Normative Concepts

The use of normative vocabulary is ubiquitous in both ordinary and philosophical discourse. Providing a list of paradigmatic normative expressions, however, is not an easy task. Many would agree on the fact that the expressions in italics in the following sentences would have to figure in any list of normative terms: 'you *ought* not to laugh at other people's misfortunes', 'a jury *should* not convict a person without evidence for his guilt', 'we *ought* to help old ladies in distress whenever we can', 'the fact that X has been badly hurt in the accident is a *reason* to call for an ambulance', and similar.

But while there is relative agreement regarding what any list of central and paradigmatic normative notions would have to include, what other expressions also qualify as normative can be controversial. Do the expressions 'P is the rational thing to believe', 'X is a good action', or 'John is brave' all contain normative language? Perhaps only some do? If so, which? And, in any case, how do we go about deciding - i.e. what makes normative language *normative*?

The difficulty of answering those and related questions points to the fact that providing a philosophically viable characterization of what ‘the normativity of all that is normative’ consists in (Raz: 1999) is far from being an easy endeavour. Many of the intuitively plausible ways to approach this task turn problematic upon closer examination.

Trying to cash out the idea of norms in terms of ‘rules’ does not allow one to capture the fact that not all (sets of) rules are normative in the desired sense, and it would seem to prejudge the issue as to whether normativity is ‘rule-governed’ in a way which some philosophers would resist.⁷

Appealing to the notion of value is equally controversial: some philosophers claim that the fact that something is valuable, an action for instance, does not by itself provide a properly situated agent with reasons for performing it (Williams; 2001: 96).

The view that the *motivating* dimension of the normative is this latter’s distinctive feature (Mackie: 1977, Williams: 1980, Joyce: 2001) is similarly problematic and has been rejected by many who have argued, for instance, that the concept of a normative reason does not bear the tight links to motivation that this approach requires (Shafer-Landau: 2003; Svavarsdottir: 1999). Moreover, insofar as some philosophers think that normative vocabulary can be used to qualify beliefs and emotions, as well as actions, it is not clear that the emphasis on motivation provides a particularly felicitous way of capturing the normativity of *all* that is normative. Arguably, the concept of motivation is most often used to refer to a relation between an agent and *an action* of some sort. And whether believing and feeling are things that, like actions, can be said to be under an agent’s voluntary control is a much disputed question.⁸

Similar limitations would affect an account of normativity in terms of ‘deontic concepts’ – such as those of requirements, prohibitions, permissions, duty, etc, at least where these were understood as requiring that the responses or states of agents that these concepts are applicable

⁷ I have philosophers of a particularist bend in mind (Dancy: 2004), and those who claim that rationality is ‘uncodifiable’ (Child: 1993).

⁸ See, for instance, James (1979), Shah (2002), Hieronymi (2006), Frankish (2007), Southwood and Chuard (2009).

to be ‘sufficiently under the voluntary control’ of agents (Alston: 1982). Such an account would also fall short of accounting for the meaningfulness of a question such as: ‘are requirements (permissions, prohibitions, etc.) of kind X normative?’. To this extent, it would fail to capture a sense of ‘normative’ that is at work in many contemporary debates in philosophy (Broome: 1999, Dancy: 2009).

Lastly, trying to get some leverage from the contrast with *the descriptive* in order to bring to light what ‘the normative’ consists in would seem to be equally unhelpful. Meaningful expressions such as ‘you ought to F’, ‘It is rational to F’, ‘A is a reason for F-ing’ behave – at least *prima facie* – like assertions. In this regard, they resemble expressions such as ‘Fido is on the mat’, ‘the sky is blue’, and ‘the bus is late’. *Qua assertions*, sentences of this kind can be taken to express beliefs; and beliefs are *descriptions* of what is the case. This observation would seem to tell in favour of considering normative claims as a species of descriptive claims – much as beliefs about the colours of objects are – rather than as a class of claims that can be usefully contrasted with the claims we call descriptive. One might want to defend the view that, despite their appearance, normative sentences are not *assertoric*, and re-establish the distinction between the descriptive and the normative in this way.⁹ But this move requires substantive assumptions about the semantics of normative language and – many would claim – it can hardly be taken to provide a neutral characterization of the normative domain itself. Short of making this or analogously controversial assumptions then, the strategy of characterizing the normative in terms of the contrast with *the descriptive* is a problematic one. Firstly, it is not clear that this contrast ‘carves at any interesting joint’, so to speak, with respect to the task at hand. A second worry is that, even if the contrast were at all helpful, the strategy in question would risk being circular: if by descriptive one merely meant ‘not normative’, this approach could hardly deliver a particularly enlightening account of the normative itself.

⁹ One would also have to specify what kind of non-assertoric role normative expressions play – so as to distinguish them from other non-assertoric expressions that are not normative.

2.2. Normative Concepts through Normative Reasons

In the previous section I drew attention to the controversial nature of some of the most influential attempts to indicate which features of normative concepts make for their normativity. That an account is controversial does not, of course, entail that it is incorrect. On the other hand, the very idea of a single correct account of the kind at issue here might be mistaken. It could be that the task of isolating a single criterion that allowed one to determine which concepts are normative is a misconceived one. Perhaps there are many different senses of normativity at work, and even if they partially overlap, none of them might be such as to be representative of all. This would be tantamount to thinking that normative concepts are family-resemblance concepts.

Yet, it would be helpful if we could come up with some univocal way of characterizing normative concepts. If we are to conduct a comparative study of normative discourse across the domains in which normative concepts are used (practical, evaluative, epistemic, etc.), doing so would enable us to be clear about what kind of expressions we should be on the look-out for, and what it is that we should be investigating. If it turned out that there cannot be a single criterion for normativity, then an account of the kind in question would partly consist in fixing the meaning of ‘normative’ in a way which captures certain core features of normative concepts and which allows progress in the kind of investigations we are about to pursue. Moreover, it would be an advantage of one such way of proceeding if (1) the criterion it relied on was compatible with as many as possible of the different senses in which the concept has been used and (2) if it was as neutral as possible with respect to substantive and debated issues in the philosophy of normativity.

In what follows below, I will provide an account of normative concepts which is of this kind. Before outlining what it consists in, however, I suggest that we go back to our original task and try to be clearer about what ‘providing a characterization of the normative’ requires. The questions ‘which concepts are normative?’ and ‘what is it for a concept to be normative?’ can be

taken to set the agenda for slightly different tasks.¹⁰ For instance, they could be taken as a request for a *reductive account* of normative concepts – i.e. for a definition of normative concepts provided in non-normative terms. The questions could also be taken as the request for the identification of (one or more) normatively *basic concepts* and for the specification of the kind(s) of relation a given concept would have to bear to them in order for it to count as normative. Concepts of this kind would be such that all other normative concepts could be defined in terms of them (together with non-normative concepts), while they could not be defined in terms of the others. They would be ‘uniquely primitive’, as one might put it.¹¹ It might well be that the only way to get a grip on the normative must end with paradigmatically normative notions, and that the only account of normativity one can provide is circular in this sense. If this were the case, an account of the normative could still be ‘enlightening’ by identifying the most basic normative concepts, exploring the conceptual links they bear to one another, and showing how other normative concepts can be ‘derived’ from them.

I have suggested that we take the questions ‘which concepts are normative?’ and ‘what is it for a concept to be normative?’ as a request for as neutral and uncontroversial a characterization as possible of the normative domain. I take it that it would be an advantage for this characterization if it were neutral not only with respect to the substantive issues I have briefly mentioned in the previous section, but also about whether we can single out a normative concept which is *uniquely primitive*, or provide a specification of normative concepts in *non-normative* terms. This would help us proceed with the kind of investigations that I set out at the beginning of this chapter – i.e. to raise questions such as ‘what does it mean for a certain epistemic concept C to be normative?’ – with some degree of clarity and in a way which is both mindful of philosophical work in the practical domain and careful not to take sides on (too many) controversial issues in this field.

¹⁰ My discussion here is indebted to Robertson (2009).

¹¹ One such concept would be – for instance – the concept of a normative reason in what we may refer to as ‘buck-passing’ accounts of the normativity of all that is normative.

With this purpose in mind, I will proceed by taking the concept of a (normative) reason as a paradigmatic normative concept. I will then tease out some important conceptual links between this concept and other central normative notions, and suggest that we go about identifying which (other) concepts are normative by relying on the paradigm normative notion in question in the way(s) I will specify below. My aim is to use the concept of a reason as a heuristic device, so to speak; a particularly useful tool for circumscribing and exploring the normative domain. Whether its usefulness in this respect is an indication of its being ‘the most fundamental’ or the ‘uniquely primitive’ normative concept is an issue about which I wish to remain neutral.

2.3. Normative Reasons

The term ‘reason’ is not always used in a normative sense. Often, it is used to introduce what we might call non-normative explanations, as illustrated by the following sentences: ‘the reason why I was late is that I missed the bus’, ‘the reason why the washing machine is not working is that there is no electricity’ or ‘the reason why John raised his voice was that he was angry’. Normative uses of ‘reason’ are often (Scanlon: 1998, Dancy: 2009, Cuneo: 2007) distinguished from purely explanatory,¹² non-normative, ones by appeal to the idea that a normative reason is a consideration that *counts in favour of*/or *against* that for which it is a reason. Reasons in this sense typically *count in favour of/against* responses of a certain kind on part of properly situated agents. In what follows, I will restrict my attention to reasons as *favourers* – i.e. to what we might call ‘positive’ reasons – for the sake of simplicity. What I will say about reasons construed in this sense, I will also take to apply to ‘negative’ reasons, i.e. considerations against responses of a specified kind.

The idea of reason as *favourers* is not only helpful to distinguish normative from merely explanatory reasons. It is also instrumental to bringing to light the ‘relational structure’ of the

¹² I use ‘purely’ explanatory here (and ‘merely’ below) to signal the fact that reasons can be both explanatory and normative as in ‘the reason why John helped that old lady cross the street was that this was what he had *most reason* to do’.

concept of a normative reason itself. Without prejudging any substantive ontological issue, we can say that the concept of a normative reason behaves *as if* it referred to a complex relation. Underlying typical (normative) reason-ascription statements – whether this is made explicit or elliptically conveyed by the reason statement itself – is a basic structure of the following kind:

Consideration F favours response S of agent A in circumstances C.¹³

which we can also express as follows:

R (F, S, A, C)

The term ‘reason’, even when it is employed in a normative sense, is often used ambiguously. It is sometimes used to say that things are related in a certain way; i.e. that the *reason-relation* ‘relates’ the relata F, S, A, C. When we say of an agent that he has responded correctly to some reasons he had, arguably, we are saying that the agent in question was appropriately sensitive to *the fact that a certain response was called for (of him) in given circumstances*. One way to put this is by saying that the agent was sensitive *to the obtaining of the reason relation* between a certain fact and a given response. Perhaps more often in ordinary usage, we use ‘reason’ or ‘reasons’ to refer to *reason-giving* or *reason-providing* facts – i.e. to the considerations (F, G, etc.) that favour a given response of an agent/believer. One way to look at reasons in this sense, is by thinking of them as whatever we would be offered in answer to the questions of the kind: what is the reason for S-ing? What gives or provides A with a reason to S? To these considerations, or reasons in this sense, I will also refer to as the *ground(s)* of the reason-relation.

Reasons in the normative sense I am exploring are *pro tanto*. That is to say: the reasons that each of any of a number of grounds or *reason-giving* facts – F, G, H, etc. – provides can come in degrees and they can vary in strength or, as it is sometimes put, in ‘weight’. They can also differ

¹³ Things are more complicated, and a more adequate specification of the structure of reasons would require adding more variables to the very basic structure that I have introduced here. For instance, normative reasons are reasons *at a time*, and *to a certain degree*, etc. See, for instance Skorupski (2009).

‘in direction’, i.e. in the nature and content of the response which they favour. In virtue of these features of theirs, reasons can often be combined (‘added’, so to speak), and they can conflict. There can be reasons of different strength for and against a given response of an agent and when this is the case, determining what there is *most reason* or *decisive reasons* for him to do – for instance – will typically require ‘weighing’ reasons against one another. The strength and direction of the relevant reasons will similarly be what decide whether there are *sufficient/insufficient* reasons for given responses on parts of specific agents, or what there is *overall* reason for them to do.

Thinking of normative reasons in terms of ‘considerations in favour’ allows for a relatively uncontroversial initial ‘grip’ on this concept. Any further purchase on the concept however, will call for a more systematic investigation. It will have to build on the many ways in which we speak and think about reasons in this normative sense in ordinary and less ordinary contexts. And it will require raising questions – and introducing distinctions – which would allow us to make this concept ‘sharper’ and to arrive at a better understanding of it. Here, unsurprisingly perhaps, we quickly enter controversial philosophical territory. Philosophers have disagreed over what kinds of things are eligible to count as reasons – whether they are facts, mental states, or (true) propositions, for instance. Similarly, they have disagreed on whether reasons depend on the contingent set of pro-attitudes the agents for whom they are reasons hold at any given time; or about whether a fact (or proposition, or mental state) must be *accessible* for it to count as a reason and, if so, about what this relation of *accessibility* amounts to. Again, they have disagreed about whether the reason relation is conceptually primitive. And the list could continue.

The relevant literature also abounds in distinctions between different *kinds* of normative reasons, all of them intended to shed light on the concept at issue. Reasons have been ‘spoken of’ in many ways: *subjective* and *objective*, *veridical* and *non-veridical*, *internal* and *external*, *wide-scope* and *narrow-scope*, *apparent* and *actual*, the *reasons an agent has* and *the reasons there are (for that agent)*, to mention but a few. Opinions have differed regarding which of these distinctions are useful and

which misleading; and about how exactly the distinctions which are important to preserve are to be understood, how they are related, which of them are more fundamental, and so on.

If the concept of a normative reason is to be of any use to us here, we must find a way to proceed in the midst of such complexity. We can do so by side-stepping the controversies and disagreements the solution to which does not affect our investigation; and by ‘fixing’ the controversial aspects that might arise, thus restricting the focus of the discussion to a notion of normative reason which has been made more precise in this way. I will return to say a bit more in defence of this way of proceeding shortly. Let us take a normative-reason statement of the following kind as our example:

F is a reason for A to S (in C).

I will take the notion of a normative reason here to have the following features. Firstly, reasons are *propositional* in form. This means not only that the kinds of considerations which provide reasons can always be referred to by means of ‘that’ clauses, but also that any adequate specification of reasons must be sensitive to the ways in which things are presented or described. Thus, when Jane has had a heart attack, it is *that Jane has had a heart attack* which provides or gives me a reason to call an ambulance. Or it is *that the universe is expanding* that provides a reason to think that the big bang theory is true. As we will see, some philosophers think that these kinds of considerations alone are not, by themselves, sufficient to provide reasons for the responses in question. I will take it that – even if these philosophers were correct in thinking that this is the case – whatever else needs to be added to these considerations in order for them to favour those responses will also be propositional in form. Similarly, on the *description-sensitive* account of reasons I am adopting here, that A is *the winner of the boxing heavyweight championship of the world* might be a reason for me to admire him. However, that A is the father of so-and-so need not be,

even on the assumption that both things are true of A, and that each of the two expressions I have used is a legitimate way of referring to him.

Secondly, reasons do not entirely depend on – nor are they determined by – the *non-factive* mental states of agents. Thus, for instance, whether F is a reason for A to S does not depend on whether A believes *that F*, or on whether A believes *that F is a reason for S*. I take this claim to entail that for F to be a reason for A to S, F must be true. Reasons in this sense are, as many would put it, *factive*. If we operate with a notion of fact which equates facts with true propositions (Skorupski; 209: 115), then we can also say that ‘reasons are provided by facts’ (Parfit: 2011).

That reasons are *factive* in the sense I have just specified is compatible with thinking that occurrences of non-factive mental states can, sometimes, be themselves reason-providing, as when we say that the fact that some dictator plans to pursue some cruel aim fostered by, say, racial prejudice is a reason to believe that he is morally depraved; or that the fact that A is under the spell of recurring hallucinations is a reason for A to seek treatment of some sort (Anscombe; 2005: 115, Clarke-Doane; 2012: 316-317). I also take the claim that reasons are not *entirely* determined by the non-factive mental states of agents to be compatible with the idea that reasons might nonetheless *partly* depend on mental states of this kind; for instance, on the fact that the agents (for whom they are reasons) have pro-attitudes – preferences, desires, etc. – of some sort. I will return to say more about the pro-attitude dependence (or independence) of reasons in the following chapters.

Before moving on to the next section, I want to spend a few words to defend the way I have chosen to proceed here. If I have restricted and clarified in a somewhat arbitrary way the concept of a normative reason I will be relying on, this fact need not prejudge the value of the investigations which follow. Arguably, much of what I will have to say in the course of this and the following chapter does not depend crucially on the aspects of this concept which I have ‘fixed’ in the discussion so far. For instance, it seems to me that the main argument of this thesis can be adapted to various alternative frameworks, i.e. to views of normative reasons which are

different from the ones I have adopted here. In some instances this will be more obvious than in others, although I cannot give this matter the full attention it deserves here. In any case, the notion of a reason which I am here adopting as ‘standard’ reflects views which are held by many in the philosophical literature on normativity. Thus, even if one were to question the extent to which the discussions which follow are universalisable to cover other approaches, their value would nonetheless be great.

With the concept of a normative reason now in hand, we are ready to put it to work in the search for criteria for the normativity of concepts. I will turn to this task in the next section.

2.4. Normative Concepts and the question ‘Is X normative?’

Having now become clearer about the central concept of a normative reason,¹⁴ I suggest that we use this concept in order to address the question of which (other) concepts are *normative* concepts. Let us consider any concept C. I will assume that if C is the concept of a normative reason, then C is a normative concept. In the case in which C is not the concept of a normative reason, then we can formulate two different conditions on the basis of which we might go about establishing whether C is a normative concept. These conditions can be seen to correspond to two different *senses* in which the question ‘Is C normative?’ can be understood. The first I will call the ‘essentialist’ sense; the second the ‘reason-implying’ sense. Likewise, I will refer to the first condition as to the ‘essentialist’ condition, and to the second as to ‘the reason-implying’ condition.

It seems to me that each of those senses and conditions ought to be mentioned in an account of normativity which attempted to capture and preserve (or ‘fix the meaning of’) the central ways in which normativity has been ‘spoken of’ in the philosophical literature. Thus I suggest that we

¹⁴ Unless otherwise specified, any occurrence of ‘reason’ or ‘reasons’ from now on must be taken to express the concept of a *normative* reason.

take a concept C to be a *normative* concept iff it is either the concept of a reason or it fulfils either of these two conditions:

Condition I/Essentialism: Concept C *bears essential ties to the concept of a reason*. This can happen in any of three ways: (1) 'The concept of a reason'¹⁵ must be mentioned in providing the definition of concept C (2) Concept C must be mentioned in providing the definition of the concept of a reason (3) The concept of a reason and concept C must each be mentioned in providing the definition of the other. If C fulfils this condition, I will say that C is normative *in the essentialist sense* (Wedgwood: 2007).

Condition II/Reason-Implication: Concept C is *a reason-implying concept*, in the sense that, when C occurs in a proposition P which is of the form *X is C*, P is such as to imply a proposition Q which is of the form *A has a reason to S*. In other words, the concept of something which is C is the concept of something which is such that – in virtue of it being C – it *provides* a reason for a given response S on the part of an appropriately situated agent A. If C fulfils this condition, I will say that C is normative *in the reason-implying sense* (Parfit; 2011: 31-5).

Before I elaborate further on these conditions, I need to clarify the following. It is possible to disagree with how I have suggested we go about identifying which concepts are normative. In particular, some might think that concepts are normative when – and only when – they (are either the concept of a normative reason or) fulfil Condition I alone; in other words, that we ought to reject Condition II, perhaps because it is in some way misleading, or because it does not capture any interesting way in which we talk about or use normative concepts (Wedgwood: 2009, Horwich: 1998). I think this view is mistaken and part of what I will say below will also serve to

¹⁵ Together with other, non-normative concepts perhaps, such as those conveyed by the use of logical vocabulary, for instance.

illustrate why this is so. In any case, my overall strategy does not depend on settling this question. Ultimately, my point is to bring to view the utility of the concept of a normative reason for thinking about the normativity of other concepts. And the concept of a reason will still play this role if we chose to restrict the way we go about specifying which concepts are normative by means of Condition I alone.

This said, let me elaborate some more on what these criteria amount to and on how they are related. It is important to notice that Condition I does not require that the concept of a reason be a *uniquely primitive* normative concept. That is to say, the concept of a reason need not be such that both (1) it is not to be defined in more basic normative terms and (2) all other normative terms are to be defined in terms of this concept together with non-normative ones. The condition is then compatible with the thesis that there are other *more primitive* or *equally primitive* normative concepts. In this latter case, satisfying the condition would only require that the concept of a reason and (any other) normative concept C must each be mentioned in giving an account of the meaning of the other, so that the two concepts would be essentially interdependent, without either of the two being reducible to the other. An example might help illustrate this point. Let us assume that the following two bi-conditionals are true:

F is a reason for/against ϕ -ing iff *a rational agent*, in normal conditions, upon becoming aware of F would be disposed to take F into consideration when deciding whether to ϕ .

and:

F is a reason for Agent A in favour/against ϕ -ing iff by ϕ -ing as a response to F A would exemplify any of a series of epistemic/moral/practical *virtues* or *merits/vices* or *demerits*.

And let us also take it that each of them captures *essential ties* between the concepts mentioned on either sides of the bi-conditional. Given what we have assumed bi-conditionals of the kind above can be interpreted in different ways. They can be taken as providing a definition of the concepts of rationality and the virtues in terms of the concept of a reason. Or they can be taken to provide definitions of the concept of a reason in terms of that of a rational agent and of the virtues, respectively. Or again, they can be interpreted as making explicit certain essential conceptual links between ‘equi-primitive’ concepts. Condition I is to be understood as compatible with each of these readings.

It is important to draw attention to the fact that a normative concept being *primitive* does not entail its being *irreducibly normative*. A concept is irreducibly normative when an account of what it is to be that concept – if it can be provided at all – must appeal to concepts which are themselves normative. Clearly, not all *primitive* normative concepts need to be *irreducible* in this way. For instance, one can think that the concept of a (normative) reason is a uniquely primitive concept, and explain other normative concepts in its terms. But this is compatible with thinking that – in the moral domain – what it means to say that F is a reason for doing P is, say, that F is conducive to an increase in the well-being of a large number of individuals, etc. Or that – in the epistemic domain – what it means to say that F is an epistemic reason for believing that P is that F raises the probability of P being true. In each of these cases, the concept of a reason is *reduced* to other (I am assuming) non-normative concepts. If this is the case, a so called ‘analytic naturalist’ understanding of the concept of a reason is compatible with the idea that this concept qualifies as a *uniquely primitive* normative one.

The question ‘Is concept C normative?’ (Is goodness normative? Is evidence normative? Is justification normative? etc.) is sometimes not meant as a question about the existence of essential ties between C and the concept of a reason. It is rather best heard as a question as to whether a given concept is *reason-implying* in the sense I have specified. This brings us to the

question of the relation between the first and the second condition for the normativity of concepts which I have provided.

Clearly, a concept can fail to fulfil Condition I and still manage to fulfil Condition II. How this is so can be illustrated by the following examples. One might think that the concept of goodness does not bear the essential ties to the concept of a reason that satisfaction of Condition I requires; for instance because they believe that saying that something, i.e. an action, is (morally) good is equivalent in meaning to saying that the action in question contributes to an overall increase in happiness¹⁶ for a great number of people. Or perhaps because it is believed that the concept of goodness is un-analysable and, in particular, that it cannot be explained – in the ‘way buck-passers’ claim (Scanlon: 1998, Parfit: 2011) – in terms of the concept of a reason (Crisp: 2006). Thus, on either of those views, goodness would not count as a normative concept by Condition I. Holding either of these views on goodness is, however, compatible with believing that the concept of goodness is such that saying that an action is good *entails* the claim that some agent A has a reason for doing it.¹⁷ If it were true that goodness is a reason-implying concept in this sense, goodness would still qualify as normative by my Condition II. Similarly – as we will see in the next chapter – one can think that the concept of evidence does not bear the essential ties to the concept of a reason that the first condition requires of normative concepts. To specify what it means to say that F is evidence for P, we might not need to use the vocabulary of reasons: perhaps saying that F is evidence for P is equivalent in meaning to saying that (the occurrence of) F increases the probability of P’s truth. Yet we can still believe that the claim that F evidence for P implies the claim to that some agent A has a reason to believe that P. If this were the case, the concept of evidence would not be normative by Condition I, but it would still count as normative by Condition II.

¹⁶ I am assuming that the concept of happiness is not an irreducibly normative one.

¹⁷ Or one can think that the concept of ‘murder’ does not bear any essential ties to the concept of a reason, and yet believe that it is a necessary truth that if a given act counts as an instance of murder, then this provides one with reasons against performing/in favour of preventing an act of this kind.

Clearly, some concepts can fulfil both conditions at once. This will be the case, for instance, if one subscribes to a *buck-passing* account of goodness – i.e. if one believes that the concept of goodness reduces to the concept of something's having some feature that provides reasons for a response of some specifiable kind of an agent. On this account, the concept 'good' will be normative by Condition I. It will also be normative by Condition II insofar as claiming that something – say an action – is good will imply that some agent has a reason for doing it.

Drawing attention to the two criteria for the normativity of concepts I have formulated is instrumental to disambiguating two very different senses in which the question 'is C normative?' is often asked of concepts in the context of philosophical discussions. Disagreements as to whether 'thick' notions such as 'courageous', 'generous', 'intelligent', 'perceptive', etc. (or 'cowardly', 'mean', 'insensitive', etc.) are *normative* and as to what their normativity consists in can be clearly made sense of in the framework I have provided. So can debates regarding the normativity of rationality (Broome: 1999), the normativity of evidence (Leite: 2007, Kelly: 2007) or, for instance, the normativity of values (Williams: 2001).¹⁸ It also provides one with the tools for showing how there are often substantive questions to ask regarding concepts the normativity of which is often – and often quite misleadingly – taken for granted. For what the normativity of these concepts consists in can be understood in very different ways, and bringing those to view can help us see substantive disagreements regarding both the nature of those concepts and the nature of the properties that they purportedly refer to. The issue as to whether the notion of epistemic justification is a normative one being a prime example of a case of this kind. But this will become clearer in the next chapter, where I turn to look at the normativity of the concepts we use in epistemology.

¹⁸ Bernard Williams (2001) believes that the fact that something is valuable does not entail that there is reason for an agent to do it. If this claim was true, then values would not be normative by either of the criteria I have suggested. It seems to me an advantage of the criteria I have adopted that they allow for a position like Williams's – i.e. one in which values and the issue of the latter's normativity can come apart.

Chapter 2

Locating Normativity in the Epistemic Domain

1. Introduction: Epistemic Normativity

A discussion of epistemic normativity must proceed on the basis of some idea of what is distinctive of the sub-class of normative concepts which we want to mark off as *epistemic*. This might not be as easy a task as it looks. In what follows, I will briefly explain why I think the most common way of absolving this task is not obviously the most adequate. I will thereby suggest that we adopt a different strategy: one which is at the same time less controversial than the one which is commonly used and particularly fit for the specific goals of my investigations in this thesis.

In the previous chapter I drew attention to the centrality of the role that the concept of a normative reason can be made to play in identifying the normativity of all that is normative. I will thereby focus my discussion of epistemic normativity on the concept of an epistemic reason. By being clear about the concept of an epistemic reason we will then be able – on the blueprint of the criteria (Conditions I and II) I have provided in the previous chapter – to go about identifying which (other) concepts in epistemology are normative and what their normativity consists in.

When I first discussed the concept of a normative reason, I said that one helpful way to elucidate it is by saying that reasons are *considerations in favour* of certain responses on part of agents or believers. In coming to draw a distinction between epistemic and practical (including moral) reasons, it is thus tempting – and quite customary¹⁹ – to make it rest on the *kind of response*

¹⁹ See for instance Millar (2009), Skorupski (2009), Street (2009).

which reasons of each kind typically favour. One obvious way to do this, for instance, is by thinking of practical reasons as *reasons for actions* and of epistemic reasons as *reasons for belief*. However, this way of drawing the required distinction is defective in many respects. It is so – it can be argued – to the extent that it does not do full justice to the complexity of normative discourse both in the practical and the epistemic domain.

For one can hold that the range of responses that are favoured by practical reasons is not exhausted by actions, and that it includes emotions and beliefs among other things. One can perhaps have moral – or prudential – reasons to be *kind*, for instance, or moral reasons to *believe* that one's best friend is not a liar. Similarly, the range of responses that are favoured by epistemic reasons would not appear to be exhausted by beliefs, but rather also to include emotions and actions. Perhaps there are epistemic reasons not to feel angry in contexts in which this emotion is 'intellectually obfuscating' (Cuneo; 2007: 72), or to refrain from engaging in certain courses of action, when doing so impairs one's possibility to represent reality aright. Again one might think that the mental state of *hoping* that something is the case can sometimes be appraised by using normative vocabulary of a distinctively epistemic kind, in the sense that epistemic reasons would, in certain circumstances, bear on whether we have reasons to hope that something is the case: as when we say that A has no reason to *hope* that P is the case because P can never be true; or because A is not justified in believing that P can ever be true. Similarly, *adopting a hypothesis* can also be subject to appraisal which invokes reasons of an epistemic kind, as when an army general assumes that the enemy will attack from a certain side, where this assumption serves the purpose of better organising a defensive strategy (Skorupski; 2009: 114); and the circumstances are such that the general has evidence for the fact that the enemy will attack from the opposite side from the one he is assuming they will attack from. Here it would appear to be the case that the practical reasons the general has are constrained by (negative) epistemic reasons; and that we can appeal directly to those latter epistemic reasons in order to support the claim that assuming what the general did, in the circumstances I described, was (epistemically) irrational on his part.

If what I have just said is correct, it would testify to the fact that a wider range of responses than actions and beliefs can be favoured (or be told against) by practical and epistemic reasons respectively. But this kind of consideration also goes some way towards showing that the set of responses favoured by epistemic and practical reasons might well overlap, and that reasons of each kind can favour, in the appropriate circumstances, responses such as actions, beliefs, and emotions, etc. of various kinds (Zagzebski: 1996). Some philosophers have even gone as far as believing that there are ‘hybrid’ reasons, i.e. reasons for given responses on part of agents which are both *epistemic* and *practical* at the same time. Reasons to be fair in evaluating the views of others – it has been argued – are an instance of reasons of this kind. (Cuneo; 2007: 77).

There is a strategy that goes some way²⁰ towards defending the distinction between practical and epistemic reasons, when this is drawn on the basis of the kind of response reasons apply to – i.e. actions vs. beliefs. It consists in showing that there cannot be – appearances to the contrary – practical reasons for belief. This is because practical reasons only apply to responses which are under the voluntary control of agents (namely, actions), and beliefs are not of this kind.²¹ This approach is often supported by claims to the effect that alleged cases of practical reasons for belief turn out on analysis to be cases of reasons for action in disguise (Skorupski: 2009). Thus *there might seem to be* a practical reason for A *to believe* that he is healthier than he is regardless of his having evidence to the contrary, if this increases his chances of healing or allows him to live his last days in a happier state. But the reason in question, it is suggested, is *really* a reason *to get oneself to have the belief* in question. And reasons *to get oneself to believe* – as opposed to straightforward reasons *to believe* – are a kind of *reasons for action*.

This way of defending the ‘response-based’ characterization of epistemic reasons is often run together with the idea that the only reasons for belief there might be must be grounded on there

²⁰ It does not go *all the way* because the strategy at issue calls into question the claim that there can be practical reasons for belief. But a full defence of the ‘response-based’ view will also have to show that there cannot be epistemic reason for responses of agents other than beliefs, such as – for instance – emotions and actions, etc.

²¹ Besides Skorupski (2009) see also Alston (1988) for a general understanding of normativity that lends support to a strategy of this kind.

being evidence for the belief in question (Skorupski: 2009). This entails that any alleged reason for believing that is not sensitive to evidence cannot qualify as a reason for believing at all. If anything looks like a reason for believing which contravenes this requirement then, again, it must turn out to be a reason for *getting oneself to believe* instead.

Each of these commitments can be – and has been²² – questioned. And while we might think it plausible that believing is not in any way under the voluntary control of the believer, it is perhaps less plausible to think that it is a requirement of genuine reasons for belief that they be responsive to evidence in the way this approach requires. The example of ‘the friend and the liar’ I have mentioned above can be used to illustrate this point. The reasons to believe that my best friend is not lying seem to be largely independent²³ of the evidence I might have for believing so. Moreover, they would seem to be genuine reasons for belief. To the point that, arguably, if I were *to get myself to believe* that my friend is not lying in order to believe that he is not lying, I would not be a good friend. One way to put this is by saying that I would not be appropriately responsive to the (moral, presumably) reason to believe that he is not lying. It is precisely in the spontaneity of the belief in question – one might say – that the satisfaction of the requirement of friendship in this case resides²⁴.

Pending resolution of these controversies, it would be an advantage of a criterion for the individuation of epistemic reasons that it be as neutral as possible with respect to them, and as mindful as possible of the complexities of epistemic-normative discourse that I have drawn attention to. As we saw, the kind of considerations that – arguably – can be adduced in favour of (or against) holding beliefs of various kinds can be extremely diverse. This makes it the case that

²² The literature on ‘doxastic voluntarism’, i.e. on whether agents can believe ‘at will’, is rather extensive; for an exchange of (opposite) views on this topic see, for instance Hieronymi (2009) and Setiya (2008). Similarly, the question of whether there can be genuine pragmatic reasons for belief has received a lot of attention in recent philosophical literature, where it is often discussed under the heading of the ‘wrong kind of reasons’ problem, and also in connection with the distinction between ‘state-given’ and ‘object-given’ reasons. For an overview of the terms of those debates, see for instance, Schoeder (2012) and Reisner (2009).

²³ One does not need to make a case for *complete independence* here. It would suffice, in order to make the case for this fact, that at least some of the considerations in favour for holding a certain belief are not entirely based on (or commensurate to) considerations regarding which evidence one for that belief one has at his disposal.

²⁴ One might think that once we have established the plausibility of this view in the morality/friendship case, the way for the possibility of having genuine practical reasons for belief more generally has been paved.

identifying which reasons are to count as epistemic *merely* in terms of the attitude they favour (i.e. beliefs, as opposed to intentions, desires, etc.) is not very helpful; especially so if one's goal is to investigate the normative dimension of *traditional epistemological discourse*.

It remains true that central epistemological notions which are candidates²⁵ for the title of normative concepts – such as those of justification, warrant, entitlement, and the like – are primarily used to qualify *beliefs* (as in ‘P is justified/unjustified’) or *believers* with respect to the kind of beliefs they hold (as in ‘A is justified/unjustified in believing that P’), in the sense that – whether or not we believe that those notions can also be used to evaluate other mental states (such as emotions or intentions) – they are rather uncontroversially at work in the evaluation and appraisals of beliefs; and that they are most often used in this way. Central epistemological notions which are candidates for the title of normative concepts are also such that they bear some distinctive relation to the concept of truth. As Cuneo puts it:

Epistemic merits such as being justified, entitled, warranted, a case of knowledge, and so forth are almost always understood to be alethic merits [...] – merits the presence of which in an agent's attitude indicates that the content of this attitude is true, likely to be true, or that he has done a good job of trying to get at the truth. (2007: 58)

As I have indicated above, it is possible to disagree as to whether the *only* merits (or demerits) of beliefs (or believers, etc.) are to be understood in this way; for instance by thinking that some of the merits/demerits of beliefs are *pragmatic*. One way to make sense of merits/demerits of this latter kind, as we saw, is in terms of there being genuine *pragmatic* reasons for belief, where those reasons are such as not to bear the relation to truth – whatever it is – that is distinctive of central epistemological notions. Yet even if one believes that there are genuine pragmatic reasons for belief, it is hard to deny that *at least some* reasons for belief will bear the distinctive relation to truth that the passage I quoted from Cuneo attempted to capture.

²⁵ The reasons why I say that these notions are ‘candidate’ normative notions – rather than, for instance, ‘paradigmatic examples’ of normative notions, will become clear in what I go on to say below. In brief, these are that on some of the accounts of epistemic-normativity that I will review, those notions might end up *not being normative* after all.

Because of this, if what we are after is a criterion for the identification of epistemic-normative concepts that rely on a notion of reasons which is distinctively *epistemic*, then it would make sense to focus on a concept of reason which is also such as to allow us to capture the ‘connection to truth’ that central epistemological concepts display.

My suggestion²⁶ is that we specify what reasons count as *epistemic* reasons by building on the fairly uncontroversial facts that I have just drawn attention to above; and that we take *epistemic* reasons to be what I will call *evidence-based reasons for belief*. These are reasons which count in favour of a belief that P *at least partly* in virtue of the fact that they are *based on* what evidentially supports the truth of P, in the sense that if there is a reason to believe that P, and if this reason is to count as an *epistemic* reason, the following will be the case: (1) the presence of a reason of this kind will entail that there is *evidence* of some kind or other for the truth of P and (2) the evidence there is for P will enter the grounds of the reason for believing that P. We can also put this by saying that epistemic-reasons for believing that P *at least partly depend* on there being evidence for P.

In other words, we can tell whether a fact F is an epistemic reason for believing that P – I am suggesting – if when we ask ‘what is the reason F for believing that P?’ part of the correct answer will have to be ‘That E is the case’, where fact E is/provides evidence for the truth of P and fact E is either identical to or part of (complex) fact F.²⁷ Again, consider the following exchange: ‘What reason is there for believing that John is guilty of murder?’ ‘That John’s fingerprints are on the murder-weapon’. If it is true that the fact that John’s fingerprints are on the murder-weapon is evidence for John’s being the murderer, and if this fact is one of the grounds of the reason for believing that John is the murderer, then the reason in question will be an *epistemic* reason.

The notion of an *evidence-based reason for belief* is itself open to importantly different interpretations, and I will return to discuss these in some details in the next sections. For the

²⁶ This follows Street (2009), but also departs from her account in important ways. I will return to discuss these matters at some length in the last chapter of this thesis.

²⁷ This qualification – as well as the qualification that epistemic reasons *at least partly* depend on evidence – may seem rather obscure at this stage. The reason for those qualifications will become apparent in Section 3 of this chapter.

time being, however, we can see how this concept can be initially put to work in providing a first characterization of *epistemic-normative* concepts on the blueprint of the more general criteria for the normativity of concepts which I have already outlined. We can say that a concept EC is an epistemic-normative concept iff either (1) it is the concept of an evidence-based reason for belief or (2) it fulfils either of the following two conditions:

Condition I/Essentialism – Epistemic: Concept EC bears essential ties to the concept of an epistemic reason.

This can happen in any of three ways: (1) The concept of an *epistemic reason* must be mentioned in providing the definition of concept EC (2) Concept EC must be mentioned in providing the definition of the concept of an *epistemic reason* (3) The concept of an *epistemic reason* and concept EC must each be mentioned in providing the definition of the other. If C fulfils this condition, I will say that C is an *epistemic-normative* concept *in the essentialist sense*.

Condition II/ Reason-Implication – Epistemic: Concept EC is a *reason-implying concept*, in the sense that, when EC occurs in a proposition P which is of the form *X is EC*, P is such as to imply a proposition Q which is of the form *A has an epistemic reason to S*. In other words, the concept of something which is EC is the concept of something which is such that – in virtue of it being EC – it *provides an epistemic reason* for a given response S on the part of an appropriately situated agent A. If EC fulfils this condition, I will say that EC is an *epistemic-normative* concept *in the reason-implying sense* (Parfit; 2011: 31-5).

This first characterization of which concepts are to count as normative in an epistemic sense goes some way towards allowing us to adjudicate whether crucial epistemological notions are epistemic normative notions, and what their epistemic normativity consists in. For instance, suppose that we believe that the concept of epistemic justification is such that the following states a conceptual truth about it:

J: If belief P is justified (in the epistemic sense) then there are sufficient/strong evidence-based reasons to believe that P.

On this account, epistemic justification qualifies as an *epistemic normative concept*, in that it entails the concept of an evidence-based reason. That the concept of justification is a normative concept in this sense is, however, not a trivial result. Consider the case of philosophers who think that the notion of epistemic justification is to be understood in terms of reliability. That is to say, philosophers who hold that for a belief to be justified – roughly – is for that belief to be the result of a process or method of belief acquisition that yields true beliefs in a large enough number of occasions and in normal circumstances. And suppose that those philosophers also think that a claim to the extent that a belief is justified in this sense, by itself, says nothing about there being any reason (evidence-based or not) for holding it (Papineau: 2006, Street: 2009). On this interpretation of epistemic justification, the concept would neither bear essential ties to nor entail the concept of an evidence-based reason. Thus, it would not count as an epistemic normative concept by either of the conditions I specified.

In a similar way, it is easy to see how we can go about testing ‘thicker’ epistemological concepts for epistemic normativity – such as the concepts we employ when we call a belief ‘delusional’ or ‘preposterous’ or ‘accurate’, or a believer ‘confused’, ‘hasty’, or ‘inaccurate’. On the basis of the conditions I have provided, we can also allow for the fact that the concept of a normative reason may be employed in an epistemic sense even when the mental state or attitude it concerns is not belief. For reasons for or against hoping, or hypothesising, or being angry will also count as epistemic reasons – and thus as epistemic normative concepts – if they are related to the concept of an evidence-based reason for belief in the way our criteria specify. The specific ways in which the conditions are fulfilled might vary from case to case. In the case of misplaced hope, the reason against hoping is epistemic in virtue of the fact that there is no evidence-based

reason for believing that the content of the hope at issue can ever be true. In the case of hypothesising, the reason against assuming that the enemy will attack on a certain side depends on there being no evidence-based reason for believing that it is true that the enemy will attack from that side. In the case of anger, the epistemic reason against holding that feeling depends on the fact that anger might get in the way of the believer appropriately responding to evidence-based reasons for believing that something is true, and so on.

The way I have suggested we understand the notion of an epistemic reason and, hence, that we go about identifying which concepts are *epistemic-normative* concepts presents some clear advantages. Firstly, it is compatible with thinking that there might be other kinds of reasons for belief – i.e. practical (moral, prudential, etc.) – which might be based on considerations other than those evidentially supporting the truth of the beliefs in question. Yet at the same time it restricts the focus of our attention in a way that is sensitive to the nature of traditional epistemological discourse. More specifically, it goes some way towards preserving the intuition that epistemic reasons and candidate epistemic normative concepts must bear some distinctive link to the notion of truth.

Secondly, it is neutral with respect to whether or not believing is a voluntary action or a voluntary activity of sort. Perhaps it is the case that we can *decide* whether or not to believe things. Epistemic reasons to engage in this kind of action, then, will be reasons for believing based on evidence in a way which is analogous to that in which prudential reasons for action, for instance, are reasons for action which are based on considerations of self-interest.

Thirdly, it neither takes side nor depends on the outcome of recent debates as to whether believing has a ‘constitutive aim’ and, if so, as to what this aim is (Shah: 2003, Velleman: 2000, Shah and Velleman: 2005, Wedgwood: 2002). Even philosophers who reject the idea that belief is constitutively regulated by truth will hardly deny that at least some reasons for belief are of the evidence-based kind. Thus they should have no quarrel with my suggestion that we take epistemic reasons to be reasons of this latter kind.

Fourthly and lastly, it allows us to make sense of the idea that epistemic reasons might apply to responses other than beliefs – in the way I have indicated.

While the notion of an *evidence-based reason for belief* can allow us to make significant progress in circumscribing the domain of epistemic-normative concepts, it is itself open to importantly different interpretations. To these there correspond very different interpretations of epistemic normativity more generally. In order to see this, I must say a bit more about the concept of evidence and about how the concept of evidence and that of an epistemic reason are related. I will turn to this task in the next two sections.

2. Evidence

In the previous section, I suggested that we take epistemic reasons for believing that P to be considerations in favour of believing P which are such that they at least partly appeal to what evidentially supports the truth of the content of the belief in question. I said this was to be understood as requiring that a claim as to there being an *epistemic reason* for believing that P entail a claim about there being *evidence* of some sort for the truth of P itself.

It is important at this stage to be clear about the notion of evidence I take to be at work here. ‘Evidence’ is a term of art in epistemology, the philosophy of science, and the philosophy of probability. And often disagreement between competing theories in these fields is mirrored by (or consists in) different interpretations of what the notion of evidence itself amounts to. At first, then, it might not seem such a good idea to link the notion of an epistemic reason to what is, in many respects, another philosophically controversial notion. However, it seems to me that we do not need to take sides on too many vexed questions in order for our enquiry to benefit from reflecting on the connection between epistemic reasons and evidence and that, on the other hand, the pay-offs of proceeding in this way are high.

The concept of evidence I will rely on is, arguably, one which seems to be at work both in ordinary and in much philosophical and scientific discourse. In what follows I will (1) exemplify the concept by providing a brief list of sentences in which it figures and (2) draw attention to (or fix) some of its central features. Consider the following sentences:

E1: That A's fingerprints are on the murder weapon is evidence for the truth of the proposition that A is the murderer.

E2: That there are dark clouds in the sky is evidence that it will rain soon.

E3: That the fossil record exhibits certain characteristics²⁸ is evidence for the truth of (the set of propositions which comprise) evolutionary theory.

E4: That Darwin's theory is true is evidence for the fact that animals were not created by God to fit the environment they live in.

Each of these sentences is of the form:

E*: that E is evidence for H.

The notion of evidence at work in these sentences is such as to display the following features. Firstly, E can be evidence for H even if H is false. In other words, *misleading* evidence is still to count as evidence. Secondly, whether E is evidence for H does not entirely depend upon nor is entirely fixed by anyone's non-factive mental states. More specifically, it does not depend on anyone's beliefs about E or H or their relationship. Thirdly, for E to count as evidence for H, E must be true.²⁹ We can express this second and third point by saying that evidence is *factive*. Lastly, it seems to me that the notion of evidence at work in E1 to E4 is compatible³⁰ with thinking that in cases in which E entails H, E is to count as evidence for H. One way to put this

²⁸ For instance the gradual changes in the morphology of fossils across different time-strata in rock formations, etc.

²⁹ This qualification entails that evidence of the kind at issue here is *propositional* – i.e. such that we can always refer to it by using 'that' clauses. See Williamson (1997: 724-5).

³⁰ Contrary to what Achinstein believes, see Achinstein (1978: 23).

is by saying that *deductive evidence* is still evidence of some sort. Some might disagree (Achinstein: 1978) on this point. If they do, I will have no quarrel with them. Not much of what I will have to say in what follows depends on defending this latter point.

Notice how the concept of evidence at work here – perhaps unsurprisingly³¹ – bears many striking similarities to the concept of a normative reason I have been operating with. Both the term ‘reason’ and the term ‘evidence’ can either be used to refer to a relation (the *reason relation* and the *evidential support* relation, respectively) or to one of its *relata* (what provides the reason and what counts as evidence, respectively). Both the reasons there are for a given response and the evidence there is for the truth of a belief can differ in direction and strength, and they can compete and be in conflict. In the case of epistemic reasons the *strength* of epistemic reasons for given belief is also at least partly determined by – and directly proportional to – the strength of the evidence there is for the truth of the belief in question.³²

A wide range of expressions can also be used to qualify – in similar ways – both reasons and evidence: of reasons and evidence it is sometimes appropriate to say that they are ‘sufficient’, ‘overwhelming’, ‘decisive’, ‘conclusive’ etc. Again, the expression ‘consideration in favour’ can often be used both to indicate the evidence there is for the truth of a given proposition and the reason for holding a particular belief, as in: ‘the fact that John’s fingerprints are on the murder weapon is a consideration in favour of believing that he is guilty of the crime’ and ‘the fact that John’s fingerprints are on the murder weapon is a reason to believe that he is guilty of the crime’. Similarly, we can speak of the evidence *an agent has* in contrast to the evidence *there is for an agent* (Schroeder: 2008, Lord: 2010), as well as of evidence being only *apparent* evidence, in a way which mirrors the use of these expressions to qualify normative reasons. Lastly, some of the core features of the concept of evidence I have drawn attention to above are shared by the notion of

³¹ Given that some philosophers think that the concept of evidence and the concept of an epistemic reason are one and the same concept. See my discussion in the next section.

³² I don’t say ‘entirely’, because if one believes that epistemic reasons also depend on the pro-attitudes of agents (as in Humean Instrumentalist readings of epistemic reasons), the strength of these reasons will also presumably depend on the strength of the pro-attitudes they depend on. See the next section.

a normative reason which I have been working with. In the sentences I have listed above, that which provides the evidence – like that which provides reasons on my understanding of reasons – is *propositional* in form. And reasons in the sense I have been discussing are also ‘factive’ in the sense that they are not completely fixed by the non-factive mental states of agents and that in order for the proposition that F to be a reason for response S, F must be true.

It is precisely because of these strong links between the two concepts – I claim – that the notion of evidence is particularly useful when thinking about epistemic reasons.

In making the concept of evidence I am relying on more precise, I have had to take sides on some controversial and debated issues in epistemology. ‘Evidential internalists’ (and some kind of internalists about justification more generally), will deny that evidence is factive in the way I have assumed.³³ Similarly, the claim that (all) evidence is propositional in form is a disputed one and has been the object of some disagreement.³⁴

On the other hand, I did not take sides on other pressing questions surrounding the understanding of the notion of evidence itself, such as what the *having* or the *possessing* of evidence amounts to; or how the notion of evidential support itself is best understood; whether or not, for instance, it can be analysed in terms of probability (Kelly; 2008: 11-12); and if it can, what understanding of probability would be relevant to shed light on the evidential support relation itself.

It seems to me that the value of the present investigation is not compromised by those assumptions and omissions. Versions of the distinctions I am going to discuss below can be reformulated by taking different notions of evidence (and of reasons) into account, while nothing of what I will say in what follows depends on settling the issues I have left open. With those provisos in mind, we are now ready to see how thinking about the relation between

³³ See, for instance, Bonjour (1999) and Audi (2001) for representative positions of this kind. See, on the other hand, Williamson (2000) and McDowell (1982, 1995) for representative evidential externalist views.

³⁴ See, for instance, Williamson (2000: 194-200) for a defence of the claim that evidence is propositional in form, and Plantinga (1993) and Conee and Feldman (2004) for a defence of the claim that at least some evidence is non-propositional.

evidence and epistemic reasons is instrumental to bringing to the fore very different interpretations of epistemic reasons.

3. Three Concepts of Epistemic Reasons

3.1. Epistemic Reasons as Evidence-Based Reasons: Some Conceptual Distinctions

Borrowing an expression from Parfit (2011:503), I will use the label ‘alethic concepts’³⁵ to refer to the concept of evidence (in the sense I have specified above) and to other concepts which are used equivalently to it. Arguably, expressions such as ‘F makes P (more) likely’, ‘F makes P (more) probable’³⁶ or ‘F supports P’ or ‘F confirms P’, and the like are often – in some contexts at least – used interchangeably with expressions such as ‘F is evidence for P’. When they do, they express what I am calling ‘alethic concepts’. I also will call ‘alethic properties’ the properties which are referred to by alethic concepts.³⁷ Similarly, I will take ‘alethic facts’ to be the kind of facts which are referred to by sentences which contain terms expressing alethic concepts.

Establishing the link between the notion of an epistemic reason and what I have called alethic concepts affords one way of making explicit the link between the kind of reasons we are interested in and the notion of truth. To have evidence of the kind I have indicated for a belief or proposition is to have evidence for their truth; and reasons for belief which depend on evidence are thus sensitive to considerations which bear on the truth of beliefs themselves. The link in question also allows us to capture an uncontroversial conceptual relation between epistemic reasons and some (other) key epistemic concepts (i.e. alethic concepts).

³⁵ One might think this choice of words to be a bit infelicitous given the close association of the adjective ‘alethic’ (via its etymology) to the notion of truth. In the sense of ‘alethic’ which I will be adopting, the notion of truth would not count as an *alethic* concept, and this is – admittedly – strange. I am merely following Parfit here, and adopting a terminology which has been used in the context of a discussion in many respects similar to the one I am about to follow.

³⁶ That the expression ‘F makes P (more) probable’ can – in certain contexts – be used as equivalent to ‘F is evidence for P’ need not entail that – indeed leaves it quite open whether – evidence is to be ultimately explained in terms of probability.

³⁷ I am not assuming at this stage that ‘alethic concepts’ are genuinely referring concepts. Only that when they do refer, they refer to properties I call alethic, etc.

More importantly still, it provides a common departure point and a shared ground between distinct and competing accounts of the nature of these reasons.³⁸ It is thus instrumental to drawing attention to very different ways in which the concept of an epistemic reason can be understood. They correspond to different ways to construe the relation between what I have called ‘alethic concepts’ and the concept of an epistemic reason itself. In what follows, I will turn to the task of illustrating these different interpretations. In doing so, I will rely on the alethic concept of evidence both for reasons of convenience and because of the central role this concept has played in my initial characterization of epistemic reasons themselves.

In Section 1 of this chapter I have suggested that we understand epistemic reasons as reasons which are based on – or depend on – evidence, and I have used the expression *evidence-based reasons* as an alternative way to refer to reasons of this kind.

Consider now the claim below:

ER: F* is an epistemic reason for A to believe that P in circumstances C.

where F* is shorthand for the fact or facts which we take to provide sufficient grounds for there to be an epistemic reason for A to believe that P in C.

The understanding of epistemic reasons I am operating with, as we saw, assumes that (1) ER conceptually entails that there is evidence for A to believe that P and that (2) among the facts which comprise F* is some fact or facts E which is such as to be or provide evidence for the truth of P.

One way to put this is by saying that if we were asked: ‘what epistemic reasons are there (for A in circumstances C) for believing that P?’ ‘The fact that E’ will have to be (at least) *part* of the correct answer. The truth of (1) and (2) above, however, leaves some important issues open, and

³⁸ For instance, some philosophers think that epistemic reasons are dependent upon an agent’s having cognitive goals of a certain kind, others deny this. But both camps would be happy to agree that what they are talking (and disagreeing) about are reasons which are based on evidence in the way I will specify.

it is here that important conceptual disagreements about epistemic reasons can come to the fore. One way to bring to light what the disagreement in question amounts to is by asking: *what else* – if anything at all – must enter the grounds of the (epistemic) reason relation in order for one to have a *sufficient* specification of these grounds? In what follows I discuss three very different ways of answering this question.

3.1.1. The Instrumentalist view of Epistemic Reasons (Epistemic Instrumentalism)

The instrumentalist view of epistemic reasons takes epistemic reasons to be a kind of instrumental reasons. This view affords a distinctive way of understanding the relation between alethic concepts and the concept of an epistemic reason by appealing to the conceptual structure of instrumental reasons claims.³⁹ The core idea in epistemic instrumentalist approaches of the kind I am considering is that believing what there is evidence for is *instrumental* to achieving some specifiable goal or end, where the goal or end is typically a *cognitive* one – such as the goal of acquiring true beliefs, or justified beliefs, or knowledge, etc.⁴⁰ – although it need not be.

Within broadly instrumentalist approaches to epistemic normativity it is helpful to distinguish two varieties which I will refer to as *Humean* and *Non-Humean* Instrumentalism respectively.

According to the first (Humean Instrumentalism), if there is an epistemic reason F* for A to believe that P, a full specification of the grounds for such reason will have to include: (1) some fact E which is evidence for the truth of P and (2) the fact that agent A has a contingent pro-

³⁹ The view I am calling ‘epistemic instrumentalism’ bears obvious resemblances to views which are sometimes discussed in the philosophical literature under the heading of ‘epistemic teleology’ or ‘epistemic consequentialism’ (Berker: 2013a and 2013b). Although it is most natural for proponents of views which I am calling ‘epistemic instrumentalist’ to be also committed to ‘epistemic consequentialism’, I believe the two views to be distinct. I do not here have the space to discuss this matter at any length and with the attention it would deserve.

⁴⁰ One might think that on this view – given that epistemic-reasons can be entirely assimilated to instrumental reasons – they will end up being practical *reasons* of some sort. I think this is mistaken. On some varieties of non-Humean instrumentalism, the reasons we have for believing what we have evidence for are, ultimately, provided by the fact that *knowledge is an intrinsically valuable state to be in*. It is not clear why a goal of this kind should count as a *practical* one. If it isn’t, then it is not clear why the reasons which are instrumental to achieving this end should count as practical *reasons* of some sort either.

attitude of some kind (i.e. a desire) towards some goal or state of a specifiable kind⁴¹ and (3) the fact that the goal or state in question is (likely) to be achieved by A, if A were to form the belief that P (i.e. the belief for which E is evidence). The goal A has a contingent pro-attitude towards can, but need not be a cognitive goal.

On some variety of Humean instrumentalism, for instance, A will have a reason to believe that P only if *A desires to believe that P if P is true*. On others, A's reason to believe that P will depend on the fact that believing what there is evidence for is instrumental to acquiring true beliefs, and that achieving this latter goal is itself instrumental to furthering some non-cognitive goal A might contingently desire to pursue: such as the goal of maximising his overall desire-satisfaction. In this latter case, we can say that the instrumental reasons to believe what one has evidence for ultimately depend on the satisfaction of a non-cognitive goal.

Philosophers who hold the Humean instrumentalist view of epistemic reasons believe that epistemic-reasons will have the following structure:

HI: [The fact that X-ing is a means to Y-ing and that A has a pro-attitude towards Y-ing] provides a reason for A to X.

According to instrumentalists of the Humean variety, if there is an epistemic reason for A to believe that P, A will not only have evidence E for P, but he will also be such as to value some goal G (cognitive or non-cognitive) which would be 'furthered' by his forming the belief that the evidence supports. When Humean instrumentalists are asked to provide a full-specification of the grounds which are sufficient for there to be an epistemic reason for A to believe that P, they

⁴¹ For instance: to form a true belief as to whether p or not, to have knowledge of whether p, to form a justified belief as to whether p, etc.

will mention the fact that E is evidence for P, the fact that believing what there is evidence for is instrumental to achieving G, and the fact that A has a contingent pro-attitude towards G.

Philosophers like Sharon Street, David Papineau, and Robert Nozick (among many others) can be read as holding views that fall under this broad kind.⁴²

Non-Humean instrumentalist about epistemic normativity will instead subscribe to an analysis of epistemic reasons according to which these reasons have the following structure:

NHI: [The fact that X-ing is a means to Y-ing and that A has a reason for Y-ing] provides a reason for A to X.

where the reason for Y-ing is independent of the presence in agent A of a contingent pro-attitude towards Y-ing itself. Non-Humean instrumentalist views will differ with regards to the account they provide of the grounds of the reason for Y-ing. They will also differ with respect to their specification of the goal evidence-based reasons would be instrumental to achieving. For instance, some will think that we have *moral reasons* to pursue *knowledge* or truth (Clifford: 1877);⁴³ others that *knowledge* is valuable because of the role it plays in practical reasoning (Hawthorne: 2004); or that there are *prudential reasons* for forming *as large a number as possible* of true beliefs; yet others that truth (Zagzebski: 2003) or knowledge, or understanding (Kvanvig: 2003), etc. are valuable in themselves, and that this fact provides reasons for forming beliefs in such a way that the beliefs we form are likely to exemplify those properties. When non-Humean instrumentalists are asked to provide a full-specification of the grounds which are sufficient for there to be an epistemic reason for A to believe that P, they will mention the fact that E is evidence for P, the fact that believing what there is evidence for is instrumental to achieving a certain (cognitive or

⁴² See the discussion in Kelly (2003) and Berker (2013a and 2013b) for a much longer list of philosophers which can be read as subscribing to epistemic-instrumentalist views of the kind I am discussing here.

⁴³ See Cuneo (2007: 79): 'I think therefore there is a sense in which one can say that epistemology is a branch of ethics'.

non-cognitive) goal G, as well as the fact that there are pro-attitude-independent reasons for A to achieve the goal in question.

3.1.2. The Evidentialist View of Epistemic Reasons

Evidentialists about epistemic-reasons believe that although claims about evidence (such as ‘E is evidence for P’) and claims about epistemic reasons (such as ‘F* is an epistemic-reason to believe that P’) are conceptually distinct claims (i.e. they are not equivalent in meaning), claims of the former kind conceptually entail claims of the latter. Unlike instrumentalists about epistemic reasons, when evidentialists are asked to provide a full-specification of the grounds which are sufficient for there to be an epistemic reason for A to believe that P, they will not appeal to means/ends relations between believing what there is evidence for and achieving a certain goal G⁴⁴. They will merely appeal to the fact *that E is evidence for P*. This fact alone will be sufficient to secure the presence, for A, of an epistemic reason to believe that P. We may put this by saying that, according to the evidentialist, evidence is intrinsically reason-providing (and that the reasons evidence provides are epistemic reasons). If we use square brackets to mark off the grounds for the epistemic-reason (there is for A) to believe that P, philosophers who hold this view believe that:

EV: [The fact that E is evidence for P] is/provides an epistemic reason to believe that P.

Parfit defends an interpretation of the relation between (the concepts of) evidence and epistemic reasons along these lines in *On What Matters*⁴⁵ (2009: 202-10). The same view seems to be

⁴⁴ Unless this goal is that of *forming the beliefs there evidence for*. In this case evidentialism will collapse into a form of instrumentalism of some sort – in a similar way in which most theories which subscribe to the intrinsic value of some property F can be turned into instrumentalist theories of some kind or other, when the means/ends relation is construed as including the instantiation relation.

⁴⁵ There Parfit says that epistemic reasons are provided by/grounded in what he calls *alethic* facts, i.e. facts about evidence or about what is likely to be true.

implicitly endorsed by Achinstein (1978: 27)⁴⁶ and, arguably by Conee and Feldman (2004: 257-8) and Berker (2003a) among others.

3.1.3. Evidence-as-Reasons Interpretations: Analytical Epistemic Naturalism and the Evidence as Irreducibly Normative Views

There is yet another way to understand the relation between epistemic-reasons and evidence which is compatible with the claim that epistemic reasons are evidence-based reasons in the sense I have specified. I will call this the ‘evidence-as-reasons’ interpretation of epistemic-reasons. According to this view, the concept of evidence and the concept of an epistemic reason are *equivalent in meaning*. In other words, saying that E is evidence for P and saying that E is/provides a(n) (epistemic) reason to believe that P amounts to saying one and the same thing. Thomas Kelly (2006) seems to have an understanding of the relation between evidence and epistemic reasons of this kind when he writes:

The concept of evidence is closely related to other fundamental normative concepts such as the concept of *a reason*. Indeed, it is natural to think that ‘reason to believe’ and ‘evidence’ are more or less synonymous, being distinguished chiefly by the fact that the former functions grammatically as a count noun while the latter functions as a mass term. (Kelly: 2006)

Unlike evidentialists, proponents of the evidence-as-reasons interpretation of epistemic reasons will not believe that the concept of evidence and that of an epistemic reason are distinct. Thus they will say that the facts which provide the grounds for the epistemic reason to believe that P are *the same facts* which evidentially support the truth of P.

In other words, if proponents of the evidence-as-reason view are asked for a full specification of the grounds which are sufficient for there to be an epistemic reason for A to believe that P, they will say that *fact E* is the reason (for A) to believe that P. If we use square brackets to mark off the grounds for the epistemic-reason (there is for A) to believe that P, philosophers who hold this view believe that:

⁴⁶ See Achinstein’s ‘principle of reasonable belief’ (1978: 27).

ERV: [fact E] is/provides an epistemic reason to believe that P.

We can bring to view the crucial difference between evidentialists and evidence-as-reasons theorists here by saying that the evidence-as-reasons theorists believe – while evidentialists deny – that facts about evidence and facts about epistemic reasons are identical. To illustrate this point on behalf of what I am calling the ‘evidentialists’ Parfit writes:

The alethic concept [*is evidence for*] is quite different from the normative concept *gives us a [...] reason to believe*. (2011: 504)

According to evidentialists, what follows from a conceptual analysis of the notions of (epistemic) reasons and evidence is that epistemic-reason facts and evidence-facts are so related that (1) epistemic-reason facts are distinct from evidence facts and (2) epistemic-reason facts are necessarily and exclusively grounded in evidence-facts themselves.⁴⁷

Let us now go back to the evidence-as-reasons analysis of epistemic-reasons. Believing that the concept of an epistemic-reason and the concept of evidence are *equivalent in meaning* – or as Kelly puts it ‘more or less synonymous’ – is compatible with holding different views as to whether either of these concepts could be appealed to in order to *explain or analyse* the other.

It is, for instance, possible to believe that the concept of evidence is *not an irreducibly normative concept*. One way to do so is – for instance – by believing that the notion of evidence is to be understood in terms of the notion of objective probability (where this notion is not itself understood in irreducibly normative terms). The notion of an epistemic reason could then, in turn, be analysed in terms of the notion of evidence thus understood. According to this view, all that we would be saying when we say that E is an epistemic reason to believe that P is – for instance – that given that E, the probability that P will occur has risen above a certain threshold. If we take claims

⁴⁷ It is, of course, possible to agree with such an analysis of the relation between epistemic-reason facts and evidence-facts and to believe that *there are no epistemic-reason facts*. In other words, it is possible to be an error-theorist about epistemic-reason claims thus understood. This position will be similar to the position of error-theorists about morality who – for instance – believe that moral concepts are *categorical* and *irreducible* to concepts of other kinds, but that they fail to refer to anything.

about probability to refer to some kind of natural facts,⁴⁸ a view of the kind I have just sketched will be an instance of what Parfit calls ‘analytical epistemic naturalism’ (2011: 509). A philosopher (among many perhaps) who endorses a view of this kind is Schiffer (2003); Lenman (2008) also discusses the possibility of occupying a position of this kind with regard to the analysis of epistemic reasons.

But it is also possible to believe that evidence *is an irreducibly normative concept*; for instance by believing that any analysis of evidence which is not provided, ultimately, in terms of reasons for belief, will fail to do justice to central features of this concept. We might call this position the ‘evidence as irreducibly normative’ view. Horwich (1982) and Kelly (2007) seem to subscribe to a conceptual analysis of epistemic reasons/evidence along those lines. Believing that *the concept* of evidence is irreducibly normative in this way leaves it open whether or not this concept refers to *irreducibly normative facts*. For one could believe that evidence is irreducibly normative in this way and yet believe that claims about evidence refer to claims about probability of the kind I have exemplified above. A position of this latter kind would be similar to non-analytical versions of naturalism in ethics.

4. Importance of the Distinctions for the Project of Locating Normativity in Epistemology

My discussion of the conceptual relation between the notion of evidence and that of an epistemic reason was designed to bring to light some often neglected issues surrounding the use of normative concepts in epistemology. In particular, it was meant to illustrate that there are very different ways of understanding the concept of an epistemic reason which are compatible with my initial characterization of epistemic reasons as *evidence-based reasons*.

⁴⁸ It does not matter for the sake of the point I am making here that we have a precise understanding of what facts are natural in the relevant sense at work here. Let us just assume that whatever natural facts are taken to be, they will not be irreducibly normative facts.

It is customary to find paradigmatic instances of claims about epistemic reasons formulated in the following terms:

E2: The fact that John's fingerprints are on the murder weapon is a reason to believe he is the author of the crime.

or

E3: The fact that Jane is a heavy smoker is a reason to believe that she will not be healthy for long.

But we are now in a position to see that when this is the case, agreement on the truth of these and similar claims can conceal significantly different views regarding the nature of the reasons which are being talked about. Bringing these differences to light is important. This is because to each of the distinct interpretations of epistemic reasons I have reviewed, there will correspond very different ways of understanding whether a concept is normative *in the epistemic sense* and what the epistemic normativity of *epistemic normative* concepts will consist in. Settling those issues in an unambiguous way will depend not only on establishing (1) whether concept C bears to the concept of an *evidence-based reason* the kind of relations which either of the conditions I have provided specifies. But it will also depend on being clear about (2) how the concept of an *evidence-based reason* is understood. Suppose, for instance, one believes that the following conditional expresses an essential feature of the concept of epistemic rationality:

ER: If agent A is epistemically rational, then A is (in normal conditions) disposed to respond correctly to the epistemic (i.e. evidence-based) reasons he attends to.

Clearly this makes the concept in question normative (it fulfils Condition I - Epistemic). Yet the understanding of this normative concept will differ significantly, if for instance, one is a *Humean*

instrumentalist about epistemic reasons rather than an *analytical naturalist* about the same concept. In the first case, epistemic rationality will be an instance of (Humean) instrumental rationality, and the task of establishing whether an agent is epistemically rational or irrational will require ‘looking into’ the set of contingent pro-attitudes which he has at any given time. In the second case, on the contrary, this task will be independent of taking into account those elements of the psychological makeup of the agent.

Analogously, one might agree in believing that the concept of epistemic justification is normative because one believes that the following bi-conditional states an essential truth about this concept:

EJ: Belief P is justified iff there is a strong/sufficient epistemic (i.e. evidence-based) reason to believe that P.

Here, as in the case of epistemic rationality above, the understanding of what the normativity of this concept amounts to will vary significantly depending on the account of epistemic reasons one favours.

Suppose instead that one believes that the notion of epistemic justification is such that the following bi-conditional is true:

EJ*: Belief P is justified for agent A iff there is strong evidence for A in support of the truth of P.

Will the truth of EJ* entitle one to believe that the concept in question is an *epistemic normative* concept? This will again depend on the view of epistemic reasons one adopts. For instance, on either the *Analytical Epistemic Naturalist* and the *Evidence as Irreducibly Normative* views, the concept of justification at work in EJ* will be a normative concept. That P is justified conceptually entails

that there is evidence for it and evidence claims are conceptually equivalent to epistemic reasons claims on either of these views. Thus the concept of justification fulfils Condition II – Epistemic and it qualifies as an epistemic *normative* concept. Justification will similarly be a normative concept on an *Evidentialist* reading of epistemic reasons. This is so if entailment relations among concepts are construed as transitive. For claims about justification entail claims about evidence (via EJ*) and – if we assume that evidentialism is true – claims about evidence will entail claims about reasons. On the other hand, if we consider once again a Humean instrumentalist reading of epistemic reasons, the truth of EJ* will *not* entitle one to believe that justification is a normative concept. For, on this view, claims about evidence alone do not license claims about reasons. For them to do so, as we saw, one must add that the agent whom the evidence is for, has a pro-attitude towards some goal or other which is instrumentally served by believing what he has evidence for.

On the blueprint of the cases I have examined above, it is easy to see how one might go about formulating and settling questions regarding the epistemic normative status of many other central concepts employed in epistemology, such as those of evidence itself, probability, confirmation, entitlement, warrant, reliable method of belief-formation, epistemic virtue(s), and the like.

Separating out the differing views on the nature of epistemic reasons will not only be relevant to addressing conceptual questions of the kind I have discussed. It will also be clearly relevant to evaluating meta-normative debates in epistemology. For it will make perspicuous how there can be different kinds of things about which one can occupy each of a variety of distinct meta-epistemological views. Because of this, it will help in the evaluation of competing views by guiding us towards the sort of considerations that might be relevant for assessing them.

Epistemic realists⁴⁹ who subscribe to *Humean Instrumentalism*, for instance can be expected to support their views by very different arguments from the ones employed by epistemic realists who understand epistemic reasons along *Evidentialist* lines. In light of this fact, arguing against the first rather than the second view will require very different strategies. Analogously, assessing realist views about epistemic normativity on the third kind of reading I have provided will – presumably – require discussing the nature of facts about evidence in a way in which assessing *Instrumentalist* or *Evidentialist* views of Normativity will not. But we can hardly see this, unless we are clear about the fact that the views I have just mentioned would count as normative realist views *of significantly different kinds*.

How different interpretations of epistemic reasons interact with different meta-normative views about those reasons will become clearer in the next chapter. In it I will be occupied with charting a map of the largely unexplored field we may call ‘meta-epistemology’ and with discussing realism about epistemic normativity in some detail.

⁴⁹ The definition of epistemic realism I will operate with does not require versions of this view to be committed to the *objectivity* of epistemic-normative facts; hence, the possibility of a *Humean-instrumentalist* version of epistemic realism. I return to discuss this topic in the next chapter.

Chapter 3

Varieties of Normative Realism in Epistemology

1. Introduction

In my discussion so far I have been mainly occupied with conceptual issues. Firstly, I provided a general criterion for the identification of normative concepts in which the notion of a normative reason played a crucial role. I did not take a stance on whether the concept of a normative reason was to be regarded as ‘most basic’ or ‘uniquely primitive’. Nonetheless, I argued, focussing our discussion on this concept had obvious advantages. If the concept of a normative reason is *most basic*, this would be the obvious thing to do. If it turned out not to be, thinking about reasons can still provide the tools for thinking about other (more) primitive, or equally primitive, normative concepts.

Secondly, I defended and adopted a strategy for indicating which normative reasons are to count as *epistemic* reasons. I suggested that a promising way of thinking about epistemic reasons is in terms of the notion of an *evidence-based* reason. I then reviewed different interpretations to which the concept of an epistemic reason thus understood is open. I did so by showing that there are very different and incompatible ways of understanding the notion of a reason to believe which is *based on evidence*. Finally, I used the general criterion for the identification of normative concepts which I had provided and the concept of an epistemic reason to formulate a criterion for singling-out the sub-class of normative concepts which are to count as *epistemic-normative* concepts.

The whole point of the investigation thus far has been that of coming to a better understanding of how normative concepts apply to – and come into – epistemic discourse. This

is a task that can hardly be neglected if one is to formulate and evaluate arguments against realist views of epistemic normativity: before doing so, we must have a clear sense of *what* epistemic realists are being realists *about*. To be sure, and as crucially, we must also have a clear sense of what the *realism* at issue amounts to. This is the question I will address in the present chapter. I will begin by outlining a ‘rough guide to meta-epistemology’. This will provide a first characterization of what epistemic realism is and what it is to be contrasted with. I will then turn to look at epistemic realism in more detail and distinguish different varieties of it. All along, I will be often relying on – and borrowing from – vocabulary and distinctions which are common currency in much recent and contemporary meta-ethical literature.

Drawing on the work of the previous chapter, it will perhaps come as no surprise to find that epistemic realist views of various kinds – as well as a number of other meta-epistemological positions – can be occupied ‘at different places’. There are different things one can be an epistemic realist (or anti-realist, etc.) about, and these will correspond to the different interpretations of epistemic reasons which I have distinguished. Once we appreciate this, we will be finally in a position to outline and evaluate an argument that relies on ‘Darwinian considerations’ in order to call into question realist views in meta-epistemology. More specifically, we will be in the position to be clear about *which* epistemic realist views, if any, are allegedly put under strain by those considerations. We will also know where to look in order to evaluate whether this is indeed the case.

2. Meta-Epistemology

The variety and complexity of positions that can be occupied with regard to meta-normative issues in epistemology makes it very hard not to draw lines which might appear arbitrary from some perspective or other, or not to label as distinct positions that – were other criteria to be adopted – would fall under the same heading. To this extent, and unsurprisingly, meta-

epistemology mirrors the cognate field of metaethics. The map I will draw here is designed to remain as close as possible to uncontroversial uses of key philosophical terms. At the same time, it will be instrumental to bringing to the fore issues that, in my opinion, have hardly been articulated in discussions of epistemic normativity so far. Much of the assessment of the impact that Darwinian considerations might have on the tenability of positions in meta-epistemology will rest upon making these distinctions clear. In what follows, I will restrict my attention to the concept of an epistemic reason.⁵⁰ I will take paradigmatic epistemic sentences to be first order,⁵¹ positive, and atomic sentences of the kind:

ER: Fact F* is a reason for A to believe that P (in circumstances C).

The classification I will adopt divides the territory of meta-epistemology into views that can be distinguished from one another with respect to their acceptance or rejection of the following four theses:

Psychology Thesis (PT): At least some epistemic sentences are the expression of beliefs.

Semantic Thesis (ST): At least some epistemic sentences are truth-apt.

Alethic Thesis (AT): At least some epistemic sentences are true.

Existence Thesis (ET): There are epistemic facts.

On the basis of these four theses, we have a way of classifying positions in meta-epistemology on the blueprint of what are recognizably familiar positions in metaethics. One would be hard put to

⁵⁰ Much of what I will have to say below, however, also applies to ‘epistemic sentences’ of other kinds. Take concepts which are epistemic-normative in virtue of satisfying the first of the conditions I have specified. If we use ‘epistemic term’ to refer to a word or expression that expresses an epistemic normative concept in this sense; and ‘epistemic sentence’ to refer to a sentence that contains at least one epistemic term, then much of the discussion to follow can be easily extended to encompass first-order, positive, and atomic epistemic sentences of these other kinds.

⁵¹ In the sense that epistemic sentences of the kind at issue do not feature in ‘embedded contexts’ of various kinds, such as belief-contexts or claims in which reasons are themselves said to be sources of reasons, etc.

find matches between the meta-epistemological positions that *can* be occupied with respect to each of the four theses above and those which have *actually and explicitly* been occupied by philosophers writing about normative issues in epistemology.⁵² While this provides a clear indication of the extent to which the field of meta-epistemology has been neglected, it ought not to prevent one from acknowledging the varieties of positions which are also available here.

One way of drawing the distinction between realist and anti-realist views in epistemology – the one which I will adopt here – takes the realist/anti-realist divide to amount to a fundamental disagreement with regard to the existence of epistemic facts.⁵³ In what follows, I will consider acceptance of the Existence Thesis⁵⁴ to be a necessary and sufficient condition for a view to qualify as ‘realist’ in the sense in which I am using this expression. Rejection of this thesis, on the other hand, will be a necessary and sufficient condition which all the views I will call ‘anti-realist’ must be taken to satisfy. Both realism and anti-realism thus construed are best seen as families of views. Each can come in many varieties, and one way of distinguishing those versions of realist and anti-realist views will be on the basis of whether they accept or reject the truth of each of the Psychology, Semantic and Alethic Theses. We can assign names to these views by redeploying vocabulary that is already in use in traditional metaethical taxonomies. The references I will provide in brackets are to philosophers who have occupied positions in metaethics analogous to the ones I am allowing for in meta-epistemology. To be sure, ethics and

⁵² Philosophers who see themselves as subscribing to *naturalist* programs in epistemology are quite often oblivious to the fact that there are different ways of being naturalists about key epistemological concepts or paradigmatic epistemological facts. This is so on each of the many ways in which the property of ‘being natural’ has been understood. I take it that one of the outcomes of the discussion of this chapter is that of making this important point perspicuous.

⁵³ A different and alternative way of drawing the realist/anti-realist distinction is common currency in much metaethical literature. It consists in adopting as the defining feature of moral realist views their commitment to the thesis that moral facts are ‘objective’ or ‘attitude-independent’ in some to-be-specified sense. My definition of epistemic realism says nothing about the ‘objectivity’ or ‘attitude-independence’ of epistemic-normative facts. In this respect, it departs from one standard way in which the term ‘realism’ is used in the context of many metaethical discussions. I will return to discuss this point in the next section.

⁵⁴ Standard realist views in metaethics will accept the moral analogues of *all* the four theses which I have listed above. Thus it is perhaps natural to expect that most versions of epistemic realism would also be committed to belief in the Psychology, Semantic and Alethic theses. Given my definition of epistemic realism, commitment to these theses will be a *contingent* feature of epistemic realist views. One advantage of defining epistemic realism in this way is that it allows for the possibility of (epistemic) ‘realist expressivist views’; see below and Copp (2011a).

epistemology are not the same thing. Thus it would be a mistake to assume automatically that philosophers who are anti-realist of one kind or another about practical and moral reasons will also be anti-realist about epistemic reasons. In what follows, I will not do so. Nonetheless, providing such references can help shed light on the nature of the meta-epistemological positions I will outline. It can act as a pointer, so to speak, towards some of the ways in which the details of these positions might be filled out. This, in turn, might provide – at least in some cases – an indication of the reasons there might be for wanting to hold one position rather than another.

While, on my taxonomy, all anti-realist views will reject the existence thesis, it is possible for anti-realists to hold different attitudes with respect to the psychology, semantic, and alethic theses. Some anti realist positions – error theories or nihilist theories – will accept the Psychology and Semantic theses and reject the Alethic Thesis (Mackie: 1977), as will some ‘fictionalist’ accounts.⁵⁵ Others – traditional expressivist views – will reject all of these three theses (Ayer: 1936). Yet others views will be Cognitivist and Anti-Realist (these views are sometimes referred to as ‘cognitivist irrealist’) and will consist in the acceptance of each of the Psychology, the Semantic, and the Alethic theses (Wright: 1995, Skorupski: 1999 and 2002); while ‘quasi-realist’ forms of expressivism (Blackburn: 1993, Gibbard: 2003) will accept each of the Psychology, Semantic, and Alethic theses on an appropriately ‘minimalist’ and ‘deflated’ understanding of the notions of ‘belief’, ‘truth-aptness’ and ‘truth’ that figure in each.⁵⁶

Paradigmatic forms of realism will consist in the acceptance of each of the four theses above. Nonetheless, more ‘hybrid’ forms of realism have been defended in metaethics, and it is plausible to think that correlate positions can be occupied in meta-epistemology. These views – realist and

⁵⁵ This will be those varieties of fictionalism which claim that epistemic sentences are assertoric, that they express beliefs, but they are not literally true: rather, they have some other surrogate property such as that of being ‘true-in-the-moral-fiction’. Not all fictionalists need be committed to this analysis. A different brand of fictionalism about epistemic discourse, like forms of fictionalism about different areas of discourse, will consist in thinking that epistemic sentences do not express beliefs at all but rather some other attitude such as ‘make-belief’ or ‘pretence’ (Walton: 1990; 1993), or ‘acceptance’ (VanFraassen: 1980).

⁵⁶ On an appropriately minimalist understanding of the notion of fact, these views can also be thought of as accepting a version of the Existence Thesis: we might call this weaker thesis the ‘Quasi-Existence’ Thesis.

expressivist – can be seen as combining acceptance of the Existence Thesis with rejection of the Psychology, Semantic, and Alethic theses. They hold that there are normative facts, while also believing that epistemic sentences with the surface form of sincere assertions typically express non-representational (conative) states of mind (Copp: 2011) and that, because of this, they are neither truth-apt nor (a fortiori) true. Varieties of realism can also be distinguished according to their differing views on the nature of normative facts. I will return to discuss this topic at some length in the next sections. Before doing so, however, I want to say something about the way I have chosen to characterise normative realism in this chapter.

3. A Note on the Definition of Normative Realism

3.1. Different Definitions of Normative Realism

One important thing to notice about the taxonomy I have adopted is that it does not count ‘constructivist’, ‘subjectivist’, or ‘response-dependent’ views of various kinds as falling automatically on the anti-realist side.⁵⁷ While my way of carving up of the distinction between realism and anti-realism is not unusual⁵⁸ (Robertson: 2009, Sayre-McCord: 1988, Fitzpatrick: 2008), it differs from an alternative way of proceeding which is also standard in much metaethical literature. Sharon Street characterises these distinct ways of getting at the distinction as follows:⁵⁹

On one common understanding, realism about morality and normativity more generally is the view that (1) moral or normative claims are capable of being true and false; and (2) some of these claims are true. If we understand realism this way, then metaethical constructivism counts as a brand of realism – as indeed do an extremely wide range of views, including even a simple subjectivism according to which what’s good for a person is whatever that person thinks is good. But it is also fully in

⁵⁷ See Robertson (2009) for a similar way of carving the distinction at hand. Street also acknowledges the viability of drawing the realist/realist distinctions in these terms in (Street 2010 and 2009), although this is not the way Street herself chooses to frame it.

⁵⁸ In the sense that it bears significant resemblance to – and shares ‘the spirit’ of – one standard way of drawing this distinction. The fact that I rely *exclusively* on the Existence Thesis to define realist views is – I admit – somewhat unusual. I have already indicated some reasons as to why this seems to me a good way to proceed.

⁵⁹ For a much more detailed discussion of this and related issues see Fitzpatrick (2008).

keeping with philosophical tradition to understand the ‘realism/anti-realism’ debate as concerning more than this, and in particular as concerning the question of *mind dependence*. [...]. On this understanding the key point at issue between realist and anti-realist is the answer to the central question of Plato’s *Euthyphro* (in rough secular paraphrase), namely whether things are valuable ultimately because we value them (antirealism), or whether we value things ultimately because they possess a value independent of us (realism). In the final analysis, in other words, is normativity best understood as conferred or recognised? (Street: 2010: 70)

It is easy to spot the similarity between the first of the ways of understanding normative realism which Street discusses in the quotation – i.e. in terms of the thesis that *at least some normative beliefs are true* – and the way I chose to define normative realism – i.e. in terms of the thesis that *there are normative facts*. And one might wonder whether the alternative way Street mentions would not be a better way to proceed in the context of an enquiry of the kind I am engaged in here.

The idea of the *mind-dependence* of normativity which Street articulates in the quote above is amenable to two very distinct kinds of readings. On the first reading, it is a claim about normative concepts and what they entail. On the second, it is a claim about the nature of normative facts. On either of those readings – it seems to me – there is something to be said against using the notion of mind-dependence as a tool for carving up the distinction between realism and anti-realism about normativity. I cannot discuss this matter at any length here. In order to motivate my way of proceeding, suffice it to say the following. I believe that many of the questions that are traditionally asked in order to classify forms of normative realism into various kinds – such as the question whether normative facts are *causally efficacious*, or whether they are *reducible/irreducible* to other kind of facts – can also legitimately be asked when facts of this kind are understood to be *subjective* (i.e. *mind-dependent*) in many of the ways in which this latter notion can be plausibly spelled out.⁶⁰ And it is an advantage of a taxonomy that does not build *objectivity* (i.e. *mind-independence*) into the definition of normative realism that it can contribute to make this fact perspicuous. In the end – as will become apparent – nothing substantive about

⁶⁰ This will become apparent from my discussion in the next sections of this chapter.

what I have to say in what follows will depend on adopting the specific way of defining normative realism which I adopt.

3.2. Normative Realism and Normative Facts

There is another way in which my characterization of realist and anti-realist views about normativity may be thought to be less than satisfactory: it is somewhat coarse-grained and it does not do full justice to the complexity which can sometimes accompany the correct classification of some meta-normative positions. For instance, where the line between realist and anti-realist views is drawn will depend crucially, in some cases at least, upon how the notion of ‘fact’ in the Existence Thesis is understood.

A ‘robust’ reading of this notion (Copp: 2011), will yield a different carving of the distinction from a broader and more liberal interpretation of the same, i.e. one which is compatible with the thought that some of the things which can enter facts as constituents and which facts can be about exist in some ‘non-ontological’ sense (Parfit; 2011). Likewise, it can be hard to draw the distinction between the claim that some things exist in a ‘non-ontological sense’ and the claim that we can talk of certain kind of facts obtaining in a ‘minimalist’ understanding of ‘fact’, ‘existence’, and ‘obtaining’ (Gibbard: 2003, Blackburn: 1993). When this is the case – if one is to assess whether a specific view falls onto either side of the realist/anti-realist divide – the content of the Existence Thesis will have to be made more precise.

Rough as it may be however, the map I have sketched so far does afford an initial characterization of epistemic realism, while also determining its location with respect to other distinct meta-normative views in epistemology. In the remainder of this chapter I will turn to consider epistemic realism at some length and in greater detail.

4. Varieties of Epistemic Realism I: Irreducibility, Objectivity, Causality

We are familiar from traditional taxonomies in the metaethical literature with the idea that realism about practical normativity can take many forms. These are typically told apart on the basis of *what kind of facts* normative facts are taken to be. In particular, moral realists (and realist about practical normativity more generally) typically disagree with regard to whether they believe that moral facts are objective, or whether they are reducible to non-moral facts, or whether they are causally efficacious. And varieties of moral realism can be distinguished according to which side of each of these questions they come down to.

Different versions of moral realism can be expected to have counterparts in the varieties of forms which realism about *epistemic* normativity can take. Here too, we must allow for the fact that one important dimension alongside which epistemic realist positions can vary has to do with the *kind of facts* they take epistemic-normative facts to be. More specifically, of these epistemic normative facts we can ask: are they *irreducibly* normative? Are they *objective*? Are they *causally efficacious*? It is open to epistemic realists to answer each of these questions differently. Thus one way to go about classifying different kinds of epistemic realism will take into account whether any given epistemic realist view accepts or rejects each of the following three theses:

IRREDUCIBILITY: At least some epistemic-normative facts are *irreducible* to non-normative facts.

OBJECTIVITY: At least some normative facts hold independently of the set of an agent's contingent pro-attitudes, such that there can be a normative reason for an agent to X even though the conclusion that there is this reason in no way follows, as a logical or instrumental matter, from the agent's own contingent set of pro-attitudes in combination with non-normative facts.⁶¹

⁶¹ My understanding of what the *objectivity* of reasons amounts to here, and of what its denial (subjectivity) entails yields a distinction between *objective* and *subjective* reasons which corresponds to the way in which the distinction between *categorical* and *hypothetical* reasons is often drawn in the philosophical literature. (See for instance Cuneo:

CAUSALITY: At least some epistemic normative facts are *causally efficacious*.

Rejection or acceptance of each of the theses above can be variously combined in epistemic realist views. To be sure, some of these combinations will be more obvious than others. But this fact ought not to obscure the complexity of the ways in which these theses can come together. Indeed, as will become clear in the course of this thesis, a correct evaluation of the impact of Darwinian considerations for meta-epistemology lies precisely in bringing to the fore the full complexity of the range of epistemic realist views which can be occupied here. In order to appreciate this, we must say more about what each of the theses above amounts to and about how they can be combined.

When, in the last chapter of this thesis, I will turn to discuss Street's evolutionary arguments against epistemic realism, I will defend the claim that the best way to read those arguments is by seeing them as directed at epistemic-realist views which are committed to the IRREDUCIBILITY thesis. Because of this, it is of great importance to be clear about what the IRREDUCIBILITY thesis amounts to, and how IRREDUCIBILITY is related to the theses I have named OBJECTIVITY and CAUSALITY. I will turn to an examination of these topics in the next sections. While my main focus will be on the elucidation of IRREDUCIBILITY my discussion will also clarify what OBJECTIVITY and CAUSALITY might amount to, when these are understood as theses about the nature of epistemic normative facts.

4.1. Irreducibility

IRREDUCIBILITY is a claim about the nature of normative facts. It says that facts of this kind are a *sui-generis* kind of facts, in that they are not reducible to facts of other kinds. In what follows, I will articulate a general framework in which what the irreducibility of normative facts

2007: 38-9]. This distinction between *objective* and *subjective* reasons also corresponds to the distinction which is often drawn between 'internal' and 'external' reasons, on the understanding of these expressions that originated with Bernard Williams's seminal paper (1981).

amounts to can be made perspicuous. I will then use this framework to elucidate the nature and variety of non-reductive realist positions about epistemic normativity. In doing so, I will build on the conceptual distinctions I have introduced in Chapter 2 and make use of the meta-epistemological map which I have sketched in section 2 of this chapter.. I will take normative facts to consist in the instantiation of normative *properties*. As I have done so far, my focus will be on *reasons*. Thus the (paradigmatic) normative property the discussion to follow is intended to be about is the relational property of *being a reason*.

To say that normative facts are irreducible is to say that a *reductive account* of facts of this kind cannot in principle be provided. But what is it to provide *an account* of the intended kind, and what would it be for an account of this kind to be a *reductive* one? By an account of the kind at issue here, I will mean *a constitutive account*. Following Wedgwood (2007), I will take a constitutive account of some property of kind P^{62} to be provided by a necessary bi-conditional which is such that it specifies the *real definition* or states the constitutive *essence*, i.e. the *nature* of the property in question. And I will call a constitutive account *reductive* if and only if the right hand side of the bi-conditional (the one in which the nature of the thing that is being defined is specified) is built up out of terms for properties and relations the nature of which can be explained without making any reference to any property of the kind which is being explained. I will take normative properties to be properties which are non-accidentally referred to by normative expressions or concepts.⁶³ And I will take non-normative properties to be those properties which are such that (1) they are non-accidentally referred to by non-normative expressions or concepts (2) a

⁶² Henceforth, I will be speaking of properties, rather than facts, for the sake of convenience. Much of what I will have to say here will also apply – in obvious ways – to normative facts. If one were to hold the view that there are no normative properties, but *only* normative facts, then what I have to say can also easily be reformulated in a way which is compatible with this view.

⁶³ This way of individuating normative properties and facts does not take into account the possibility of views of normativity which are both realist and expressivist. According to these views, normative expressions and concepts, strictly speaking, do not refer to anything. Thus, the identification of normative facts and properties on the basis of what normative expressions refer to cannot work here, and it will have to be adjusted to accommodate those views. One way one could do so is by saying that a property is normative iff it is the property which causes the non-cognitive state which is typically expressed by normative expressions and concepts. But the details of an account along these lines will differ according to the specific nature of the realist expressivist views in question. We can safely ignore these complications here and in what follows in the chapter.

constitutive account of these properties can be provided in terms of expressions or concepts that do not refer to any normative property and also do not refer to any other entity a full account of which would in turn have to refer to some normative property.⁶⁴

On the basis of the terminology I have introduced so far we can say the following. To claim that it is possible to provide a reductive account of normative properties is to say that it is possible, at least in principle, to provide an account that specifies what it is for something to be that property (or to have that property) purely in terms that stand for non-normative properties, together with topic neutral terms such as the logical vocabulary. The IRREDUCIBILITY thesis, then, will amount to the claim that an account of this kind cannot in principle be provided.

As we have seen, reductive accounts share three important features. Firstly, they take the form of necessary bi-conditionals (NECESSITY). Secondly, they are not circular (NON-CIRCULARITY). Thirdly, they provide a specification of the nature (the constitutive essence) of the property which figures on their left-hand side (ESSENCE). NECESSITY, NON-CIRCULARITY, and ESSENCE are necessary conditions which a bi-conditional must satisfy in order for it to count as a successful reduction. Thus for a *candidate* reductive account to fail to be a *successful* reductive account, it is sufficient that it fail to satisfy at least one of the conditions above. This suggests that there are three distinct ways in which any given allegedly reductive account can in principle be challenged. It is useful to draw attention to these because – in line with the goal of the present enquiry – they allow us to see three distinct ways in which a non-reductive realist about epistemic reasons can defend irreducibility in the face of proposed reductive accounts of normative properties. Noticing this fact also brings to the fore that the realist's acceptance of IRREDUCIBILITY might take quite different forms and be motivated by very different kinds of considerations. These will come down to different ways of understanding

⁶⁴ Condition (2) needs to be added because, arguably, normative properties can be non-accidentally referred to by non-normative expressions, as in the following case: 'For all X, X has the property of being wrong iff x has the property that I am now actually thinking about' (on the assumption that 'the property that I am now, etc.' is a rigid designator and, of course, that the property I am now actually thinking about is the property of being wrong).

the relation between normative facts on the one hand, and alethic facts (i.e. ‘evidence-facts’) and – in some cases – facts about the pro-attitudes of agents, on the other.

4.1.1. Challenging NECESSITY

One way to resist reduction goes via the denial of the necessity of the bi-conditional that is alleged to spell out the reduction. This can happen in either of two ways (inclusive ‘or’). Consider the following candidate reductive claim:

NEC (x is R iff x is N)⁶⁵

Where R stands for some normative property and N for a non-normative property; N can be either simple or complex – according to whether it is or it is not built out of non-normative properties together with logical vocabulary, etc. A claim of this kind can be resisted by showing that x can be R without being N. Or by claiming that x can be N without it being R.

By appealing to this strategy, for instance, a non-reductive epistemic realist might resist the claim that facts about epistemic reasons *reduce* to facts about the pro-attitudes of agents together with facts about what means are likely to satisfy them. That is to say, they can resist the idea that what it is for A to have an evidence-based reason to believe that P, is just for A to have a desire to believe that P if P is true and for there to be evidence (accessible to A) for the truth of P. Believing what there is evidence for is a means to acquiring true beliefs. Thus the epistemic reason for A to believe that P is ultimately an instrumental one, and one which depends on the nature of the pro-attitudes a given agent happens to have. One way to resist this Humean-instrumentalist, realist and reductivist account of epistemic reasons goes via denying that means/end facts of the kind I have indicated, together with facts about the pro-attitudes of

⁶⁵ R stand for some normative property (i.e. *being a reason*); N for a non-normative property (either simple or complex – i.e. being built out of non-normative properties *together with* logical vocabulary, etc.).

agents, are either necessary or sufficient to secure the presence (i.e. to entail the existence) of epistemic reasons. An epistemic realist might do so by defending the claim that epistemic reasons do not depend on the pro-attitude of agents in the way the truth of the alleged reductive account would require.

This is not *the only way* to resist reduction of epistemic reasons by appealing to the strategy of claiming that NECESSITY has not been satisfied. For one might suggest that epistemic-reasons are reducible to facts about probability instead (see below). And in this case, a defence of IRREDUCIBILITY along the lines I am considering here will go through providing arguments to the extent that the (kind of) probability facts which are indicated as providing the alleged reduction are neither necessary nor sufficient for the presence of epistemic reasons themselves.

4.1.2. Challenging NON-CIRCULARITY

The second strategy for resisting reduction consists in claiming that, despite appearances to the contrary, the alleged non-normative properties that are appealed to in the specification of normative properties are *in fact* themselves normative. What this amounts to saying is that the correct constitutive account of the alleged non-normative properties that ‘reduce’ the normative ones cannot be given without itself ultimately having to mention normative properties. The specific nature which a defence of IRREDUCIBILITY along these lines will take will depend on the specific nature of the suggested reduction.

Take the suggestion that the property of *being an epistemic reason* ‘reduces to’ the property of *evidentially supporting*. On this view, what is it is for a fact F to be an *epistemic reason* for believing that P is for this fact to *evidentially support* the truth of P. An epistemic realist might defend IRREDUCIBILITY in this case by defending the thesis that it is the property of *being evidence for* which is to be explained and accounted for in terms of the property of *being an epistemic reason for*,

rather than vice-versa. Were this to be the case, the suggested reductive account would violate the non-circularity condition that genuine reductions must fulfil.

Or take again the suggestion that facts about epistemic reasons reduce to facts about the pro-attitudes of together with facts about what is instrumental to fulfilling them, along the lines of the Humean proposal I mentioned above. One way to resist reduction by appeal to the violation of NON-CIRCULARITY here could go via a defence of the claim that ‘the mental is normative’: that a specification of the essence of any kind of mental state will ultimately require appeal to normative properties of some kind or other (Wedgwood). If this were true, the suggested reductive account would fail. At best, if at all, it will succeed in reducing epistemic reasons (in the sense I have been speaking about them) to reasons or normative facts of some other kind.

4.1.3. Challenging ESSENCE

A third strategy that it is possible to deploy in challenging a reductive account of normative properties is as follows.⁶⁶ The anti-reductionist will grant that the bi-conditional that is alleged to provide the reductive account is indeed necessary. He will also grant that the properties referred to in the right-hand side are genuinely (i.e. not only *prima facie*) non-normative. Nonetheless, he will argue, the bi-conditional at issue does not yet amount to the announced reduction. This is so because the mere necessary co-extensiveness of normative and non-normative properties cannot be taken to provide – in the particular case at hand – *a specification of the essence* of the normative properties themselves.

One way to illustrate what a defence of IRREDUCIBILITY along these lines might amount to can be provided by reflection on a form of argument which is often deployed in defence of non-naturalist realist positions in metaethics and can be found in the writings of Dancy (2006),

⁶⁶ See Shafer-Landau (2003: 98) for acknowledgement of a position of the kind I am about to discuss on the relation between normative and non-normative properties and facts.

Parfit (2011), and McNaughton and Rawling (2003)⁶⁷ among others. Copp (2011:44) summarises the version of this of argument in McNaughton and Rawling (2003) as follows:

A reductive naturalist must propose a theory that identifies each normative property R with some particular natural property N. But [...] if such a proposal is to be at all plausible, it must at least be the case that when something is R, it is R because it is N. But if so, then a thing's being R is not identical to its being N. The 'because relation' is asymmetrical; it cannot hold between a fact and itself [...]. If an action is wrong, it is wrong because of some of its natural properties, and this means that wrongness is distinct from these natural properties.⁶⁸

Arguments such as the one in question can help us see a way in which suggested reductive accounts of normative facts can be resisted, even in the face of a successful specification of a category of non-normative facts which are necessarily coextensive with them. These arguments appeal crucially to a distinction⁶⁹ between *normative facts* and what, with Dancy and Parfit, we might call *normatively significant* facts. In the specific case of normative reasons – for instance that F is G is a reason for Q – the distinction at hand would correspond to that between facts which provide the grounds for a specific reason (F is G) and the facts which consist in the instantiation of the reason relation between a given ground (F is G) and a specified attitude (Q) – such as the fact *that F is G is a reason for Q*. With the 'ground' of the reason corresponding to the *normatively relevant* fact, while the (whole) fact *that F is G is a reason for Q*, corresponding to the *normative fact proper*, so to speak. Proponents of the kind of argument I am considering here hold that expressions such as 'in virtue of' or 'because of' or 'makes' which are at work in claims of the form:

X is R *because of* X being N [F is a reason for Q because F is G]

⁶⁷ Arguably, the (different) argument in Enoch (2007) in favour of the view he calls 'robust metanormative realism' can also be taken as providing a defence of IRREDUCIBILITY of the kind I explore here. For – it seems to me – that argument is meant to be compatible with the assumption that a candidate reductive account of normative facts satisfies the NECESSITY and NON-CIRCULARITY conditions. I do not have the space to address this matter here.

⁶⁸ I have changed the letter M (for 'moral') to the letter R (for 'normative'/'reason') to fit what the quote is about with my way of setting things up in this chapter. What Copp calls 'a reductive naturalist' corresponds – in the vocabulary I have been using – to someone who believes that normative facts and properties can be reduced to non-normative facts and properties.

⁶⁹ Another way to read arguments of this kind would take them to *establish* such distinction.

or:

X is R in *virtue of* its being N [F is a reason for Q in virtue of its being G]

or:

X being N is that which *makes* X R [F being G is what makes F a reason for Q]

refer to the *normative relation* which holds between a given (kind of) normatively significant fact and (the kind of) attitude (Q) for which this (kind of) fact is normatively significant. Because of this, we might say these expressions are here employed in a ‘substantive normative sense’ (Copp; 2011: 45). Normatively significant facts and what I have called ‘normative facts proper’ are *distinct* kinds of facts on this analysis; the former are *a component* of the latter. Thus one can hardly think of providing a reduction of the latter in terms of the former. This would hold even on the assumption that claims of the form:

X is R iff X is N

are true, and that they satisfy the NECESSITY and the NON-CIRCULARITY conditions. Going back to the quotation above, it is not at all clear that merely drawing attention to the kind of ‘asymmetry’ of the ‘because/in virtue of relation’ which would hold between facts of kind R and facts of kind N is sufficient to establish the claim that R-facts and N-facts must be distinct, and thus to resist reduction of the former to the latter. After all, expressions such as ‘in virtue of’ or ‘because of’ are very often employed in what we might call an ‘essentialist sense’. That is to say, successful reductions are often *stated* by means of vocabulary of this very kind, as in the following cases: ‘X is water *in virtue of* X being H₂O’; ‘X is a circle *because* it is a line enclosing an area in a plane and there is a point within the area that is equidistant from every point on the line’. The kind of relation which ‘because/in virtue of’ express in those cases is also ‘asymmetrical’: X is water in virtue of X being H₂O but, arguably, X is *not* H₂O in virtue of X

being water, and so on. Thus merely drawing attention to the kind of ‘asymmetry’ considerations of the kind McNaughton and Rawling rely on is not likely, by itself, to persuade the reductionist who believes that R-facts reduce to N-facts of the misconceived nature of their reductionist attempt. Neither is the reductionist in question likely to be persuaded, for that matter, by someone merely *stating* that ‘because of’ in claims such as ‘X is R because of its being N’ is to be read in the ‘substantive normative sense’ (Parfit; 2011: 299). However, the converse is also true. The anti-reductionist will resist the claim that the only – or the best – way to understand claims such as ‘X is R in virtue of X being N’ is by going the reductionist route. And he will not be swayed by the fact that the property of *being N* and *being R* turn out to be necessarily coextensive.

Whatever one’s position on the force and prospects of arguments of the kind I have reviewed – Copp (2011) calls them ‘arguments from because’ – discussion of them plays a valuable role in the context of my analysis of the IRREDUCIBILITY thesis. It shows that, even granting the NECESSITY and NON-CIRCULARITY of certain candidate reductive claims, there is substantive disagreement to be had with regard to whether those alleged reductions are successful. As I have indicated, these will come down to substantive disagreement between reductionists and anti-reductionists with regard to how to understand expressions – such as ‘in virtue of’ and ‘because of’ – which figure crucially in those claims. Disagreement of this kind will correspond to disagreement about what kind of relation must be thought to hold between normative facts and the non-normative facts which they are necessarily coextensive with. Pending resolution of these issues, we can say that the fact that the NECESSITY and NON-CIRCULARITY conditions are satisfied by a given candidate reductive account need not, by itself, warrant belief in the fact that the reduction in question is successful.

The strategy for resisting reduction (by denying ESSENCE) which I have just exemplified is quite general. Thus it is unsurprising to find that it can be put to work against reductive accounts of normativity in epistemology. In order to see how this strategy can be put to work here, let us

go back to the suggested reductive accounts of epistemic-reason facts which I have already considered.

The first case attempted to reduce facts about epistemic reasons to facts about pro-attitudes of agents of a specific kind together⁷⁰ with specific kind of facts regarding what is instrumental to satisfy those attitudes.⁷¹ On the scenario we are considering here the non-reductive realist will agree on the claim that epistemic reason facts and facts of the kind which are mentioned in the alleged ‘instrumentalist reduction’ are necessarily coextensive. He will also grant that the facts mentioned in the alleged reduction are genuinely non-normative. However, he will still go on to say that the reduction at hand will still not succeed. He will say that it is, in general, quite plausible to think that if any normative fact obtains, this is so because some non-normative fact or other obtains. There is a reason to run out of the building *because the building is on fire*, a reason to cry *because one’s close relative just died*, and so on. We should expect epistemic reasons to be no different. On the account we are considering, every time an epistemic-reason fact obtains, the same kind of (complex) non-normative fact will provide the reason in question. There is a reason for A to believe that P *because* A desires to believe that P if P is true and *because* there is evidence for P, and *because* believing what one has evidence for is a means to forming true beliefs. However, the non-reductive epistemic realist will claim that the facts which provide the epistemic reasons in each case and the fact that they provide epistemic reasons are *distinct* facts. If we think back to the relational structure of epistemic-reasons facts (Chapter 1), one way to put this line of argument is by saying that the suggested reduction mistakenly conflates the whole relational fact (That G is a reason for P) with one of its *relata*, i.e. with the ground G of the relation itself. If the aim was to provide a reduction of the whole (complex, relational) fact then this strategy can hardly succeed, even on the assumption that NECESSITY and NON-CIRCULARITY are satisfied.

⁷⁰ Such A’s desire to believe that P if P is true.

⁷¹ I.e. facts such *that believing that P upon coming across evidence E for P is a means/instrumental to the acquisition of true beliefs*.

The second case attempted to reduce facts about epistemic reasons to facts about evidence, where these latter are not understood as being themselves irreducibly normative. Here again, a non-reductive realist can agree on the claim that epistemic reason facts and facts about evidence are necessarily coextensive. He can nonetheless insist that the fact that, say, *E is evidence for H* and the fact *that E is evidence for H is a reason to believe that H* are distinct facts. They are facts about different things, with a different ‘subject matter’ (Dancy: 2006) and with different ‘constituents’. Because of this, one can hardly expect that facts of the former (non-normative) kind will be in a position to provide a successful reduction of facts of the (latter) normative kind.

4.2. Irreducibility and Objectivity

At this point of our discussion, it is worth drawing attention to some important aspects of the relation between IRREDUCIBILITY and OBJECTIVITY. The relevant thing to stress here is that, while commitment to IRREDUCIBILITY can and often will typically be accompanied by a commitment to OBJECTIVITY in epistemic realist views, commitment to OBJECTIVITY is neither entailed by nor does it need to be an essential feature of non-reductive realist views. The reverse is also, if perhaps a lot less surprisingly, the case: commitment to IRREDUCIBILITY is neither entailed by nor it is an essential feature of epistemic realist views which are committed to OBJECTIVITY.

IRREDUCIBILITY is the claim that normative facts are *neither identical with nor reducible to* non-normative facts of some specifiable kind. OBJECTIVITY is the claim that normative facts are *independent from* contingent facts pertaining to the pro-attitudes of agents. To the extent that claims about the *dependence* (or independence) of facts of kind S upon facts of kind B are distinct from claims about the *reducibility* of facts S to facts B, it is not surprising that OBJECTIVITY and IRREDUCIBILITY will be distinct theses. In other words and generally speaking, whether normative facts are pro-attitude-independent is *orthogonal* to whether they are irreducible.

How IRREDUCIBILITY and OBJECTIVITY are different theses can perhaps be more perspicuously elucidated by drawing attention to how their truth-values can come apart in concrete cases. The following scenarios illustrate the non-equivalence of the two theses.

Scenario One: it is possible to hold that there is a reason for A to Q only if the conclusion that there is this reason follows as a logical or instrumental matter from A's own contingent set of pro-attitudes in combination with the non-normative facts, and yet still believe that the fact that there is such a reason for A is an irreducibly normative fact. This gives IRREDUCIBILITY without OBJECTIVITY. In the epistemic case, for instance, it is possible to hold that the fact that F is evidence for P is a reason for A to believe that P, only if A has the pro-attitude of desiring to form a true belief as to whether P and to claim – further – that the reason in question is instrumental in kind; that is to say, that it holds in virtue of the fact that believing what one has evidence for is instrumental to the acquisition of true beliefs. On this analysis, that there is a reason for A to believe that P on the basis of F *follows instrumentally* from A's desire to form a true belief as to whether P; thus A's reason to believe that P is not *objective* in the sense at work in the OBJECTIVITY thesis. It is nonetheless possible to insist that the pro-attitude-dependent, normative fact in question is neither identical with nor reducible to the fact that F is evidence for P together with the fact that A has a pro-attitude of the specified kind (Scanlon: 2009a)⁷². The defender of irreducibility can, in this case, invoke the distinction between reason-giving facts (evidence fact + pro-attitude fact) and the fact that these facts are reason-giving. And he can claim that this latter fact supervenes on and is necessitated by, but is neither identical nor reducible to, facts of the former kind (Parfit: 2011, Shafer-Landau: 2003). Similarly, for instance, a realist committed to IRREDUCIBILITY can agree on the fact that instrumental reasons depend on there being a means/ends relation of the appropriate kind between the given

⁷² To these one may want to add, perhaps, the 'means/ends fact' that believing what one has evidence for is instrumental to the acquisition of true beliefs.

contingent goal of an agent and a certain course of action. Nonetheless, he can still resist the claim that the instrumental reasons in question are *reducible* to that relation.

Scenario Two: It is possible to claim that normative facts are reducible to pro-attitude-independent facts of certain specifiable kinds. This gives OBJECTIVITY without IRREDUCIBILITY.

Again looking at the epistemic case, it is possible to think that the fact that F is an evidence-based reason for believing that P *reduces to/is identical with* the fact that, for instance, the probability of P given F is significantly above 0.5 and that F is the case. That is to say, epistemic-reason facts reduce to probability-facts of certain specifiable kinds. In this case, epistemic-normative facts would be reducible in my sense, and yet clearly not pro-attitude-dependent.⁷³

4.3. Irreducibility and Causality

Like IRREDUCIBILITY and OBJECTIVITY, IRREDUCIBILITY and CAUSALITY can also come apart, and not always in obvious ways. The importance of distinguishing these theses and investigating their relation can hardly be overestimated.

One popular way of characterising ‘naturalistic’ approaches to normativity which can be found in the philosophical literature defines ‘naturalism’ in terms of *reductionist* claims of various kinds, for instance in terms of the thesis that normative facts are *reducible* to non-normative facts, or that normative discourse and practices can be explained in terms of *non-normative facts* (Gibbard: 1990 and 2003, Blackburn: 1993) or that – to put it in David Papineau’s words ‘there are no *sui-generis* normative facts’ (Papineau: 2006). Another way in which the label ‘naturalism’ is frequently used in the context of discussions of normativity is as a means to distinguishing realist views which are committed to the *causal efficacy* of normative properties from those realist views which are not. (Parfit: 2011, Wedgwood: 2007). Being clear about the relation between

⁷³ I am assuming that the probability-facts at work here are not ones which are themselves ultimately to be understood in terms of normative facts. I am also assuming that, if one were to adopt a subjective Bayesian account of probability, the ‘degrees of belief’ which figure centrally in this account will not count as pro-attitudes of any kind.

IRREDUCIBILITY and CAUSALITY will shed valuable light on how those different ways of understanding naturalism about normativity are related. It will also provide some tools for disambiguating the often generic avowals of ‘naturalistic’ views about normativity in epistemology which can be found in the philosophical literature.

Lastly, a discussion of how IRREDUCIBILITY and CAUSALITY are related is all the more urgent given the role played by causal/explanatory considerations in much debated arguments against various forms of normative realism in metaethics (Harman: 1977 and 1986, Sturgeon: 1986 and 1988, Enoch: 2007). Discussing these theses will be a crucial preliminary step to investigating the counter-parts of those arguments and debates in the epistemic domain.

A thorough examination of the relation between the irreducibility and the causal efficacy of epistemic normative properties and facts raises difficult issues which I cannot here address in full and with the attention they would deserve.⁷⁴ The following remarks, however, should suffice to alert the reader to some of the complex ways in which the two theses can be related, when it comes to the domain of epistemic normativity.

We are familiar from traditional taxonomies in metaethics with the idea that claims about the irreducibility of (moral) normative properties to non-normative properties are compatible with either accepting or rejecting the claim that irreducible normative properties are causally efficacious. Realist positions in metaethics which are often referred to as non-reductive naturalist (Boyd: 1988, Sturgeon: 2006, Oddie: 2005) can be viewed as being committed to moral versions of both the IRREDUCIBILITY and CAUSALITY theses. On the other hand, views which are often labelled as non-naturalist realist (Parfit: 2011, Enoch: 2007, Dancy: 2006) can be

⁷⁴ I have shown how there are different ways to understand and defend IRREDUCIBILITY. It is quite reasonable to expect that how one understands IRREDUCIBILITY will affect the nature of one’s position on the CAUSALITY thesis. A defence of the causal efficacy/inefficacy of normative facts on an account of IRREDUCIBILITY that grants that normative facts are necessarily co-extensive with non-normative facts of specifiable kinds might look very different from a defence of the causal efficacy/inefficacy of normative facts in an account of IRREDUCIBILITY that denies that normative facts and non-normative facts are so related. As I said, I cannot discuss this matter further here.

understood as also holding on to IRREDUCIBILITY, while at the same time rejecting the CAUSALITY thesis.

In the spectrum of normative realist positions which can be occupied in epistemology, we can expect IRREDUCIBILITY and CAUSALITY to come apart – or not to come apart – in similar ways, and to do so for reasons which will be similar to those which are adduced to defend non-naturalist realist or non-reductive naturalist positions in metaethics, respectively. This is indeed the case. Wedgwood's discussion of normativity in *The Nature of Normativity* (2007) is meant to cover both the practical and the epistemic sphere. The general framework of his discussion reflects the central belief that, as he put it, 'the normativity of the intentional is the key to metaethics' (2007: vii). Because of this, and because he does not address the case of epistemic normativity separately it is often difficult to extract his views on some of the topics I have discussed.⁷⁵ However, it is quite clear that the position he outlines and defends in his book is a kind of normative realism which subscribes both to IRREDUCIBILITY and to CAUSALITY. Wedgwood writes:

It would seem [...] that there is no prospect of reducing mental properties and normative properties to any more basic family of properties. Mental and normative properties are metaphysically irreducible. (2007: 184)

And he goes on to add:

The argument I gave for the normativity of the intentional has a further implication about the metaphysical character of normative facts and normative properties. Specifically, it implies that normative facts and normative properties are *causally efficacious*. (2007: 185, italics mine)

Parfit's treatment of epistemic reasons in *On What Matters* (2011), on the other hand, is rather explicit in portraying epistemic-reasons facts as being strongly supervenient on – while being distinct from – what he calls 'alethic facts' (i.e. evidence-facts). He is likewise explicit in claiming that normative facts of the kind at issue are 'not part of the causal fabric of the universe'. That is

⁷⁵ It is hard, for instance, to extract Wedgwood's views on the relation between the notion of evidence and that of an epistemic reason.

to say, Parfit's realism about epistemic reasons is clearly committed both to the IRREDUCIBILITY thesis and to the denial of the CAUSALITY thesis.⁷⁶

If believing in IRREDUCIBILITY is thus compatible with either accepting or rejecting CAUSALITY, so is believing that IRREDUCIBILITY is false. This is perhaps less obvious and less likely to be noticed. But it would be a mistake to think that once epistemic-normative facts are reduced to non-normative facts, the issue of whether or not they are *causally efficacious* is then settled 'by default', so to speak. That is to say, it would be wrong to assume that the reduction of epistemic-normative facts to non-normative facts entails that the epistemic normative facts thus reduced are 'part of the causal fabric of nature'. To illustrate this point, consider the following. One possible reductionist view regarding facts about epistemic reasons says that these facts are to be understood in terms of (i.e. as being reducible to) facts about evidence, where facts of this latter kind are not themselves thought to be irreducibly normative. Assume, as many philosophers are inclined to hold,⁷⁷ that facts about evidence are, in turn, to be analysed in terms of facts about probability.⁷⁸ For instance, that F is evidence for P can be taken to come down to the fact that F being the case raises the probability of P being true, or, alternatively, to the fact that F being the case, the probability of P being true is/raises above a certain threshold N.⁷⁹

If facts about epistemic reasons are ultimately to be understood in this way, then the issue whether they are causally efficacious comes down to the issue whether *probability-facts* of the kind I have exemplified are causally efficacious. And this latter question is one which is far from being straightforwardly settled. An answer to it will, arguably, depend on what one takes 'probability-facts' to be. Intuitively, on some interpretations of probability – such as subjectivist, frequentist, and propensity interpretations – it will be easier to provide some account of their causal efficacy than on others – such as so called 'logical' interpretations.

⁷⁶ For a similar version of non-reductive realism about normativity in general see also Enoch (2007).

⁷⁷ See, for instance, the discussion in Achinstein (1978) and Cloos (unpublished manuscript).

⁷⁸ If, for instance as Achinstein maintains, one were to find it more plausible to believe that evidence is to be understood in terms of both probability and explanation instead, what I will have to say below still applies.

⁷⁹ A view of this kind is discussed (and criticised) in, for instance, Achinstein (1978).

For instance, if probabilities are analysed in terms of *degrees of belief* (as in subjective Bayesianism) and if beliefs (typically in combination with desires) are thought of as causally efficacious (they cause agents to behave in certain ways, they causally explain why agents behave as they do, etc.), then it is easy enough to see how an account of the causal efficacy of facts about probabilities might go. It is also easy to see how we can make sense of the claim that agents can be ‘causally affected’ by evidence, on this interpretation of probability. For to say that evidence F supports proposition P will be tantamount to saying that the discovery of F causes an agent’s *degree of belief* in P to raise (or to raise above a certain threshold N). (See Lipton; 2004: 103).

On the other hand holding on to CAUSALITY will be more difficult on (and perhaps a much less likely commitment of) other accounts of probability, such as those for example which hold – as logical theories of probability do – that probabilities are logical, objective, and necessary relations between propositions, and that they can be known *a priori* (See Plantinga; 1993: 144). If what I have said so far is correct, it shows that the *reducibility* of normative facts, by itself, entails nothing about their *causally efficacy*. Thus the denial of IRREDUCIBILITY is, in principle, quite compatible with either accepting or rejecting the truth of the CAUSALITY thesis.

5. Varieties of Epistemic Realism II: Instrumentalism, Evidentialism, and Evidence-as-Reasons Views

In the previous chapter, I drew attention to the fact that claims about epistemic reasons are often ambiguous or elliptical. I then discussed how the ambiguity at issue can hide substantive disagreement regarding the notion of an epistemic reason. I also, if briefly, drew attention to how this can lead one to confuse meta-epistemological views (and meta-epistemological debates) that ought to be kept distinct. It is not surprising then that a second element on the basis of which epistemic realist positions can be classified is determined by the interpretation of epistemic reasons they are committed to.

We can expect that on each of the instrumentalist, evidentialist, and evidence-as-reasons interpretations of *epistemic-reason claims*, epistemic realists can hold very different views regarding the nature of *epistemic-normative facts*. From my discussion of epistemic realism so far it will have already emerged that the question ‘*what kind of facts are epistemic-normative facts?*’ can elicit disagreement among epistemic realists not only with regard to their positions on the IRREDUCIBILITY, OBJECTIVITY, and CAUSALITY theses. For epistemic realists can and will also disagree on *what they are being realists about*, so to speak. That is to say they will disagree on the kind of ‘structure’ and on the kind of ‘constituents’ which they will take normative facts to have.

5.1. Realism and Instrumentalism

Take a paradigmatic epistemic-reason claim of the form:

ER: F* is an epistemic reason for A to believe that P (in circumstances C).⁸⁰

Remember also that on the understanding of epistemic reasons I am operating with (1) ER conceptually entails that there is evidence for A to believe that P and that (2) among the facts which comprise F* is some fact or facts E which is such as to be or provide evidence for the truth of P (See Chapter 2, Section 3.1). And let us begin by considering realism about epistemic reasons construed along the lines of epistemic instrumentalism. Standard versions of realism about epistemic reasons thus understood will believe that claims about epistemic reasons, when true, are made true by (normative) facts. But regarding the nature of the facts that make these normative claims true, epistemic realists of this kind can differ widely. They will all grant that facts about epistemic reasons will be complex facts with facts of the form *X is instrumental to Y* among their constituents (i.e. believing what there is evidence for is instrumental to being in

⁸⁰ Example: That the butler’s fingerprints are on the murder weapon is a reason to believe that the butler committed the crime.

cognitive mental state C^{81} /believing that P on the basis of E, when E is evidence for the truth of P is instrumental to being in C). But they will disagree, for instance, over whether epistemic-reason facts must *also* have facts about the contingent pro-attitudes of the agents (whom the reasons are for) among their constituents; or whether, instead, they will have to include the fact that those agents have a pro-attitude-independent reason of some sort or other for being in cognitive mental state C.

We can expect epistemic realists who are committed to a Humean Instrumentalist analysis of epistemic reasons to take a different position on those issues than Non-Humean epistemic realists. In other words, epistemic realists of these respective kinds will disagree over whether epistemic-normative facts of the kind we are considering are *objective* in the way the truth of the OBJECTIVITY thesis requires. Analogously, it is possible for epistemic realists committed to an instrumentalist reading of epistemic reasons to disagree over whether the normative facts in question are *irreducible*. I have already illustrated some of the different ways in which this can happen.

In general, whatever non-normative facts N one suggests for a reduction of normative facts R, there will always be scope for defending the thesis that facts of kind R supervene upon and are necessitated by facts of kind N, but are nonetheless distinct from these.⁸² That is to say, it will always be possible to defend IRREDUCIBILITY by appealing to the distinction between *reason-providing facts* and the fact *that those reasons-providing facts are reason-providing*. To the extent to which one finds this strategy doubtful or question-begging (Copp; 2011: 46), there will also always be room for re-iterating reducibility of R-facts to N-facts.

Lastly, I think we must allow for the possibility that epistemic realists who are committed to instrumentalism can disagree over whether normative facts are *causally efficacious*. I have already

⁸¹ C can be variously understood. It can be ‘believing that P if P’, ‘having a justified belief as to whether P’, ‘knowing that P’, etc.

⁸² This line of thinking seems to be what underlies the confidence of many non-natural realists about normativity regarding the fact that no non-normative properties and facts can possibly be produced which will be such as to provide a successful reduction of normative facts themselves (Dancy: 2006, McNaughton and Rawling: 2003).

commented on the lack of clarity that often besets the terms of the disagreement regarding the issue of the causal efficacy of normative facts. Without delving too deep into this topic, we can illustrate the point I want to make here as follows.

Suppose that, upon coming across (and attending to) strong evidence E for the belief that P, agent A forms the belief that he has strong evidence for P. Add to this that A happens to have a strong desire to form a true belief as to whether P, and that he is aware of this desire. In these circumstances and upon reflecting on the situation he is in, A forms the belief that he has a strong reason to believe that P. Let us assume, lastly, that a Humean-Instrumentalist analysis of epistemic reasons is correct⁸³ and that realism about epistemic reasons is the correct meta-epistemological view. If one were to ask ‘why has A formed the belief that he has a reason to believe that P?’ it would seem that a correct way to answer it would be by saying: ‘because he had a reason to believe that P’. By looking at the form of this explanation, it is not *prima facie* implausible to think that it is a *causal* one⁸⁴ – to think, that is, that the fact that *caused* A’s normative belief was a normative fact. On a reductionist account that identifies facts about reasons with means/ends facts and facts about pro-attitudes of agents, it is not too difficult to see how this intuition about causality can be preserved.⁸⁵

However, that this intuitive and at least initially plausible view is the correct one can be a matter of some controversy. Take a Humean Instrumentalist who also subscribes to IRREDUCIBILITY by appealing to a distinction of the kind I have discussed above. He will grant that the facts that E is evidence for P, that A desires to form a true belief as to whether P, and that forming beliefs on the basis of evidence is instrumental to acquiring true beliefs are *sufficient* to secure the presence of an epistemic reason to believe that P. Given that he agrees on

⁸³ The version of Humean Instrumentalist which I am assuming correct here, is the one that says that an agent’s epistemic reasons for forming beliefs of a certain kind depends on his having a desire to form beliefs of those kind in the case they are true.

⁸⁴ There is an important difference between the two (causal) *explananda* here.

⁸⁵ Mental states are, according to many, causally efficacious. And how means/ends facts of the kind needed for this account can be part of a causal story is not too hard to imagine. The complete and detailed story to tell here might not be a simple one, but it does not look like it will be one that it is too hard, in principle, to provide.

the Humean instrumentalist analysis of epistemic reasons, he will also grant that those facts are necessarily co-extensive with facts about epistemic reasons themselves. However, he will hold that while those facts *provide the grounds* for the reason to believe that P, they are not identical with that reason. For, as we saw, the facts which provide reasons and the fact *that those facts provide reasons* are distinct facts on this analysis.

Normative realists who appeal to a move of this general kind to resist reductionist views of reasons are also often committed to the denial of the *causal efficacy* of normative facts thus construed. Indeed, in the work of philosophers who hold ‘non-naturalist’ views of normativity, this anti-reductionist move (or something like it)⁸⁶ is also at the same time, taken to provide *an argument against the causal efficacy* of normative facts themselves.⁸⁷ If facts about epistemic reasons, on this analysis, are believed not to be causally efficacious, then A’s belief that he has a reason to believe that P cannot have been *caused* by the fact that A had a reason to believe that P.

Thus, that realists agree on a Humean instrumentalist analysis of epistemic-reason facts need not by itself settle their agreement on issue of the causal efficacy of normative facts. Whether those facts are ‘part of the causal network’ is a substantive issue which will have to be addressed separately.

5.2 Realism and Evidentialism

The variety of positions is similarly complex when it comes to epistemic realist views which are committed to the other interpretations of epistemic reasons which I have discussed, with the following exception. It seems to me that debates regarding whether normative facts are *objective* or *subjective* ought not to arise on either the evidentialist or the evidence-as reasons views, at least if ‘objective’ and ‘subjective’ are understood in terms of the acceptance or rejection, respectively, of my OBJECTIVITY thesis. Whatever one’s favoured interpretation of epistemic reasons is, it

⁸⁶ See Enoch (2007: 44).

⁸⁷ What the argument is, and whether it is a good argument, is an entirely separate matter. I will return to say something about this at the end of the next section.

is safe to assume that *the* notion of *evidence* they depend on (or that they are identified with, on what I have called the ‘evidence as reasons view’) is objective in this specified sense. That is to say, whether or not there is evidence F for a given proposition P and, thus whether any given agent A can come to possess evidence for P, does not depend on A’s contingent pro-attitudes – such as, for instance, A’s desire to believe that P if P. This way of understanding evidence is quite uncontroversial. It not only captures a central platitude that governs our application of the concept of evidence in ordinary and scientific contexts.⁸⁸ It is also typically taken for granted by many philosophers who reject OBJECTIVITY about epistemic reasons (Papineau, Street, Nozick, among others).

Notice that believing that evidence is *objective* in this sense leaves many other features of the concept of evidence entirely open. Rejecting OBJECTIVITY about epistemic reasons is entirely consistent with – and largely independent from – holding any of a variety of differing accounts of evidence. It is, for instance, compatible with believing in what we would classify as strong forms of *epistemic relativism*, such as, for instance, the view that whether F is evidence for P depends on whether a given community or social group believes this to be so. Insofar as one believes claims about evidence and claims about reasons to be equivalent in meaning (evidence-as-reasons views), or to be such that the latter entails the former (evidentialist views), this feature of evidence will have to be shared by epistemic reasons as well. Lastly, if one thinks – as ‘standard’ realists do – that epistemic reasons claims are made true by facts, these facts will also presumably have to inherit the *objectivity* of the reason claims they make true.

Realists on either of these accounts of epistemic reasons, however, can still disagree over whether epistemic-reason facts are such that IRREDUCIBILITY and CAUSALITY are true of them. I will discuss how this can happen on evidentialist interpretations in this section. I will turn to evidence-as-reasons view in the next. According to the evidentialist view of epistemic

⁸⁸ That evidence is objective in this sense is not among the many features of the concept of evidence which have been at the centre of much philosophical controversy (see Kelly: 2006 and 2008 for an overview of what those features are).

normativity, epistemic reason claims are distinct from, but entailed by, claims about evidence. More precisely, according to this view, it is a conceptual truth that facts about evidence – such as that F is evidence for the truth of P – provide the grounds for epistemic reasons – such as the reason to believe that P. As we have already seen, this view can be made perspicuous by using square brackets as follows:

EV: [The fact that E is evidence for P] is/provides an epistemic reason to believe that P.

Standard forms of realism about epistemic reasons thus construed will hold that claims of this form are made true by facts. Epistemic anti-realist views will deny that this is the case.

Despite the fact that, on this analysis, claims about evidence and claims about epistemic reasons are taken to be conceptually distinct claims, it is open to epistemic realists here to think that claims of each kind are nonetheless made true by the same (sort of) facts. If this is the case, they will believe that the claim that *the fact that F is evidence for P provides a reason to believe that P*, when true, is made true by *the fact that F is evidence for P*.

The position I am considering is, in this respect, similar to realist meta-ethical views which are often referred to as versions of ‘non-analytical naturalism’. On those views, claims about moral facts and claims about natural facts are conceptually distinct, and yet refer to facts of the same kind – i.e. ‘natural facts’ on some specified understanding of what facts of this kind amount to. On the view I am envisaging here, claims about epistemic reasons and claims about evidence are also conceptually distinct, and yet refer to facts of the same kind, i.e. facts about what evidentially supports what.

It seems to me plausible to assume that, if we were to take this view, we would also have to believe that the evidence-facts which true epistemic reasons claims will refer to will not be *irreducibly* normative facts. If this is the case, the view I am considering will be a *reductivist* view.

And if what I have said above about OBJECTIVITY is correct, the facts which make epistemic reason claims true will also be *objective*.

Of the issues I have been discussing so far, what thus remains open is whether those facts are also causally efficacious. I have already said (in the previous section) why I think that a successful reduction of normative facts to facts about evidence non-normatively construed does not by itself settle whether the CAUSALITY thesis is true or false of these facts. Thus I will take it that being an epistemic realist of the variety I have just discussed is, in principle, compatible both with believing that normative facts are causally efficacious and with believing that they are not. Settling this question will have to await a specification of what one takes facts about evidence to be.

But this is not the only way to be an epistemic realist on the evidentialist understanding of epistemic reasons. For an epistemic realist might instead take the *conceptual* difference between claims about evidence and claims about reasons to point to a *real* difference holding ‘in the world’, so to speak. On this view, the fact *that E is evidence for P* would be distinct from the fact *that E is evidence for P is a reason to believe that P*. *That F is evidence for P* would provide the ground for the reason to believe that P, thus entering the epistemic-reason relation as one of its *relata*. Epistemic normative facts would thus be distinct from facts about evidence, while nonetheless supervening on and being necessitated by them.

It is easy to see how this analysis of (the structure of) epistemic-reason facts can afford a way to defend IRREDUCIBILITY with regards to these facts. As we have seen, the most plausible candidate for providing a reduction of facts about reasons, under the interpretation of epistemic-reason claims I am discussing, are facts about evidence. But on the analysis of epistemic-reason facts I am considering, the attempts at reducing reason-facts to evidence-facts will be resisted by claiming that such reductionist attempts will rest upon a failure to appreciate that facts of the respective kinds have a ‘different subject matter’, so to speak (Dancy: 2006), and that whatever appeal the alleged reduction holds, it will be due to a misleading conflation between reason-

providing facts (i.e. evidence-facts) and the fact that these reason-providing facts are reason providing. To put it as Scanlon does, by identifying facts about reasons with non-normative facts, a reduction of the kind in question would rather seem to eliminate their *normativity* altogether (Scanlon: 2009a).

I have already noted how realists about normativity in general who appeal to a move of this kind in order to resist reductionist views of reasons across the board are also often committed to the denial of the *causal efficacy* of normative facts. However, merely noting that facts about evidence and facts about epistemic reasons are related in the way this analysis suggests ought not, by itself, to settle the issue of the causal efficacy of the latter one way or another.

Dancy seems to admit this much when he writes:

[One] could accept the suggestion that normative facts are facts with a distinctive subject matter but could now say that if this is all that the difference between the normative and the non-normative consists in, there is still no reason to suppose that normative facts cannot be natural. Why couldn't we think of facts about the practical [read: epistemic] relevance of other facts as natural – especially if they remain capable of contributing to the causal order? [...]. The only characterization of the natural that the normative does satisfy is the causal one – and that only possibly, the matter being very contentious. (Dancy; 2006: 135)

Thus I will take it that even if a defence of IRREDUCIBILITY along the lines I have explored here is successful, the position one holds on CAUSALITY cannot be taken for granted.. Much of what one thinks with regards to those issues will depend on their position on the causal efficacy of the evidence-facts which provide the supervenience basis for epistemic-reason facts on the model I am considering. I have already noted how settling whether evidence-facts are causally efficacious is far from being a straightforward matter, let alone one upon which one can assume there is some 'default position'.

Whether IRREDUCIBILITY entails CAUSALITY will also depend on the position one is inclined to take on the causal efficacy of supervenient properties more generally. Suppose

evidence-facts turned out to be causally efficacious. If one subscribed to the (unqualified) principle that properties and facts which supervene on properties and facts that are causally efficacious are themselves causally efficacious (Shafer-Landau: 2003), then normative facts would inherit the causal powers of their supervenience basis and CAUSALITY would be true of them. Or one can think that the best explanation of why beliefs with a certain content (i.e. that E is evidence for P is a reason to believe that P) are true requires appeal to normative facts of the form I am discussing here; that beliefs are ‘facts in the natural world’ and that whatever we must appeal to in order to explain facts in the natural world must itself be part of the ‘causal network’ of nature. Whatever side of these questions one is inclined to favour, they raise *substantive* philosophical issues. Pending further philosophical investigation, the answers to them ought not to be taken for granted.

5.3. Realism and Evidence-as-Reasons views

Let us turn, lastly, to consider what I have called ‘evidence-as-reasons’ interpretations of epistemic reasons. According to each of the two possible variations of this position – what I have called ‘analytical naturalist’ and ‘evidence-as-irreducibly normative’ views (see Chapter 2) – epistemic reason claims are equivalent in meaning to evidence claims. More precisely, epistemic reasons claims *are* claims about what is evidence for what (and vice versa). According to this view, to say that *E is evidence for P* is equivalent in meaning to saying that

ERV: [fact E] is/provides an epistemic reason to believe that P.

Standard realist views about epistemic reasons thus construed will hold that evidence/epistemic reasons claims, when true, are made true by facts. Again, this leaves open questions regarding the nature of the facts which make claims of the kind at hand true. I have already commented on the

implausibility of thinking that the facts which would make reasons claims true on this interpretation can be *subjective* in any way which could motivate rejection of the OBJECTIVITY thesis.

However, it is plausible to think that epistemic realists who subscribe to the ‘evidence-as-reasons’ view can disagree both over IRREDUCIBILITY and over CAUSALITY when these are understood as theses about (normative) facts of the kind at hand. Here again, the obvious candidate for reduction for normative facts would be facts about evidence understood in some not-irreducibly-normative way; for instance in terms of a frequency, or propensity, or logical theory of probability.⁸⁹ Surely, whatever one thinks about the prospects for a reduction along these lines will affect one’s views on the causal efficacy of normative facts, in the way I have already noted.

Depending on their details, some not-irreducibly-normative accounts of evidence-facts – in terms of probability, or in terms of probability and explanation combined perhaps (Achinstein: 1978) – will be easier to reconcile with the possibility that facts of this kind will be causally efficacious than other accounts. Thus, as I have already pointed out, it would be a mistake to assume that, were a reductive account of normative facts along the lines I am discussing to be successful, the truth of CAUSALITY would thereby have been established. On the other hand, and as we have already seen, resisting whatever not-irreducibly-normative account of evidence is put forward – i.e. holding on to the IRREDUCIBILITY thesis here – will still, in itself, be compatible with holding different views on the causal efficacy of evidence/normative facts themselves.

⁸⁹ Or in terms of the combination of probability and explanation. Also, one need not assume that evidence-facts must be given *an account* in some terms other than evidence itself. I.e. evidence-facts could be *irreducible*, and yet still not be *irreducibly normative*. The same considerations I explore above would similarly apply to the case of evidence thus construed.

6. A Spectrum of Epistemic Realist Views

We can draw together some of the strands of our discussion so far with the aid of a table. The table shows, under each different interpretation of epistemic reasons, which positions can be occupied by epistemic realists with respect to the IRREDUCIBILITY, OBJECTIVITY, and CAUSALITY theses ('Y' stands for 'acceptance' and 'N' for 'rejection' of the theses on the left-hand side of the table).

EPISTEMIC REALISM			
	<u>INSTRUMENTALISM</u>	<u>EVIDENTIALISM</u>	<u>EVIDENCE-AS-REASONS</u>
<u>IRREDUCIBILITY:</u>	Y/N	Y/N	Y/N
<u>OBJECTIVITY:</u>	Y/N	Y	Y
<u>CAUSALITY:</u>	Y/N	Y/N	Y/N

In my discussion so far, I have suggested that one useful way to classify epistemic-realist positions is on the basis of (1) which interpretation of epistemic reasons they are committed to and (2) which combination of views they hold with respect to the IRREDUCIBILITY, OBJECTIVITY, and CAUSALITY theses. I have also gone some length towards illustrating how acceptance/rejection of each of the theses I have listed, by itself, ought not to be taken as automatically entailing acceptance/rejection of any of the others. I have done so by drawing attention to some of the obvious and less obvious ways in which different stances on those theses can go together or come apart. My survey of how the different theses (on the nature of epistemic-normative facts) and interpretations (of epistemic reasons) can combine has not been exhaustive. It was not meant to be. Its purpose was rather to bring to the fore the complexity of views which ought to be up for consideration in discussions of epistemic realism.

The table above makes it easier to be alert to the wide spectrum and the broad variety of positions that – once in logical space, so to speak – would all have to count as epistemic realist views of some sort or other. It shows some of the variables that must be taken into account when one is enquiring after the content of specific epistemic-realist views. Awareness of this complex network of positions and of their relation will be of crucial importance when it comes to assessing the relative merits of various epistemic-realist views. It will clearly also be central to the formulation, discussion, and evaluation of many arguments against meta-normative position of the realist kind. An examination of one such argument will occupy the remaining chapters of this thesis.

Part II

Chapter 4

A Darwinian Epistemological Argument against Epistemic Realism

1. Introduction

Darwinian arguments against normative realism in the practical domain have been the object of much attention in recent metaethical literature. In the remaining chapters of this thesis I turn to investigate the extent to which counter-parts of those arguments can be formulated to call into question normative realism in the *epistemic* domain. As I hope to have made clear in Chapter 3, the label ‘epistemic realism’ is shorthand for a host of extremely diverse views. A discussion of Darwinian arguments against normative-realist positions in epistemology can hardly neglect this fact. For it will not only affect what the target of these arguments is considered to be. It will also determine which strategies of reply to those arguments, if any, are available. Such discussion must similarly take into account the distinctive shape that normativity takes in the domain of epistemology, particularly so when we reflect on the nature of the ties which hold between epistemic reasons and evidence, on the one hand, and between evidence and truth on the other. In what follows, I will then rely crucially (and build) on the investigations of the previous chapters.

I will proceed in two stages. The first stage will occupy this chapter and the next two chapters. In this chapter (Chapter 4), I begin by narrowing down the focus of my investigation to Darwinian arguments of a specific form; I will call these ‘epistemological’ Darwinian arguments. Arguments of this kind correspond to what are often referred to in the metaethical literature as ‘global evolutionary debunking arguments’ against normative realism (Kahane: 2011). I then provide an outline of a Darwinian-epistemological argument which targets epistemic-realism

(DEA), and provide a first intuitive characterization of how this argument is supposed to work. After this, I turn to discuss two crucial premises of the argument I have presented and devote a chapter to each. In Chapter 5 I discuss the thesis that we can provide evolutionary explanations for the epistemic-normative beliefs we hold – call this ‘the evolutionary hypothesis’. In Chapter 6, I discuss the claim that there would be a ‘lack of correlation’ between epistemic-normative beliefs and epistemic-normative facts, on the assumption that those beliefs can be explained in evolutionary terms. Investigating this claim will also provide the occasion to clarify which epistemic realist views are the intended target of the argument I focus on.

In the second and final stage of my discussion – which will occupy Chapter 7 – I engage with the Darwinian-epistemological argument against realism about epistemic normativity that Sharon Street presents in her paper ‘Evolution and the Normativity of Epistemic Reasons’ (2009). I argue that the argument which Street presents there is best viewed as an instance of the Darwinian-epistemological argument I focus on (DEA). I then go on to reply to Street’s argument on behalf of the epistemic realist. I conclude that arguments with the form of DEA do not represent a serious threat against normative realist views in epistemology.

2. Darwinian Arguments against Normative Realism

Despite having in common their reliance on evolutionary considerations of broadly similar kinds, Darwinian arguments against realist views about normativity have taken significantly different forms. Discussion of those arguments has been prominent in the *metaethical* literature. As such, it is not surprising that the focus of this discussion has been, most often, *moral* realism or realism about *practical* normativity more generally. Yet nothing in the basic structure of these arguments makes them specific to normativity in ethics or the practical sphere. And versions of the

challenges they raise have been also directed against realist views about normativity in other domains, *epistemic* normativity included.⁹⁰

In this section I discuss two broadly different kinds of Darwinian challenges to normative realism, what I will call the ‘metaphysical’ and the ‘epistemological’. My aim in doing so is to dispel possible confusions between the two, and to provide reasons as to why the latter is a worthier object of philosophical attention. In the next sections, I will then turn to discuss the epistemological challenge to normative realism in epistemology in more detail.

Challenges of the *metaphysical* kind build upon the claim that an appeal to normative facts of the kind realism believes in would appear to be redundant in evolutionary explanations of our normative beliefs and practices. This claim often comes down to the thesis that the best evolutionary explanation of those beliefs and practices requires no appeal to normative facts; and that normative facts are not reducible to facts which do figure in those explanations. On the basis of considerations of ontological parsimony (Ockham’s razor) – and/or drawing ontological consequences from the virtues of *simpler* and overall *better* theories – those arguments draw *ontological*⁹¹ conclusions. If this kind of arguments were to succeed, we would have positive reasons to believe that the kind of normative facts the realist posits do not exist (Ruse; 1986: 253-4), or no positive reason to believe in the existence of those facts (Joyce: 2013 and Gibbard; 1990: 108).⁹²

But Darwinian arguments can take a significantly different form – what I have called an *epistemological* form. Arguments of this latter kind begin by *granting* – for the sake of argument – that the normative facts that the realist posits exist. They then challenge the realist to explain how we could have (many) true normative beliefs, given that these beliefs are the products of

⁹⁰ See Street (2009), Nozick (1993), Parfit (2010), and Gibbard (1990; 2003).

⁹¹ As I have already drawn attention to in Chapter 3, Parfit’s version of normative realism – a view which he refers to as ‘Non-Metaphysical Cognitivism’ (2011: Ch.31, Sections 112-3) – is committed to the thesis that normative facts and properties exist in what he calls ‘a non-ontological sense’. The metaphysical (and epistemological) Darwinian arguments against normative realism which I discuss in this and the following sections are also meant to target this view.

⁹² Joyce (2013: 361) discusses at some length the relation between these two distinct claims. Their difference, and their relation, need not concern us here.

evolutionary forces that are indifferent or insensitive to the kind of facts that, according to the realist, make those beliefs true (Clarke-Doane; 2012: 313). Proponents of those arguments believe that the realist can't provide this explanation. On this basis, they draw *sceptical* conclusions. These conclusions might differ slightly in their formulation, but they all similarly call into question the justification for our normative beliefs and, hence, the possibility of acquiring or possessing knowledge about normative facts (Wielenberg; 2010). On those arguments, if we assume the truth of realism, we would have to believe that 'it is unlikely that our moral [read: normative] beliefs are true' (Street; 2006: 121-2)); or that our moral [read: normative] beliefs are unjustified (Joyce; 2006: 179-82); or that if our normative beliefs were true, this would at best only be the result of a 'massive coincidence' or a 'miracle' (Enoch: 2010); and so on.

Although arguments of this latter *epistemological* kind are essentially arguments *in favour of* sceptical conclusions, they are often deployed – in a broader dialectical context – to argue *against* normative realist views. The context in question is one in which competing anti-realist theories of normativity can be shown to fare significantly better at the explanatory task at which realism allegedly fails (Street: 2006), and thus to forestall the scepticism that realism would not be able to avoid.

The labels 'epistemological' and 'metaphysical' which I have used to refer to these different kinds of Darwinian strategies ought not to obscure the fact that there are some important respects in which those forms of argument are similar.

Epistemological arguments rely on the claim that the evolutionary forces that have deeply influenced our normative beliefs are *insensitive* or *indifferent* to the normative facts realism posits. I will return to discuss the meaning and plausibility of this claim at some length in the next section. Intuitively and for the time being it is helpful to think about it with the help of the following twin counterfactuals.⁹³

⁹³ I will return to discuss the role which these counterfactuals play in Darwinian-epistemological arguments at some length in Chapter 6. Notice how two obvious – but distinct – ways to read and make more precise counterfactuals of those kinds is in terms of the notions of (in-)sensitivity and (un-)safety which have been at the centre of recent

The first builds on the possibility that the normative facts might *easily* have been different from what they are. It says that had this been the case, evolution would still have led us to hold the normative beliefs we hold. Thus, even if most of our current normative beliefs were true, it would be a matter of luck that they are so.

The second relies on the claim that we could *as easily* have evolved in such a way as to hold different normative beliefs from the ones we hold, and that, in this case, it would be a matter of sheer luck if those beliefs corresponded to the kind of normative facts the realist believes in. Thus, it would be a matter of luck if those beliefs were true.

Notice how these insensitivity/indifference claims are *not equivalent* to a causal-explanatory claim of the kind that figures in the premises of metaphysical arguments. The first is about the absence of a stable correlation between normative beliefs and normative facts. The second is about the role of normative facts in the causal-evolutionary explanation for our normative beliefs. Yet the truth of the insensitivity/indifference claims is often defended by appeal to the idea that normative facts do not figure in the (best) evolutionary explanation of our normative beliefs and practices (Griffiths and Sterelny: 2010, Street: 2009, Enoch:2010, Joyce: 2013); and that normative facts are not reducible to entities and facts which do figure in those explanations. Thus, at least some versions of the epistemological argument implicitly rely on the kind of *causal-explanatory* considerations that metaphysical arguments also rely on.

Moreover, as I have already pointed out, epistemological arguments are often employed as parts of arguments which aim to establish *ontological* conclusions. When this is the case, they can be read as the first half of a *reductio* style of argument. They are the first – and crucial – step in a *modus tollens* strategy of the following kind: if realism is true, then – given ‘the facts of evolution’ – we can’t justifiably hold any of the normative beliefs we hold. But surely, the strategy goes on,

debates in epistemology. See, for instance, DeRose (2004) and Pritchard (2008). Proponents of Darwinian-epistemological arguments of the form I am discussing here typically rely on a much more intuitive understanding of the truth-conditions of these counterfactuals (see Gibbard: 1990, White: 2010, and Street: 2009). Given the point I am making in this section, the intuitive understanding will suffice.

we do justifiably hold at least *some* normative beliefs. Indeed, we hold *many* such beliefs. If it were to turn out that anti-realist theories did better at accounting for this latter fact (see above), and to avoid the undesired consequences of realism here, then we would have some reason to believe that anti-realist theories of normativity are much better theories. This, in turn, would provide reasons to believe that realism is false. Thus insofar as they are taken to articulate *a challenge for realism*, epistemological arguments can *also* be read as at least being directly associated with arguments the conclusions of which have *ontological* import. For, if the truth of realism is questioned by those arguments, so is the *existence* of normative facts.

I believe the considerations I have just reviewed partly account for why these two – broadly different⁹⁴ – kinds of Darwinian strategies are not always adequately distinguished in the literature.⁹⁵ This is somewhat unfortunate, for – once adequately formulated – the epistemological strategy would seem to raise a more powerful and quite distinctive challenge for realism about normativity.

A comprehensive and comparative discussion of the relative merits of what I have called metaphysical and epistemological Darwinian arguments exceeds the aim of this chapter. Yet, there are some obvious reasons for wanting to focus attention on evolutionary arguments of the epistemological kind. This is that metaphysical Darwinian arguments would appear either to collapse into more familiar arguments of a *non-Darwinian* kind against realism, or not to represent a serious threat for realism after all.

⁹⁴ For more on the differences between epistemological and metaphysical Darwinian strategies, see Lillehammer (2003) and Joyce (2013).

⁹⁵ This is reflected, for instance, in the fact that Joyce (2013: 356) attributes to Street an argument to the conclusion that ‘there are no objective moral facts’. Enoch reads Street as providing a paradigmatic instance of what he calls ‘the epistemological challenge to meta-normative realism’ (Enoch; 2010: 425-6). Street herself seems, at times, to oscillate between defending one version of the Darwinian challenge and endorsing the other (2006; 2009). And so does Gibbard: in *Wise Choices, Apt Feelings* he writes: ‘In my evolutionary speculation there were purported natural facts in abundance, but none singled out as peculiarly normative. [...]. We can take this speculation and see where it leads. That way, perhaps, we can illustrate *how biology could settle whether there are normative facts*. If the hypothesis holds good, we do not need normative facts to explain our making the normative judgments we do. [...]. *To suppose that there are normative facts is gratuitous*’ (1990: 107-108, italics mine). And a few pages later, in the context of the same discussion, he writes: ‘Even supposing there are normative facts, we can explain why we have the normative beliefs we do without citing any normative facts[...]. *Even if there were a special realm of normative facts, it would be gratuitous to think that normative beliefs represent facts in this realm*’ (1990: 121, italics mine).

For one, they seem to beg important questions against non-naturalist forms of normative realism, namely, those realist views which hold that normative facts are not causally efficacious. Notice how the claim that normative facts cannot enter into a certain kind of causal/scientific explanations – such as evolutionary explanations – is not likely to worry this brand of realism. In a sense, non-reductive realists *already believe* that this is the case. For if they are ready to believe that there are normative facts and that these facts are not causally efficacious, they will find it hardly surprising that facts of this kind won't figure in causal explanations *of the evolutionary kind*.⁹⁶

In any case, if this were the kind of challenge Darwinian-metaphysical arguments were meant to raise, then it is not clear how this would differ from more general, and more familiar, challenges to non-naturalist realism, such as those arguments which rely on the principle that for some kind of fact or entity to exist, they must be indispensable in (at least some kind of) causal explanations (Harman: 1977, Enoch: 2007). Call these *arguments from explanatory (in)dispensability*.⁹⁷ Moreover, it is not obvious how adding a 'Darwinian twist' to those kinds of argument would help one make progress on those other, more familiar debates.⁹⁸

Secondly, it is generally the case that metaphysical Darwinian arguments are intended to target so called *naturalist* and *non-naturalist* views of normativity alike (Gibbard: 1990, Street: 2006). Naturalist realists believe that normative facts are causally efficacious, and they often do so on the basis of the fact that, according to them, these facts do fulfil causal-explanatory requirements of the kind which warrant existential claims.⁹⁹ Of course, whether this latter claim is true has also been the object of important debates in the recent and not so recent metaethical literature (Harman: 1986, Sturgeon: 1986). But – as in the case with non-naturalist realism – it is not obvious what, if anything, appealing to evolutionary considerations can add to these other, and

⁹⁶ See, e.g. Enoch (2007), Huemer (2005), Parfit (2011), and Clarke-Doane (2012).

⁹⁷ Or arguments which are similar 'in spirit' to indispensability arguments, such as those which make existence of facts of kind N depend on their having *wide cosmological role* (Wright: 1992), or those which make the existence of N-facts depend on whether they figure in the *best explanation* of phenomena of kind P.

⁹⁸ If those other arguments are successful against the non-naturalist realist, then Darwinian arguments of the kind at issue are not needed. If they are not successful – i.e. if there were other grounds (other than causal-explanatory considerations) to believe in normative facts – then Darwinian arguments could not possibly carry much force either.

⁹⁹ See Boyd (2003), Brink (1989), and Sturgeon (1985). For a survey of defences of realism along these lines see Brad Majors (2007).

more familiar, debates. For even if normative facts were not to be *indispensable*¹⁰⁰ in the (best) evolutionary explanation of our normative beliefs and practices, and not *reducible* to facts that are, there might still well be *independent reasons* to believe that they exist. The most plausible Darwinian explanations for our capacities to form beliefs about black holes, electromagnetic fields, the big bang, molecules, etc. need not appeal to the existence of those things. Yet this is hardly the basis for believing that *there are no such things* (Copp; 2008: 190).

In general, if metaphysical Darwinian arguments of the kind at issue here are to have any force against realism, we ought to expect something like the following to be the case. Either:

(1) A strong/the strongest realist argument for the claim that there are normative facts would be that normative facts must figure in the best evolutionary explanation E1 for some phenomenon P.

or:

(2) A strong/the strongest realist argument for the claim that there are normative facts would be that normative facts need to figure in the best explanation E2 for phenomena P. E2 is *not an evolutionary explanation*. But E2 is such that, if explanation E1 for P turned out to be true, E2 would be ruled out by it. E1 is a much better explanation for P, it is an *evolutionary explanation*, and it is such as not to require any appeal to normative facts.¹⁰¹

On this latter scenario (i.e. (2)), E1 and E2 would be *competing* and *incompatible* explanations for P; and they would also be such as to appeal to different entities for explaining what they do explain (one appeals to normative facts, the other does not). The situation here would be similar to the case in which – as some philosophers and scientists believe – the truth of evolutionary

¹⁰⁰ Or have a ‘cosmological role’ that is ‘wide enough’ for this kind of facts to also figure in explanations of an evolutionary kind, etc.

¹⁰¹ Or to anything normative facts are identical with/can be reduced to.

theory would tell against God's existence (White; 2010: 583-4).¹⁰² Suppose our strongest reason for believing in God was that this is the only/best explanation for the (otherwise) mysterious fit between animal creatures' organs and habits and those creatures' natural environments. If it turned out that evolution provided the best/a much better explanation for the same phenomenon, this would appear to call into question our belief in God's existence.¹⁰³

But, as to (1), it is hardly the case that, in reviewing arguments in favour of normative realism, one comes across arguments that say that normative facts must exist *because we need to postulate them in order to provide an evolutionary explanation for P*, whether P be the current shape of our normative beliefs and practices, or anything else for that matter.

With regards to (2), much turns on the claim that the different explanations at stake there, namely E2 (the normative and non-evolutionary explanation) and E1 (the evolutionary and non-normative explanation), are *competing* and *incompatible* explanations of the same phenomena P. The first thing to notice here is that this claim is far from being obvious or *prima facie* plausible. If the analogy with our beliefs about molecules/big-bang, etc. above has any plausibility, it would rather seem to testify to the fact that E2 and E1 would be *independent* kinds of explanation. If they are independent in this sense, they can neither compete nor be incompatible. Secondly, the most obvious competing and incompatible explanations with E2 remain *non-evolutionary* explanations of the same phenomena, precisely those explanations that indispensability arguments of the *non-evolutionary* kind rely on. Thus it would seem that the most plausible argument in the vicinities of the one we are discussing would once again turn out to be the non-evolutionary kind of explanatory (in-)dispensability arguments that have already been at the centre of much debate in metaethics.

¹⁰² I am assuming – for the sake of argument – that the strongest reason for believing in God's existence is provided by the kind of considerations which the teleological argument for God's existence relies on (Paley: 2006).

¹⁰³ Of course, one must assume – as proponents of arguments against God's existence of this form typically do – that evolutionary explanations and theological explanations are *incompatible*; and this is a rather controversial assumption. I am merely relying on the God vs. evolution controversy as a useful analogy here.

Lastly, if – despite first appearances to the contrary – it did turn out to be the case that E2 and E1 are indeed competing and incompatible explanations, this would be a novel and substantive claim, a claim of the kind that would rather need to be supported by an *argument*; hardly a claim that can be left implicit or one to be taken for granted. Pending a precise articulation of this argument, Darwinian challenges of the metaphysical kind would not seem – by themselves – to carry much force.

As will become clear in the next section, epistemological Darwinian strategies do not inherit the weaknesses which affect Darwinian arguments of the metaphysical kind. The most charitable interpretation of the epistemological style of argument does not beg any interesting question against either naturalist or non-naturalist realist forms of realism. Neither does it covertly rely on – or collapse into – arguments which have already been much discussed in the metaethical literature. In this sense, Darwinian epistemological arguments would appear to raise a more serious and a *novel* challenge for realist views of normativity. I will now turn to discuss an argument of this kind; its target is realism about epistemic normativity.

3. The Darwinian-Epistemological Challenge to Normative Realism in Epistemology

3.1. The Basic Form of the Argument

The epistemological argument I will examine and discuss in the remainder of this chapter (DEA) has the following basic structure:

DEA:

(1) *Empirical Premise - Evolution*: Evolution explains why we hold the epistemic-normative beliefs we hold.

(2) *Metaethical Premise I – Epistemic Realism*: There are epistemic-normative facts.

(3) *Metaethical Premise II* – At least some epistemic-normative beliefs are true and when they are true, they are true in virtue of corresponding to epistemic-normative facts.¹⁰⁴

(4) *Epistemological premise I*: If evolution explains why we hold the epistemic-normative beliefs we hold, then those beliefs are not correlated with (i.e. they are off-track with respect to) epistemic-normative facts. This is tantamount to saying that evolution is an ‘off-track’ causal influence with respect to epistemic-normative facts.

(5) *Epistemological premise II*: If epistemic-normative beliefs are not correlated with epistemic-normative facts, then (our awareness of this fact makes it so that) our epistemic-normative beliefs are not justified.

Therefore:

(6) *Conclusion - Skepticism*: Our epistemic-normative beliefs are not justified.

Versions of this challenge are put forward or discussed in – among other places – Street (2009), Parfit (2010), and Nozick (1993). Implicitly, an argument of this form can be found in Gibbard (1990 and 2003), where what is at issue is rather realism about normativity in general. The kinds of considerations the argument relies on are also defended in the work of many evolution-friendly and naturalistic accounts of epistemic normativity, such as David Papineau’s discussion of epistemic reasons in *The Roots of Reason* (2006).

The argument as I have presented it is, however, extremely sketchy and more needs to be said in order to make its premises begin to look plausible. For one, as I have already discussed, the basic form of standard epistemic normative beliefs is something that can hardly be taken for granted (see Chapter 2). If this is so, then it is not clear what evolution is called upon to explain

¹⁰⁴ As I have already pointed out in Chapter 3, on most standard definitions of normative realism in the philosophical literature, the thesis expressed by *Metaethical premise II* would count as part of the definition of normative realism itself (for a version of this way of understanding what epistemic realist positions in epistemology amount to see Cuneo (2007)). I did not include this thesis in my definition of realism in order to allow for the possibility of normative-realist views which subscribed to an *expressivist* semantics of normative claims. It is an interesting question whether Darwinian arguments of the kind I discuss in this thesis can be thought to also raise a challenge for those latter views. I will not pursue this question further in what follows.

to start with. What realism about epistemic normativity amounts to – and hence what the target of the arguments is – is similarly far from obvious (see Chapter 3).

Moreover, the empirical premise, as it stands, is rather compressed and quite puzzling. Surely, there are many epistemic normative beliefs that evolution cannot possibly have ‘caused’ in any straightforward sense: such as, for instance, that evolutionary-epistemological arguments provide reasons to believe that our epistemic-normative beliefs are unjustified (if realism is true). If this is the case, on a rather obvious reading of its meaning, the empirical-evolutionary premise would appear to be a non-starter.

Again, what it would be for a process such as evolution to be ‘off track’, and what the reasons are for thinking that evolution is one such process – need to be made more precise. I will return to discuss these issues at some length shortly. Before doing so, however, it will be worth having an initial sense – if a very general and intuitive one – of how the argument I have outlined is supposed to work. This will also help clarify how the argument differs from other Darwinian arguments and why – as I have suggested – it improves on them.

3.2. The Nature of the Argument: An Analogy

The basic thought is the following. Suppose we held beliefs of kind C – beliefs about numbers or about colours, for instance. Let us assume that beliefs of this kind can be true and that when they are true, they are made true by facts – call those C-facts. Let us also take C-beliefs – as we normally would with beliefs about numbers and colours – to be also such as to be *correlated* to the facts which make them true in a way that ensures that, most of the time, when we judge that P (where P is a C-belief) it is indeed the case that P, and that – because of this – many of the C-beliefs we hold are true.

Correlation in the sense at issue here has counterfactual force. We often express the presence of a correlation of the kind at issue here by saying that ‘had the C-facts been different, our C-

beliefs would have been, by and large, correspondingly different'; or that 'given that the C-facts are what they are, many of our C-beliefs could not have easily failed to match them'.

Suppose it turned out that the way we form our C-beliefs had been tampered with. We are told that we have been forming C-beliefs while being under the effect of a drug that is such as to affect our mathematical/colour discrimination abilities. When it does so, it causes us to have C-beliefs the content of which is randomly determined. If this is the case, then C-beliefs bear *no correlation* to the C-facts: when we believe that P (a colour/mathematical-belief), we would have believed that P whether or not P was the case; and P is such that – even if it were true – it could easily have been false. Or, given that that Q (a colour/mathematical-fact) is the case, it would be a coincidence if we believed that Q, for we could easily have formed a belief other than the belief that Q – such as the belief that P, or R – instead.

In a situation like the one I have just described, a causal influence acting on the formation of our C-beliefs – the drug – is strong enough to affect their content, and it is such as to be insensitive or indifferent to the C-facts in the way I have just explained. The operation of the drug is such that it would be a rather surprising and accidental fact if most of our C-beliefs turned out to be true. And learning that such a causal influence has been in place ought to affect how we think of the *epistemic status* of our C-beliefs. In particular, we ought to lose confidence in the claim that our C-beliefs are justified.

This can happen in two very different ways, depending on the notion of justification in play.¹⁰⁵ On one popular notion of justification – a broadly externalist one – the kind of correlation that is undermined by the operation of the drug is *necessary* for justification to be in place. On this notion of what is required for one's beliefs to be justified, upon learning about the drug we ought to believe that the colour/arithmetical beliefs we formed after the drug had become effective were not justified. Whether we were aware of this fact or not, *the moment the drug started being effective*, our beliefs lost their justification.

¹⁰⁵ For a similar point see Kahane (2011), Enoch (2010), and White (2010).

But one might think that justification – in general – *does not depend* on a *correlation* of the kind I have discussed, and which the drug in our example undermines. On a different – internalist – notion of justification, for instance, our arithmetical beliefs are justified *by their being self-evident*¹⁰⁶ to us; while our colour beliefs are justified *by things appearing to us to be a certain way*. Even on this notion of justification, however, upon learning that our C-beliefs have been under the influence of the drug, we ought to revise our beliefs regarding their epistemic status. For *once this fact has been called to our attention*, we cannot retain confidence in our C-beliefs any longer. In this case, it is the very learning about the drug that makes it so that our N-beliefs lose their justification. On this picture, and to draw a contrast with the previous case, C-beliefs lose their justification not the moment the drug starts being effective, but *the moment we learn of the fact that a drug of the kind in question has become effective*.

Thus on each of the *internalist* and the *externalist* notions of justification, it will be true that once the fact that the drug is in operation has been called to our attention, our C-beliefs will lose their justification. And this will call for a revision of our beliefs regarding the epistemic standing of our C-beliefs. If we previously thought they were justified, we now ought to think that this is no longer the case.

We can draw on the case I have just discussed in order to get a sense of the kind of challenge Darwinian-epistemological arguments are meant to raise. The proponents of these arguments grant the realist that epistemic-normative facts of the kinds he believes in exist and that epistemic normative beliefs, when true, are made true in virtue of corresponding to those facts. They assume that the realist would want to account for the fact that – as we take it to be the case – many of our epistemic-normative beliefs are true, and that he would have a favourite explanation of some sort of why this is the case. A naturalist-realist might believe that we are *causally* affected

¹⁰⁶ The notion of ‘self-evidence’ I am relying upon here must be understood as being *non-factive*. The notion in question can be thought of in terms of accounts of *a priori* justification according to which this source of justification is fallible. See, for instance, Bonjour, L. (1988).

by these facts in the way we are causally affected, say, by the colours of objects.¹⁰⁷ A non-naturalist-realist will believe that we have *a priori access* to the kind of facts that epistemic-normative facts are taken to be, in a way which mirrors what is – on some account – our epistemic access to mathematical truths.

The proponents of the Darwinian-epistemological argument then go on to claim that the fact that evolution can be appealed to in order to explain why we hold the epistemic-normative beliefs we hold puts the realist in a situation which would be relevantly similar to that in which we find ourselves with respect to our C-beliefs in the scenario I have described above.

They hold that evolution is a causal influence on our normative beliefs which is strong enough to affect their content, and such as to be insensitive or indifferent to epistemic-normative facts of the kind the realist believes in. The work of evolutionary forces on our epistemic-normative beliefs would make it so that there is *no correlation* between epistemic-normative beliefs and epistemic-normative facts themselves. Because of this, it would be a rather surprising and accidental occurrence if most of our epistemic-normative beliefs turned out to be true.

Drawing this fact to the attention of the epistemic-normative realist, they argue, ought to affect how the realist thinks of the epistemic status of his epistemic-normative beliefs. In particular, on either internalist or externalist accounts of justification, the realist ought to lose confidence in the claim that his epistemic-normative beliefs are justified. If he previously thought they were justified, he now ought to think that this is no longer the case.

¹⁰⁷ I am – for the sake of argument – assuming that some (non-dispositionalist) version of physicalism about colour is true here (Byrne and Hilbert: 2003).

3.3. How the Argument Differs from Metaphysical-Darwinian Arguments and from other Epistemological Challenges

Notice the following things about the general structure of the argument I am discussing. Firstly, it does not – at least not overtly – rely on the claim that epistemic-normative facts are *explanatorily inefficacious* or *explanatorily redundant* in evolutionary explanations. Also, unlike metaphysical arguments, it does not draw *ontological* conclusions – that there are no epistemic-normative facts or that we have good reasons to believe that facts of this kind do not exist.

Secondly, it is distinct from epistemological arguments which claim that we can't be justified in our beliefs about (epistemic) normative facts because, for this to be the case, epistemic-normative facts must be able to *cause* our epistemic-normative beliefs and epistemic-normative facts are, according to the realist, not causally efficacious (Enoch: 2010). There is nothing specifically *Darwinian* and *novel* about this latter kind of argument. They resemble metaphysical-Darwinian arguments in this respect: the most plausible arguments in their vicinity are *non-Darwinian*, and they coincide with arguments which have already received much attention in the philosophical literature – such as those according to which the faculty of rational intuition that would be required for us to be justified in holding our normative beliefs would be mysterious. They also, obviously, do not affect *naturalist* forms of (epistemic) normative realism, while the argument I have outlined is meant to challenge naturalist and non-naturalist versions of (epistemic) normative realism *alike*.¹⁰⁸

Unlike metaphysical-Darwinian arguments and unlike the epistemological argument I have just mentioned, in the argument I have outlined the Darwinian element (the empirical premise) plays *a crucial role* and makes it so that the challenge it raises for realism is a *novel* and *distinctive* one. The argument would fail if its evolutionary-empirical premise turned out to be false. While the failure of this argument would not in any way affect the outcome of the other more familiar debates. The argument relies on the idea that, once some plausible empirical claims regarding the

¹⁰⁸ This is explicitly so – for instance – in Street (2006), where a paradigmatic instance of an argument of this kind can be found.

influence of evolution on our epistemic-normative beliefs are called to the attention of the realist, the realist ought to question his justification for these beliefs. And this will be the case whatever his justification for holding those beliefs *up to that point* might have been.

In the next chapters I turn to look at two crucial premises of the DEA. I will begin with the *empirical* premise of the argument (premise (1)) – i.e. the thesis that evolution explains why we hold the epistemic-normative beliefs we hold. I will then look at epistemological premises, in particular at premise (4), i.e. the thesis that if evolution explains our epistemic-normative beliefs, then those beliefs are not correlated to the epistemic-normative facts which realism posits.

Chapter 5

The Evolutionary Hypothesis

1. Evolutionary Explanations of Beliefs

How must we understand the claim that evolution explains why we hold the epistemic-normative beliefs we hold? I suggest we begin to make sense of it by looking once again at a different class of beliefs – colour-beliefs or C-beliefs for short – and think about what it might mean to say that evolution explains not only why we hold beliefs of this kind C, but also why we hold beliefs of kind C with a specific content (P, Q, R, etc.) in given circumstances.

Take the belief that X is red, which person A forms upon looking at a red object. It is most natural to think that in normal conditions and for a person whose visual apparatus is functioning as it should, the belief that X is red is caused *by the fact that X is red*; and that what explains this belief is *the fact that X is red*.

A more general form this explanation might take appeals to the presence, in creatures like A, of cognitive mechanisms and dispositions – call them CM – which are responsible for the ability to form true colour beliefs, and to the fact that A's forming the belief that X is red is the result of the workings of mechanisms and dispositions of this kind, in a specific instance.

The explanations I have briefly sketched above are not *evolutionary* explanations. In what sense, then, might one think that *evolution* explains why we hold beliefs such as these? Let us assume – as it is somewhat plausible to do – that the ability to discriminate correctly the colours of objects in a creature's most immediate environment increased the chances of reproduction and survival of the creatures who displayed this ability, thus conferring on those creatures an advantage over their competitors who did not display the ability in question. Let us also assume

that an ability of this kind can be passed on to following generations of organisms of the same species. If this were the case, the *increase in fitness* that this ability conferred to the creatures that had it would explain – in *evolutionary* terms – the widespread presence of this ability in a population of creatures of the kind we are considering.

Not all evolutionary explanations need be of this form.¹⁰⁹ The relevant notion at work in the kind of evolutionary explanation I am considering here is that of *an adaptation*.¹¹⁰ Let us take the following as a working definition for what an adaptation is:

Adaptation: Characteristic D is an adaptation for doing task T in a population if and only if members of the population now have D because, ancestrally, there was selection for having D and D conferred a fitness¹¹¹ advantage because it performed task T. (Sober: 1984)

And let us call performing task T the ‘natural function’ of D (Gibbard: 1990).

In the case of colours it is plausible to say that the mechanisms/set of dispositions (CM) that make for the ability to discriminate colours correctly constitute *an adaptation*, and that their *natural function* (I) is that of enabling the creatures which possess them to discriminate colours correctly.

We can further assume – as it would be plausible to do – that if the ability *to form true colour beliefs* is not identical with the ability *to discriminate colours correctly* which C confers, it bears close links to

¹⁰⁹ Other biological mechanisms and processes which are appealed to in evolutionary explanations include: genetic drift (Avers: 1989); lateral (or horizontal) gene transfer (Goldenfeld and Woese: 2007); ‘Bauplan’ constraints (Gould and Lewontin: 1973); niche construction (Day, Laland and Odlee-Smee: 2003); phenotypic plasticity (Pigliucci: 2001).

¹¹⁰ It seems to me that Darwinian-epistemological arguments against moral realism (Street: 2006) would work equally well if evolutionary forces other than natural selection were also recognised to have had a significant influence on the moral beliefs many creatures like us hold. A similar thing cannot be said for Darwinian-epistemological arguments against *epistemic-normative* realism. This asymmetry has got to do with the specific features which evolutionary explanations of our epistemic-normative beliefs must display in order for them to play a role in argument such as DEA. I cannot elaborate on this topic at any length here. I will return to discuss the nature of evolutionary explanation at work in DEA in the next section.

¹¹¹ I take the notion of fitness at work in this definition to comprise the notion of *inclusive fitness*. For one of the earliest formulations of this concept, see Hamilton (1963). For useful discussion of this concept see Sober and Wilson (1999).

the latter ability. We can thus speak, perhaps somewhat loosely, of the ability to form true colour beliefs as being itself an adaptation.

We can now see what it would be to explain in *evolutionary terms* why A forms the belief that X is red upon visually attending to a red object in normal circumstances N. It would be to say that creatures from which A descends and which – in ancestral environments – exhibited the disposition to discriminate the colours of objects correctly had a fitness advantage over creatures that did not, and that this explains the presence in A (and creatures like A) of the ability to form true colour beliefs. A's (responding to the fact that X is red by) forming the belief that X is red can then be seen as a specific instance in which this ability is exercised. The same form of explanation can account for why – in a large number of circumstances – A and creatures like A hold the colour-beliefs they hold.

2. Evolutionary Explanations of Epistemic-Normative Beliefs

2.1 Their Explanandum

The style of evolutionary explanation that I have just drawn attention to is also at work in the claim that *evolution explains why we hold the epistemic-normative beliefs we hold*. Before I say more about how the explanation at issue must be understood, it is worth being clearer about what their intended *explanandum* is.

The first thing to notice here is that, like in the case of colour-beliefs I have discussed above, what evolution is being called in to explain here is not merely why creatures like ourselves exhibit the disposition to form epistemic normative beliefs *of some sort or other*. The claim is the rather stronger one that evolution can explain why most of us are inclined to form epistemic normative beliefs *with a specific content* in given circumstances; and that it can do so in a rather large number of cases (Wielenberg; 2010: 444).

Secondly, in Chapter 2 I have drawn attention to the very different ways in which the basic structure of epistemic-normative beliefs can ultimately be understood. Clearly, settling this matter one way or another will ultimately affect what we must take the *explanandum* of evolutionary explanations of epistemic-normative beliefs to be. Among the writers who clearly think that evolutionary considerations can be drawn upon in arguing against versions of normative realism in epistemology, only Street (2009) shows some awareness of this question, and she discusses separately the evolutionary explanation of epistemic-normative beliefs of roughly the forms I distinguish (in Chapter 2). Gibbard's discussion of these matters is – as I have already pointed out – more general in form. Not only does he not distinguish between epistemic-normative beliefs of different forms, but quite often he treats epistemic-normative beliefs and beliefs about what we have reasons 'to feel and do' under the same general heading.

For the time being, I suggest that we set aside the specific differences between different interpretations of the structure of epistemic-normative beliefs, and that we focus on some more general features of the evolutionary explanation of epistemic-normative beliefs themselves. We can take beliefs of the form 'that E is a reason to believe that P' as our provisional template for standard epistemic-normative beliefs. To a first approximation, then, the claim that evolution explains why we hold the epistemic-normative beliefs we hold would amount to the following. In a large and varied number of circumstances, evolution can explain why creatures like us hold instances of beliefs of this more general form and with a given specific content – such that G is a reason to believe that Q, F is a reason to believe that R, and so forth.

Notice lastly how the claim that evolution can provide explanations of this kind is independent of – and is not meant to entail – the claim that the people who hold given epistemic-normative beliefs must also hold beliefs of a certain kind *about what these epistemic-normative beliefs are like*. In particular, the thesis I am discussing here must not be confused with the different thesis that evolution explains why we are disposed to form what we might call 'second-order beliefs' about the (epistemic) normative beliefs we hold, such as that when these

beliefs are true, they are *made true by facts of a certain kind*, or that when they are true they are *objectively*, or perhaps, *necessarily* so, etc. Some versions of Darwinian epistemological-debunking arguments follow a different route from the argument I am discussing here.¹¹² They work by defending the (allegedly empirical, evolutionary) claim that the belief in *objective/necessary/known a priori*, etc. normative facts is, to use Michael Ruse's expression, 'an illusion foisted upon us by our genes'¹¹³ (Ruse; 1986: 253). They rely upon the idea that we would have believed that there are normative facts *of these kinds* whether or not it is indeed the case that normative facts are so. The realization of this fact – these arguments claim – ought to bring with it that belief in objective/necessary/known a priori, etc. normative facts is unjustified.

What is thought to follow from this conclusion varies. If, say, we believe the *objectivity* of normative facts is 'written into the phenomenology of moral [normative] experience' (Joyce; 2103: 356), so to speak, and that it is 'an essential quality' of normative beliefs and concepts (Joyce; 2013: 357, Ruse; 2009: 507, and Ruse; 2006: 22), then it is the justification for our normative beliefs themselves which the arguments call into question. If, however, we believe that the *objectivity* (*necessity*, *a priori nature*, etc.) of normative facts is not 'written into' the normative concepts we use, then what we are not justified in believing are our second-order beliefs that normative facts are *objective* (or *necessary*, or *a priori*, etc.). And the fact that the objectivity (*necessity*, etc.) of normative facts would be the product of an 'illusion' could then be put to the service of an argument in favour of the view that normative/moral facts are *not objective* in the way we are inclined to believe.¹¹⁴

The epistemological-Darwinian argument I am discussing here is quite different from those arguments. In particular, the evolutionary premise it relies on need not (and does not)

¹¹² Those kinds of arguments are also often confused with arguments such as DEA. For Darwinian arguments of the kind in question see Ruse (1986), Parfit (2010), Nozick (2003), but also Street herself (2009).

¹¹³ Michael Ruse explicitly emphasises the adaptive importance of the *objectivity* with which moral prescriptions are infused (1986: 253).

¹¹⁴ Robert Nozick follows a strategy of a similar kind when he argues against the belief that facts about *epistemic reasons* are *necessary* and *knowable a priori*, leaving it open that they might be identical to natural facts of some kind or other (see Nozick: 2003).

incorporate any substantive second-order assumption (such as the ones I have discussed above) about what the epistemic-normative beliefs that evolution explains are like. All it says is that evolution can explain our disposition to form beliefs of the form *that E is a reason for believing that P* (i.e. that A's fingerprints are on the gun is a reason for believing that A is responsible for the crime) in certain circumstances (for instance upon coming across and attending to E). The argument then works by saying that if we were to understand the truth of these beliefs in terms of their correspondence to the kind of epistemic-normative facts the realist believes in, given the (off-track) nature of the evolutionary forces which have shaped those beliefs, those beliefs would be not be justified.

With these qualifications about their *explanandum* in mind we are now ready to look at evolutionary explanations of epistemic-normative beliefs in more detail.

2.2. The Explanations at Work

As I have already pointed out, the claim that evolution explains why we hold the epistemic-normative beliefs we hold, as it stands, does not look too plausible. For one, it is controversial whether we are warranted in believing that our remote ancestors possessed the concept of an epistemic reason (Street; 2006: 118-9, Copp; 2008: 188). Philosophers who write on these topics would seem to agree in thinking that the concept of an epistemic reason appeared quite late in the history of humankind when – as Parfit puts it – ‘some of us asked whether and how our beliefs can be justified’, and that ‘[epistemic normative] beliefs seem to have too recent an origin to have been [directly] selected for’ (Clarke-Doane; 2012: 318).

Secondly, some of the specific epistemic-normative beliefs we currently hold¹¹⁵ – such as, for instance, that the fossil record provides a (good) reason to believe in the truth of Darwin's theory of natural selection – can hardly be thought of as being themselves the direct product of

¹¹⁵ Or that we are disposed to hold after having been exposed to a certain kind of information and training.

evolutionary forces.¹¹⁶ In this respect, beliefs of those kinds are quite unlike colour (and other kinds of) beliefs. In short, they do not belong to a class of beliefs which, plausibly, it would have been advantageous to have in ancestral environments.

Thirdly, the vast majority of even the simpler epistemic-normative beliefs we currently hold is such that *the content* of these beliefs cannot be in any obvious way accounted for in evolutionary terms, in the sense that the kind of things (objects, properties) which these beliefs *are about* could not possibly have been around in the environment in which our ancestors struggled to survive and reproduce. A case in point is the belief that the fact that the TV has been switched on, it is plugged in, and it is not working is a reason to think that it is broken. Thus if the evolutionary premise of the argument we are considering has to be made more plausible, the kind of evolutionary explanation it appeals to would have to be, at best, quite *indirect* and rather *complex*.

The literature on evolutionary accounts of normativity – and Gibbard’s and Street’s writings are no exception – is not particularly helpful when it comes to specific discussions of evolutionary explanations of *epistemic-normative beliefs* and of *epistemic normativity* more generally.

The details of Street’s account rather focus on the evolution in our ancestors of what she calls ‘basic evaluative tendencies’ towards instances of *action-behaviour* of certain kinds, such as the tendency we can find in our close relatives, the chimpanzees, to experience in some ‘primitive, non-linguistic sort of fashion’ – certain things in the world as ‘*calling for* or *counting in favour of* certain reactions on their part’ (Street; 2006: 116). The examples she provides have got to do with *actions* and *motivations* of some sort or other. To this extent, she writes:

Like us, individual chimpanzees seem to experience – at some basic *motivational* level – *actions* that would promote their survival or help their offspring as in some way ‘called for.’ (2006: 117, italics mine)

¹¹⁶ This seems to be true not only if by ‘evolutionary forces’ we mean ‘natural selection’, but also other modes of evolutionary explanation which I have briefly considered in one of my previous notes.

Further, Gibbard's evolutionary explanation for the emergence of what he calls a 'capacity for normative governance' (1990: 72) in creatures like ourselves emphasises as the primary function of this capacity that of 'coordinating our *actions* in mutually beneficial ways' (Gibbard; 1990: 64-82 and 107, see also Lillehammer; 2002: 569).¹¹⁷

The following seem to me a plausible reconstruction of what both Gibbard and Street have in mind when they defend the view that evolution explains our epistemic-normative beliefs.¹¹⁸

First natural selection led to the development of some basic psychological capacities and dispositions to form beliefs (or belief-like representations) in response to circumstances of specific kinds. There are certain habits of belief formation that we have reason to believe would have been adaptive among our ancestors. In this first stage humans evolved to be disposed to form beliefs in this way. An example of a disposition of this kind is the disposition to form the belief that the next tiger will be carnivorous upon having previously observed a large number of carnivorous tigers (Street: 2009).

Second the capacities and dispositions that developed in the first stage affected the development of the conceptual apparatus of less remote ancestors. More specifically, they led to the emergence of epistemic-normative concepts and affected the content of the epistemic-normative beliefs which those ancestors came to be disposed to hold/form. As a result of this, creatures like ourselves exhibit a non-reflective tendency or disposition to form basic epistemic-normative beliefs with a given specific content. It is in this sense, then, that a disposition of the kind I have mentioned above can be appealed to in order to explain¹¹⁹ why, for instance, we are

¹¹⁷ The notion of 'planning' also plays a prominent role in Gibbard's evolutionary account of our normative discourse and practices. And while the notion of 'planning' rather straightforwardly applies to *actions* (as shown by expressions such as 'I plan to do X') it is not obvious that it applies in a similar uncontroversial way to *beliefs* (we would hardly say 'I plan to believe P').

¹¹⁸ The reconstruction I provide below follows that provided in Copp (2008) rather closely.

¹¹⁹ Street suggests that we understand explanatory claims such as these in terms of the truth of the following counterfactual: had these basic tendencies and dispositions (what she calls 'adaptive-links') been very different, our normative beliefs would also have been very different and in loosely corresponding ways (2006: 120 and footnote n.21). One can think of the relation between these adaptive links and full-fledged (if basic) normative beliefs on the blueprint of the explanatory relation which holds between the basic-mechanisms which result in the ability to discriminate colours and full-fledged colour beliefs. We would not believe that X is red, unless the output of these

disposed to form the belief that having previously observed a large number of tigers which were carnivorous provides a reason to believe that the next tiger will be carnivorous.

The account so far relies on the idea that some of the basic mechanisms or dispositions to form beliefs in certain ways are *adaptations*, and that some of the basic and unreflective tendencies to form epistemic-normative beliefs of a certain kind which we exhibit can be directly traced back to adaptations of the kind at issue. An account of this form, however, is unlikely to get one very far if it is taken too literally. As I have already pointed out, most of the epistemic-normative beliefs we are disposed to hold are not about tigers and predators, or whatever kind of things were likely to be around in ancestral environments. And the way we form beliefs in most circumstances very rarely follows rules as simple as the version of the rule of enumerative induction that Street uses in her examples (Street; 2009: 237-238). Moreover, the more sophisticated of the epistemic-normative beliefs we hold – such as, for instance, those regarding the reasons we have to believe in scientific theories – are the result of complex reflection, both at the individual level and in conversation with others.

In order to extend¹²⁰ the claim that evolution can explain why we hold the epistemic-normative beliefs we hold to beliefs of this more complex and reflective kind, we need something like the notion of an ‘indirect evolutionary explanation’. To say that there is a *direct* evolutionary explanation for a trait D is to say that trait D is *an adaptation*. To say that there is an *indirect* evolutionary explanation for trait D’ is to say that it is possible to identify some trait D which is such that (1) D is an adaptation and (2) D’ and D are linked in such a way that creatures (in the relevant population) that have D are also likely to have D’ (Wielenberg; 2010: 443).¹²¹

mechanisms was what it is. But this is consistent with believing that the output of mechanisms of this kind falls short of providing a complete explanation of fully-fledged colour beliefs. See, for instance, Fodor (2000).

¹²⁰ That Street takes her evolutionary explanations to extend to beliefs of this more complex/reflective kind is implied by what she says in both (2006) and (2009). But it is also required – it seems – by the fact that the evolutionary influences on our epistemic-normative beliefs must be taken to be *strong* and *pervasive* for an argument like DEA to work. I discuss the meaning of the claim that evolutionary influences are *strong* and *pervasive* in the next section.

¹²¹ See also Huemer (2005: 216).

Another way to express the idea that D' is amenable to an *indirect* evolutionary explanation is to say that D' is a *by-product* of a trait D which is itself an *adaptation*.

In the light of my discussion so far, we can then go back to the empirical premise of DEA and provide the following interpretation of what the evolutionary thesis it expresses amounts to:

Empirical premise – Evolution: A significant portion of the epistemic-normative beliefs we hold is such that the mechanisms/set of dispositions which result in the fact that creatures like us come to hold those beliefs can be accounted for in terms of either *direct* or *indirect* evolutionary explanations.

3. *Strength and Pervasiveness* of Evolutionary Forces

Notice the following important point about this revised formulation of the argument's first premise. The influence of evolutionary forces on our epistemic-normative beliefs is said to affect a *significant portion of them*. What this claim must be taken to amount to – I suggest – is best captured by saying that (1) the *range* of epistemic-normative beliefs that evolution explains is quite *varied* and (2) the *number* of epistemic-normative beliefs which evolution explains is also quite *large*. In other words, the influence of evolution on our epistemic-normative beliefs must be understood as being – as I will put it – rather *strong* and *pervasive*.

A qualification of the kind I have just drawn attention to is needed if the first premise of the argument – as its most natural reading suggests – must be taken to apply to epistemic-normative beliefs *in general*, rather than for instance, just to a sub-set of those beliefs.

The qualification at issue is also *required* if an epistemological argument of the kind I am discussing here is to work. The argument – if we remember – works by drawing our attention to the fact that a given class of our beliefs has been formed under a certain causal influence, and that the causal influence in question is an off-track causal influence with respect to the kind of

fact that makes those beliefs true. It then concludes that, because of this, our justification for the class of belief in question – whatever the source for this justification normally is – is undermined.

Let us go back to the example I have introduced to illustrate the nature of the argument – that of the drug affecting the formation of our mathematical and colour beliefs. Clearly, given the influence of the drug on our beliefs, and given that the influence is off-track, scepticism about colour/mathematical beliefs will depend on just *how strong* the effect of the drug is. Suppose that the influence of the drug is negligible;¹²² or that the drug affects – say – only particularly fine-grained colour beliefs; or only rather complex mathematical beliefs. In this case, the extent to which our justification for those beliefs is called into question – if at all – will have to be sensitive to these facts. For instance, if we know/have reason to believe that the drug only affects a certain sub-class of our colour/mathematical beliefs, it would seem reasonable only to question our justification for those sub-classes of beliefs. For all the other beliefs of the relevant kinds, we can go on trusting our perceptual faculties (colour-beliefs) or a priori intuitions (mathematical beliefs). We can thus say that for the influence of the drug to be such as to undermine the justification for colour/mathematical beliefs *in general*, it will have to be such as to affect *a significant portion of them*. The beliefs that the drug affects will have to be *many* and *varied in kind*. In this sense, then, the influence of the drug will have to be strong and pervasive.

‘Significant portion’, ‘many’, and ‘varied in kind’ are rather vague and highly context-sensitive notions. How large and how varied a set of beliefs of kind C need to be for it to count as ‘significant’ in the relevant sense will presumably depend upon various parameters, for instance on how exactly the target kind of beliefs is specified;¹²³ or on how high the standard of

¹²² For instance, the drug only works in 1 per cent of the cases in which it is administered; or it is only likely to affect very few of our colour-beliefs, if any at all, etc.

¹²³ For instance: how fine-grained are the colour-beliefs which belong to the target set? For what purpose/task do we want to have accurate colour beliefs? And what counts as ‘accurate’, given the purpose/task at hand? Etc.

justifications are set for the specified target of beliefs in a given context. Settling those issues in more detail need not concern us here.

For the purposes of our discussion, I suggest that we understand a causal influence CI on our beliefs of kind C as being a *strong and pervasive* causal influence on those beliefs iff, when (we know/justifiably believe that) CI is off-track, this fact is such as to call into question our justification for C-beliefs.

3.1 Stronger and Weaker Evolutionary Hypotheses and Their Degrees of Confirmation

The qualification I have discussed above – however context-sensitive are the terms it has been couched in – is far from trivial. For one, as I have just pointed out, the epistemological argument relies crucially on it. Secondly, it reveals *how strong* the evolutionary-explanatory claim contained in the first premise of the argument is. Thirdly, and perhaps most importantly, it allows us to bring to view a possible line of criticism of Darwinian-epistemological arguments which is often overlooked in the philosophical literature.

While this is not the line of criticism to the arguments which in the end I will pursue, it is a line of criticism worth emphasising nonetheless. Drawing attention to it will also serve the purpose of further elucidating the nature of the empirical premise of the DEA.

One might think that the epistemic-normative beliefs we hold – when very broadly construed – are subject to a wide variety of seemingly heterogeneous (causal) influences: ‘ideology’, economics, sociology, ‘meme’ dynamics,¹²⁴ contingent historical events, ‘fashion’, the mechanisms studied by psychoanalysis, etc.

The claim that *evolution* explains, in the sense I have discussed, our epistemic-normative beliefs entails a claim about the relative importance of evolutionary forces with respect to those other

¹²⁴ For the concept of ‘meme’ see Dawkins (1976). For the application of models and methods drawn from evolutionary theory to the study of the evolution of ideas, scientific theories, epistemic norms and culture in general, see Campbell (1974); Popper (1972); Toulmin (1967) and Hull (1988).

non-evolutionary causal influences.¹²⁵ Either evolutionary forces would be such that, when it comes to explaining *why we hold the epistemic-normative beliefs we hold*, the influence of those other forces can be assumed to be negligible (Copp; 2008: 197), or the influence of those non-evolutionary forces – however important – will have to be such that taking them into account would not affect the epistemological conclusion of an argument such as DEA.

In other words, the influence of *evolutionary* forces on our epistemic-normative beliefs must have the following feature. However many *non-evolutionary* causal influences also affect our beliefs and however strong those are, if it turned out that *evolutionary* forces were off-track, this fact alone ought to be *sufficient* to undermine justification for our epistemic-normative beliefs themselves.

That evolution explains why we hold the epistemic-normative beliefs we do is believed to be *a plausible empirical claim* by the proponents of the argument we are considering. It is precisely from the plausible empirical nature of this premise – however speculative its details might still be (Gibbard; 1990: 24-30) – that the argument receives its force.

However, it is one thing to say that evolution influenced a core of basic tendencies and dispositions to which certain of the (epistemic) normative beliefs we are inclined to hold can be more or less directly traced¹²⁶ – call this the *weaker evolutionary hypothesis*. Quite another is to say that *a significant portion* of our epistemic-normative beliefs is to be explained in terms of those or other evolved tendencies and dispositions – call this the *stronger evolutionary hypothesis*.

Notice how the degree of empirical support which we ought to assign to this latter, stronger, thesis is significantly less than that which we can plausibly assign to the former. For one, we can expect the indirect evolutionary explanations that the stronger claim must rely on to be most

¹²⁵ I am assuming that those other causal influences cannot themselves be reduced to evolutionary influences of the kind at work in the empirical premise of the DEA and the nature of which I have discussed in the previous sections.

¹²⁶ Such as, for instance, the tendency to care for our offspring, or to reciprocate the help we have received from our conspecifics (Street: 2006), etc. or, for instance, the tendency to form the belief that the next tiger will be carnivorous after having previously observed many carnivorous tigers (Street: 2009).

often quite hard to articulate in detail and to be rather speculative in nature. Hence they will be harder to test and harder to confirm.

Moreover, in the case of epistemic-normative beliefs which we cannot easily trace back to adaptive mechanisms, one must take into account the possibility that they might have *also* been affected by causal forces *other than* evolutionary ones. If we allow this to be the case, then it will be an empirical question what the relative weights of evolutionary and non-evolutionary influences on those beliefs has been. It will similarly be a matter for empirical investigation to determine whether the relation between evolutionary and non-evolutionary forces is such as to warrant the epistemological conclusions of an argument like DEA. We might expect empirical questions such as those to be very hard to settle in many circumstances. In any case, one cannot just *assume* that, when it comes to explaining those beliefs, evolutionary explanations are what we ought to look for ‘by default’, so to speak, or that evolutionary forces are strong and pervasive in the sense I have discussed.

Lastly, the idea that evolutionary explanations can explain as much as the stronger claim requires has been openly questioned by some philosophers (Nagel: 1997, Fodor: 2000, VanFraassen: 1989, White: 2010, Stich: 1990, Churchland: 1987). While I do not have the space to rehearse and evaluate their arguments here, this fact itself testifies to a certain degree of ‘peer disagreement’ regarding the scope and power of evolutionary explanations of our epistemic-normative beliefs.

If what I have said above is correct, the *weaker* evolutionary thesis rests on much more plausible empirical grounds than the *stronger* one. Yet it is something like the stronger thesis that the epistemological Darwinian argument I am discussing requires, if this argument is to be at all compelling.

Discussions of Darwinian arguments in the metaethical literature – it seems to me – often fail to pay adequate attention to the distinction I have just drawn attention to. And the dialectic surrounding whether these arguments are successful typically centres around disagreement over

what follows from the claim that evolution can explain why we hold the normative beliefs we hold.¹²⁷

My strategy of reply to the Darwinian-epistemological argument on behalf of the realist (see Chapter 7) will follow the dominant tendency in the literature and *grant* the truth of the empirical-evolutionary premise. Notice, however, that the urgency with which the normative realist will have to address – and possibly reply to – Darwinian arguments like the one I am discussing is directly proportional to the degree of confirmation that their empirical premise is believed to have.¹²⁸ After all if there were no good reasons to believe that evolutionary influences on our epistemic-normative beliefs are as *strong and pervasive* as the first premise of the argument requires, then there would be no good reasons for the realist to worry too much about the specific argument in question.

To this extent, it is worth keeping in mind that – pending adequate discussion of the issues I have raised above – the evolutionary thesis that receives the strongest empirical support from the ‘least controversial and most well-founded evolutionary speculations possible’ (Street; 2006: 112) is one that would appear to fall rather short of the evolutionary thesis that figures in the first premise of the Darwinian-epistemological argument.

4. Evolutionary Forces, Epistemic-Normative Beliefs, and True Beliefs

In this section I draw attention to two important constraints that the evolutionary explanations at work in the empirical premise of the DEA must answer to.

¹²⁷ Street – echoing Gibbard (1990) – opens her paper ‘A Darwinian Dilemma for Realist Theories of Value’ (2006) by writing ‘if the [evolutionary] facts are roughly as I speculate, here is what might be said philosophically’ (Street: 2006: 112). Responses to arguments such as the one Street articulates there, typically grant that the evolutionary facts are roughly as she takes them to be, but then they go on and disagree on what follows philosophically from that thesis. See, for instance, Enoch (2010) and Wielenberg (2010).

¹²⁸ The ‘evolutionary explanations’ of our (epistemic) normative beliefs that the first premise of the argument builds on are presented as the result of a ‘novel empirical discovery’ (Lillehammer: 2003: 569). The point I am making is that the status of the first premise has a much more ‘speculative’ character than is often openly acknowledged, ‘novel empirical discoveries’ supporting a much weaker claim than the one contained in the premise itself.

The first is dictated by the specific nature of the relationship we take there to be between epistemic-normative beliefs and many other beliefs we hold. The second is dictated by the kind of conclusion that the evolutionary-epistemological argument itself aims to establish. I will discuss them in turn.

4.1. *Desiderata and Constraints* on Explanations of Epistemic-Normative Beliefs

With regard to the first point, consider the following features which we commonly ascribe to epistemic-normative beliefs.

1. Epistemic-normative beliefs – say that E is a reason to believe that P – are often such that our¹²⁹ having those beliefs and attending to them typically *causes*¹³⁰ us to form¹³¹ *other* beliefs – such as the belief that P. Thus we believe that A is guilty of murder *because* we attend, let us suppose, to the belief that we have overwhelming reasons to believe that this is the case. Similarly, in trying to decide between which of a series of competing theories 'T', 'T'', 'T'''', etc. about Q we ought to believe, we consider what reasons we have in favour of believing each. Once we settle on the theory we (believe we) have strongest/decisive reasons for believing – say 'T' – we thereby believe such theory.
2. In normal conditions and when all the relevant epistemic-normative beliefs have been taken into account, we *also* take the following to be the case: if the epistemic-normative beliefs which guide us in forming those other beliefs (and cause us to have them) are true, then those latter beliefs (that A is guilty of murder or that 'T', etc.) are more likely to be true than if we had

¹²⁹ I am taking the 'we' (and the 'our') here to refer to what we might call 'standard epistemically rational agents'.

¹³⁰ I am simplifying here. What causes us to have other beliefs – say P – is (attending to) the belief that there are overall reasons for believing that P, and that these reasons are strong/decisive reasons in favour of holding the beliefs in question.

¹³¹ Or perhaps to 'retain' beliefs which we already had, and were wondering about or thinking about revising, etc. Clearly the causal link in question is also more complex in the sense that beliefs about reasons can cause us to stop believing what we had believed up to a certain point, or to suspend belief in certain propositions, etc. Nothing of what I say in what follows will depend on this more nuanced understanding of the causal relation between epistemic-normative beliefs and other beliefs. I am leaving this more nuanced understanding aside for the sake of simplicity.

formed those beliefs in a way which was not responsive to such reasons. In other words, in normal circumstances, figuring out *what reasons we have to believe what* and then believing it is a good way of coming to form *true* beliefs.

3. Lastly, in a large enough number of cases – with regard to a large portion of our *commonsense beliefs* about our most immediate environment and with regard to most of our *scientific beliefs*, for instance – we do take ourselves to have genuine (epistemic) reasons for holding many beliefs of those kinds. That is to say, we take our beliefs about our reasons to hold specific beliefs of those kinds to be *true*. And – as a consequence – we believe that, as far as most of our commonsense and scientific beliefs are concerned, things are pretty much how those beliefs describe them.

When philosophers claim that we can explain why we hold the epistemic-normative beliefs we hold, the facts I have listed above set *constraints* on those explanations. They specify *desiderata* for those explanations by specifying some important features of the phenomena that ought to be explained. In particular, it would appear to be a *desideratum* on those explanations that they also explain (or at least be compatible with what we take to be) the fact that many of the beliefs which our epistemic-normative beliefs cause us to have are true. Consider again many of the commonsense and scientific beliefs we hold about the world. We often form instances of beliefs of this kind *because we believe we have reasons for* them. And we believe that scientific or commonsense beliefs which we come to hold in this way are largely true.

Notice that not all possible explanations of why we hold the epistemic-normative beliefs we hold will need to be such as to fulfil this *desideratum*. Suppose we had good empirical grounds (imagine them to be rooted in ‘some recent empirical findings of sociology’) to think that our belief *that we have very good reasons to believe that humanity evolved by natural selection* could be explained entirely by the social status which the truth of the evolutionary claim would give to scientists (Gibbard; 1990: 124). Scientists dimly realized that if Darwinian explanations were accepted, they would attain the cultural role of priests. They were motivated by this prospect, and they had the

cultural influence to make it so that people would accept the belief that there are strong reasons to believe that humanity evolved by natural selection.¹³² As a result, the belief that humanity evolved by natural selection was also accepted. Suppose further that ‘sociology’ also showed that most of the epistemic-normative beliefs we have about which scientific beliefs we have most reason to hold were also similarly influenced by the ambitions, power-struggles and delusions of scientist, and that, as a result of this influence, we thereby came to hold most of the scientific beliefs we have. Assume further, as it is plausible to think, that the notion of truth we take to be the correct one (whatever it is) does not equate truth with whatever scientists say we have reasons to believe/would have us believe.

If the empirical fact regarding the *sociological explanation* of our epistemic-normative beliefs were brought to our attention, we could no longer go on believing that most of our scientific beliefs are true. The explanation of why we hold the epistemic-normative beliefs we hold would not be compatible with believing that the beliefs which we form because of them are true.

The evolutionary explanations for our epistemic-normative beliefs which the empirical premise of the DEA builds on are not of this kind. The adaptive mechanisms which – according to these explanations – ultimately give origin to our epistemic normative beliefs are pairing of circumstances and responses that are also such as to issue in us holding true beliefs for the most part. This will become clearer in Chapter 7, when I discuss the details of Street’s ‘adaptive-link’ evolutionary explanation of why we hold the epistemic-normative beliefs we hold.

4.2. ‘Selectivity’ of Darwinian-Epistemological Arguments

Being compatible with/explaining the causal link between epistemic-normative beliefs and true beliefs which we take to be at work in a large number of circumstances is more than a general *desideratum* on the evolutionary explanation that the first premise of the argument relies on. In the

¹³² I am here slightly rephrasing something Gibbard says in Gibbard (1990: 124).

specific context of a Darwinian-epistemological argument like the one I am discussing, an evolutionary explanation which fulfils this desideratum is *required* if the argument itself is to work. For the argument receives much of its power from being *selective*. It works insofar as it *only* calls into question the justification for a given class of target beliefs in particular. In the case at hand these are epistemic-normative beliefs.¹³³

The argument claims that evolutionary explanations of our epistemic-normative beliefs ought to make us sceptical about us ever being justified in holding those beliefs – on the assumption that epistemic realism is true. But an argument of this form rests upon the assumption that *not all kinds of beliefs* we hold will share in the fate of our epistemic-normative beliefs, and that the justification for many of the beliefs we hold – such as mathematical and scientific beliefs and ‘beliefs about mid-sized objects in our most immediate environment’ (Street; 2008: 217) – is not similarly called into question by evolutionary-epistemological arguments of the kind I am discussing here. To this extent, it is true of the Darwinian-epistemological argument what is also true of structurally similar Darwinian challenges against other forms of normative realism. In presenting his Darwinian argument against moral realism, Richard Joyce writes:

The dialectic within which I am working assumes that if an argument that moral beliefs are unjustified or false would by the same logic show that believing that $1+1=2$ is unjustified or false, this would count as a *reductio ad absurdum*. (Joyce; 2007: 182, n.5)

And both Gibbard and Street operate within a framework according to which, when it comes to beliefs about commonsense and about ‘core parts of modern science’ (Gibbard 2003: 263), for instance, acknowledging the influence of evolutionary forces on the faculties which produce such beliefs *does not* call our justification for those beliefs themselves into question (Gibbard; 2003: 263, Street; 2008: 216-7).

¹³³ Or moral/practical normative beliefs (Street: 2006), or normative beliefs more generally (Gibbard: 1990). Arguments of the Darwinian-epistemological kind can take beliefs other than normative beliefs as targets such as, for instance, religious beliefs (Griffiths and Wilkins: 2010). Even in those cases, arguments of this form work by being *selective* in the sense I have specified.

Indeed, according to Street, an evolutionary story can be provided which is such as to *vindicate* the reliability of such beliefs. Here is what she writes with regards to beliefs about midsized objects in our most immediate environment:

Midsized objects in our immediate environment are the kinds of things one can run into, be injured by, eat, and be eaten by. Other things being equal, then, creatures with an ability accurately to detect midsized objects in their immediate environment tended to survive and reproduce in greater numbers than creatures who lacked this ability. I am a product of this evolutionary process. This gives me reason to think my capacity to make judgments about midsized objects in my immediate environment is reliable. (Street: 2008: 217)

For both Street and Gibbard, epistemic-normative beliefs are thus *quite unlike* most of our commonsense or scientific beliefs. If we assume that epistemic-normative beliefs are true (on a realist construal of their truth-conditions), the claim that they are reliable would be ‘in tension’ with the plausible empirical hypothesis that evolutionary forces are largely responsible for us holding the epistemic-normative beliefs we hold.

Notice that if the Darwinian-epistemological argument were not selective in this way, it would merely be an instance of a more ‘routine, general sceptical worry’ and, as Street herself puts it ‘would lose much if not all of its interest’ (2008: 216). For if the argument generalised, it would seem to raise the much bigger threat of *global scepticism*, and global scepticism ought to be a worry for everyone – rather than one which is for the *epistemic-normative realist in particular* to address,¹³⁴ or one that calls epistemic-normative realism *in particular* into question.

Let us now go back to the distinction between different kinds of explanations of epistemic-normative beliefs which I have drawn attention to above. This was the distinction between:

(1) Causal explanations of our epistemic-normative beliefs which are *compatible* with us retaining the belief that many of the beliefs they cause us to have are true.

¹³⁴ Conversely, if we believe we ought not to take global sceptical scenarios too seriously, then neither ought we to worry too much about scepticism with regards to our epistemic-normative beliefs.

and

(2) Causal explanations of our epistemic-normative beliefs which are *not compatible* with us retaining the belief that many of the beliefs which they cause us to have are true.

We are now in a better position to see why the evolutionary-explanations of our epistemic-normative beliefs that the Darwinian-epistemological argument must rely on are of the first kind. If they were not, and if this fact was brought to our attention, then we ought to question the truth of the beliefs which we form because we believe we have reasons for them. If I am told that my belief *that the fossil record is a reason to believe that evolution is true* is the result of an evil demon interfering with my brain,¹³⁵ I ought not only to doubt whether I do have indeed a reason to believe in evolutionary theory. I will have to doubt that evolutionary theory itself is true, on the assumption that I have formed this belief – or that I retain it – *because I thought I had a reason for it*.

As a result, if the Darwinian epistemological-argument relied on evolutionary explanations which are of the second kind above, it would lose its selective nature. Such explanations for our epistemic-normative beliefs would call into question our justification for the large number of beliefs the epistemic-normative beliefs we hold cause us to have (or to retain, etc.). Again, if we count among those latter beliefs many of the commonsense and scientific beliefs we hold, the upshot of an evolutionary claim of this kind would come close to *global scepticism*.

Thus that evolutionary explanations for our epistemic-normative beliefs provide explanations of the kind at (1) is not only a *desideratum* on such explanations. Explanations of this kind are *required* if an evolutionary-epistemological argument like the one I am discussing is to retain its selective nature and do the work it is supposed to do. That is, if it is to lead to scepticism about our justification for epistemic-normative beliefs *only*; or be part of an argument which

¹³⁵ Assume that the content of the epistemic-normative beliefs that the demon implants in me is generated at random.

targets epistemic-normative realism *in particular*, rather than merely inducing – if anything – a more ‘routine, general sceptical worry.’¹³⁶

Being aware of this central feature of the evolutionary explanations at work in DEA is very important. For not only will it affect how we must understand the claim that evolution is an *off-track* causal process with respect to epistemic-normative facts¹³⁷ (see next section). It will also, or so I will argue in Chapter 7, ultimately point the way towards successful realist strategies of reply to the argument itself.

4.3. How these *Desiderata* Affect the Use of DEA as an Argument against Epistemic Realism

It will not have escaped notice how the considerations I have just reviewed above are not only such as to set constraints for the evolutionary explanations at work in DEA. They are also such as to set constraints, respectively, for meta-normative theories of epistemology more generally; as well as for the dialectic which sees an argument such as DEA employed as part of an argument *against epistemic realism* rather than to argue directly, so to speak, *for a sceptical conclusion*.¹³⁸

Clearly proponents of anti-realist views of epistemic-normativity will also have to ‘make good’ on the features of our beliefs about epistemic-reasons which I have listed above. For instance, even an *error-theorist* about epistemic-reasons claims – i.e. someone who believes that no beliefs about epistemic reasons are true – will be under some pressure to provide an account of the fact that we hold many of our beliefs about epistemic reasons to be *true*.

¹³⁶ Notice how the situation with epistemic-normative beliefs here differs significantly from that of normative beliefs of different form – such as *practical-normative* or *moral* beliefs. Moral and practical-normative beliefs do not bear the causal links to beliefs of other kinds (such as commonsense or scientific beliefs) which epistemic-normative beliefs bear. They bear similar causal links to *actions/intentions* instead. Because of this, the explanations of why we hold the moral beliefs we hold are not ‘constrained’ in the same way in which those for our epistemic-normative beliefs are. And they are, in general, quite removed from the possibility of generating general sceptical worries of the kind I have discussed in connection with the explanations of epistemic-normative beliefs themselves.

¹³⁷ It will have to be an explanation which is such as to be (1) *off-track* with respect to epistemic-normative facts but (2) *on-track* with respect to (many of) the beliefs which those epistemic-normative beliefs cause us to have.

¹³⁸ That is to say, the conclusion that our epistemic-normative beliefs are not justified.

One route to an error-theory of the kind at issue – we can imagine – might go via believing that what is by far the most plausible interpretation of our epistemic-normative beliefs entails that these beliefs are about irreducible normative facts; and that we have strong reasons to believe that facts of this kind are just ‘too queer’ to exist. An error theorist of this kind will still have to explain why although epistemic-normative beliefs are – strictly speaking – systematically false, we normally refer to at least some of them as *being true*. He will also have to explain why those of them which we call *true* are such that they cause us to form beliefs that, on most occasions, are more likely to be true than beliefs formed in many other ways.

Analogous considerations will apply with regard to philosophers who subscribe to *non-cognitivist* and *expressivist* analyses of epistemic-reasons claims, or to any other view that counts as *anti-realist* for that matter. In other words, in order to be able to go on using the language of epistemic-reasons (and, plausibly, epistemic-justification), anti-realist theorists of normativity will have to show that their distinctive understanding of the ‘phenomena’ of epistemic-normativity is a viable alternative with respect to epistemic realist explanations of the same phenomena.

A similar, and related point, goes with regard to the use of DEA as a dialectical tool used in arguments *against epistemic realism*. In order for DEA to be used as part of an argument against realism, DEA will have to be *selective* not only in the sense that it calls into question the justification for our epistemic-normative beliefs *only*. It will have to be selective in the sense that it calls into question our justification for our epistemic-normative beliefs *only* – or in particular – *on a realist understanding of what the truth of these beliefs requires*.

I have already discussed the idea behind this use of the argument. This is that the justification for our epistemic-normative beliefs which DEA threatens can be *restored* by shifting to a different (anti-realist) meta-normative understanding of beliefs about epistemic-reasons themselves,¹³⁹ or

¹³⁹ This seems to be the idea at work in Street (2009) and it derives from an analogous application of Darwinian-epistemological argument to challenge normative realism in the practical domain (Street: 2006).

at least¹⁴⁰ that we can account for all the important phenomena of epistemic-normativity in a way which does not lead to troublesome forms of scepticism.¹⁴¹

If it turned out to be the case that anti-realist views succeed where realist views fail, then this fact would provide an argument *against realism* by contextually providing an argument *in favour of anti-realist views* themselves.¹⁴² According to this line of reasoning those views can – while realism cannot – avoid the unwanted sceptical conclusions that DEA would lead to.

Clearly, showing how anti-realist views of epistemic-normativity can do better than realist views here – and perhaps which anti-realist view does so in particular – is far from an easy task. Yet much of the success of DEA as an argument against realism will depend on the success of those alternatives.

I will not discuss the matter of whether – and how – anti-realist metanormative theories in epistemology can do better than realist ones at coping with the kind of challenge the DEA raises. In Chapter 7, I will argue that epistemic-realists of various kinds have good ways to resist the conclusion of DEA. As a result, I will claim, on the assumption that realism is true, we do not end up with the sceptical result DEA aims to establish.

Yet the dialectical point I have just brought to view is worth keeping in mind. For even if DEA were successful, it will require serious and substantive philosophical work to show how anti-realist views can do better where realist views fail. Pending this work, we might well be stuck with the troublesome forms of scepticism that the success of DEA would bring about.

¹⁴⁰ On some views which, intuitively, ought to count as anti-realist – such as error-theories or non-cognitivist/expressivist views – we cannot speak in any straightforward sense of the *justification* for our epistemic-normative beliefs, at least on any straightforward *epistemic* sense of justification itself. On *error-theoretic* views, this is because epistemic-normative beliefs are taken to be *systematically false*. And it would be odd (to say the least) to believe that (1) a category of beliefs is systematically false and (2) at least some of the beliefs which belong to this category are such that, after all, we are justified in holding them. On standard *non-cognitivist/expressivist* views we cannot straightforwardly speak of the justification for our epistemic-normative beliefs, because epistemic-reasons judgments are not even taken to express beliefs in the first place.

¹⁴¹ As I have already discussed, the troublesome forms of scepticism are not only scepticism about our epistemic-normative beliefs, but also scepticism about all the (other) beliefs which we form (maintain, revise, etc.) because of the fact that we hold epistemic-normative beliefs of certain kinds.

¹⁴² Or in favour of a particular anti-realist view if, for instance, the view in question would be clearly and by far the best way to improve on the failures of realist views here.

It is now time to turn to discuss another crucial thesis at work in DEA: the claim that if evolution explains epistemic-normative beliefs, then those beliefs are not correlated to epistemic-normative facts (premise (4)). I will be occupied with discussing this thesis in the next chapter.

Chapter 6

Evolution and the Lack of Correlation between Epistemic-Normative Beliefs and Epistemic-Normative Facts

1. Introduction

How must we understand the claim that evolution is an *off-track causal influence* with respect to epistemic-normative facts of the kind the realist believes in? With respect to normative facts *of what kind*, if any, would evolution be an off-track causal influence? And what reasons might there be for believing that evolution is one such process? In what follows, I turn to address those issues.

2. Epistemic-Realism and the Darwinian-Epistemological Argument

Let us begin with the second question. Discussion of this point will at the same time serve to elucidate the first metaethical premise of DEA (Premise (2)). It will also help us be clearer about the target of Gibbard and Street's Darwinian anti-realist strategies; that is to say, about the kind of meta-normative view which they believe to be in tension with the facts of evolution.

So far in this chapter, I have been using the expressions 'epistemic realism' or 'normative realism in epistemology' to refer to this view. However – as I discussed at some length in Chapter 3 – these labels are far from being of straightforward application. We saw how this was in part due to the fact that the term 'realism' is used in importantly different ways in the context of meta-normative discussion. We also saw how, in the specific case of epistemic normativity, the situation is made all the more complex by the fact that expressions of the kind 'E is a reason for believing that P' (when they are taken to be about *epistemic* reasons) lend themselves to very

different conceptual interpretations, and that these interpretations, in turn, might affect *what kind of facts* the realist believes will make at least some of those expressions true.

It is thus somewhat unsurprising to find that Gibbard and Street operate with different ‘working definitions’ of normative realism (and normative anti-realism) about epistemic normativity. This is so despite the fact that each of them takes ‘normative realism in epistemology’ to be (or fall within) the target of the structurally similar Darwinian-epistemological strategies they formulate.¹⁴³ The view Street calls ‘realism about epistemic reasons’ (Street; 2009) is the view that:

Q: At least some epistemic-normative facts hold *independently* of the set of an agent’s pro-attitudes, such that there can be an epistemic-normative reason for an agent to believe that P even though the conclusion that there is this reason in no way follows, as a logical or instrumental matter, from the agent’s own set of pro-attitudes in combination with non-normative facts (Street; 2009: 214-216).

Street understands epistemic anti-realism to amount to the denial of this thesis. The specific anti-realist view she favours is the view that:

S: Epistemic-normative facts are *dependent* on the set of an agent’s *contingent* set of pro-attitudes, in such a way that there can be an epistemic-normative reason for an agent to believe that P *only if* the conclusion that there is this reason follows, as a logical or instrumental matter, from the agent’s own set of contingent pro-attitudes in combination with non-normative facts (Street; 2009: 214).

¹⁴³ These are the strategies which I have already discussed. They rely on arguments like DEA to show that epistemic-realism leads to scepticism about our epistemic-normative beliefs; and which then go on to argue for the fact that anti-realist views can avoid this troublesome result and that this is a reason to prefer those views over realism.

On the other hand, the view Gibbard calls ‘normative realism’ in general (a view which we must take to include realism about epistemic-normativity) is the view that:

Q’: There are (epistemic-) normative facts.

Notice that, on Gibbard’s definition of (epistemic-) normative realism, (epistemic-) normative realism is compatible with believing that (epistemic-) normative facts are *attitude-dependent facts* of some kind or other. The specific anti-realist view of normativity Gibbard favours – what he calls ‘norm-expressivism’ – is the view that:

G: To call something rational is to express one’s acceptance of a system of norms that permits it.
(1990: 7)

In the specific case of epistemic reasons this view comes down to the claim that:

G’: To call something – R – a reason for believing X is to express one’s acceptance of a system of norms that require treating R as weighing in favour of believing X. (1990: 163)¹⁴⁴

According to Gibbard, ‘acceptance’ is a non-cognitive state of mind. Epistemic-normative sentences do not (ultimately) express beliefs, they do not (ultimately) purport to describe anything and – if there is a sense in which we can correctly call some of these sentences true – this is not to be understood in terms of these sentences corresponding to *normative facts* of any kind.

Gibbard goes to some length in order to distinguish his ‘norm-expressivist’ position from what he calls ‘subjectivist’ views of normativity. These are views which say that:

¹⁴⁴ Gibbard’s original definition is about *reasons for action*. I have adapted his analysis to fit *reasons for belief*.

When a person calls something rational, he is reporting a state of mind: he is saying that someone – be it he, or his group, or perhaps the agent or the agent’s group – is in some particular state of mind with respect to that thing. (Gibbard; 1990: 203)

And he believes that acknowledging the influence of evolutionary forces on our normative judgments ultimately calls into question the idea that there are normative facts *of either the objective or the subjective kind* (1990: 111-117).

It seems to me that Street’s own account of epistemic normativity – what she calls *Humean Constructivism* – ought to count as a kind of *realist* and *subjectivist* view, if we were to go by Gibbard’s classifications. Notice that if this is correct, and if the kind of Darwinian-epistemological strategy that Gibbard deploys turned out to be successful against its intended target, Street’s own account of epistemic normativity would *also* be called into question by it.¹⁴⁵ This fact is in itself interesting, for it reveals a general problem with the kind of Darwinian-epistemological strategies that both Gibbard and Street adopt,¹⁴⁶ at least once we compare those strategies and keep in mind the fact that they are targeted at different meta-normative views.

A more detailed discussion of this point would require a deeper engagement with Street’s and Gibbard’s positive views than the one I can afford in the present context. One way to make progress with respect to the question I am addressing here, however, would be to focus our discussion on the epistemic-realist views, if there are any, which would figure in a Darwinian-epistemological argument that both Street and Gibbard would endorse.

¹⁴⁵ Gibbard also argues against a Humean understanding of reasons – along the lines of that which Street endorses – on independent grounds. See Gibbard (1990: 160-4).

¹⁴⁶ Both Gibbard and Street take themselves to be providing arguments which rely on the ‘most recent and uncontroversial empirical findings of evolutionary psychology’. On a plausible reading of both Gibbard and Street these arguments are of the Darwinian-epistemological kind I am discussing. However, Gibbard employs his Darwinian-epistemological argument as part of a broader argument *against* a family of meta-normative views which include the kind of *factualist* and *subjectivist* account which Street favours. On the other hand, Street employs her argument as part of an argumentative strategy which is, ultimately, *in favour of* such a *factualist* and *subjectivist* view. In particular, Street believes – and Gibbard would deny – that her own factualist and subjectivist account of epistemic normativity can answer the Darwinian-epistemological challenge that she levels against the views of epistemic-normativity that she calls ‘realist’. Clearly, they cannot both be right. The fact I have just drawn attention to opens up an obvious line of criticism of the Darwinian-epistemological strategies against realism which both Gibbard and Street rely on. However, I will not pursue this matter any further here.

2.1. The Target of the Darwinian-Epistemological Argument

In order to do so, I suggest that we focus the rest of our discussion on a set of meta-normative views in epistemology which would count as epistemic-realist by both Gibbard's and Street's standards. These views would also count as (kinds of) epistemic-realist views by the definition of epistemic realism which I have myself adopted. These are views according to which:

1. There are epistemic-normative facts.
2. Epistemic-normative facts are *objective*.

From what I have said so far, it is easy to see why views of this kind would fall within the target of both Gibbard's and Street's considerations. Gibbard argues against the view that *there are normative facts*, Street against the view that *normative facts are objective* – in the sense of '*independent from the set of the contingent pro-attitudes of agents*'. Thus both Gibbard and Street would agree in thinking that views of the kind I have just singled out would fall within the target of the kind of evolutionary strategies they deploy.

Notice how the family of epistemic-realist views of this specific kind is still quite broad. In order to see this we must go back to the taxonomy of epistemic-realist views which was the result of my discussion in Chapter 3. That taxonomy makes it clear that subscribing to the thesis that there are *objective normative facts* is compatible with subscribing to very diverse accounts of what I have called the 'structure' of facts about epistemic reasons. Hence it is compatible with epistemic-normative views which are quite different from one another.

One such view is a version of epistemic realism which combines the thesis that there are epistemic-normative facts with a *non-Humean instrumentalist* analysis of epistemic reasons. On this view, epistemic-normative facts would be identical with complex facts which can be analysed into evidence facts (that E is evidence for P), means/ends facts (if E is evidence for P then believing

that P is instrumental to acquiring beliefs which are Q), and further objective normative facts¹⁴⁷ (beliefs which are Q are beliefs which we have reasons for acquiring, where these reasons are independent of the contingent pro-attitudes of agents). Call a realist view of this kind ‘Non-Humean instrumentalist epistemic realism’.

Another such view is a version of epistemic realism which combines the thesis that there are epistemic-normative facts with an *Evidentialist* analysis of epistemic reasons and takes that analysis *at face value*. That is, it takes the conceptual structure of epistemic-reason claims thus understood to be reflected by the structure of epistemic-reason facts. On this view, epistemic-reason facts would strongly supervene on – but be distinct from – evidence facts. Call a realist view of this kind ‘evidentialist epistemic realism’.

Yet another epistemic-realist view which is also *objectivist* is the view that epistemic-normative facts are identical with evidence-facts. As I have already pointed out (Chapter 3), an epistemic realist view of this kind can be combined with both *evidentialist* and *evidence-as-reasons* analyses of epistemic-reasons. For the sake of having a convenient label at hand, I suggest that we nonetheless call a realist view of this form ‘evidence-as-reasons epistemic realism’.

As can now see more clearly, the three kinds of epistemic-realism I have listed above are very different from one another. Can we narrow down the focus of our discussion even further? One obvious way to do so would be to look to Gibbard and Street for help with this question. This is not as straightforward a task as one might think. Yet, it is not hopeless. Gibbard does not distinguish between different kinds of epistemic-realism. But nowhere does he seem to think that an instrumentalist analysis of epistemic reasons claims ought to be taken seriously.¹⁴⁸ Street does not discuss the epistemic-realist and objectivist view which I have called Non-Humean instrumentalist and which ought to count as an ‘epistemic realist’ view by her own standards. On

¹⁴⁷ The normativity of these facts could be in turn due to facts regarding the intrinsic value of knowledge/justification/understanding/true belief, etc., for instance. Facts of this latter kind would provide reasons for acquiring beliefs which instantiated those intrinsically valuable properties or states.

¹⁴⁸ On this point, see also Cuneo’s discussion of Gibbard’s views on epistemic normativity (Cuneo; 2007: Chapters 5 and 6).

the other hand, she explicitly discusses epistemic-realist views of the second and third kind I have outlined above (Street: 2009). Given these facts, I suggest we focus the rest of our discussion on the views which I have called ‘evidentialist’ and ‘evidence-as-Reasons’ epistemic-realist views.¹⁴⁹ In the next chapter I will discuss in some detail why Street thinks these views fall prey to an argument such as DEA. I will then argue that each of these versions of epistemic-normative realism can survive this argument. In other words, I will claim that if we assume that either of these forms of epistemic-normative realism is true, we will be able to resist the scepticism that proponents of DEA would seek to impose on us.

Before we are in the position to see why this is the case, however, we must address the other questions I have raised at the beginning of the previous section. What does it mean to say that evolutionary forces are ‘off-track’ with respect to epistemic normative facts construed in the way I have just made clear? And what kinds of considerations do Street and Gibbard rely on when they defend the view that evolutionary forces are of this kind? I will turn to discuss these topics in the next sections.

3. Evolutionary Forces as ‘Off-Track’ Causal Influences

3.1. Off-track: ‘Chanciness’ and ‘Predictability’

We are finally in a position to begin examining the thesis that evolutionary forces are ‘off-track’ causal influences with respect to epistemic-normative facts. In this section, I will look at how the notion of lack of correlation between epistemic-normative beliefs and epistemic-normative facts at work in premise (4) of DEA is best understood. I will take the claim that epistemic-normative belief are *off-track* (with respect to epistemic-normative facts) as equivalent to

¹⁴⁹ There are also independent reasons for focussing discussion of the argument on epistemic-realist views of those kinds. These are provided by the fact that there are good reasons to believe that the ‘instrumentalist’ analysis of epistemic reasons which instrumentalist versions of epistemic-realism depend on is mistaken. Although this is an analysis of epistemic reasons and normativity which appears to be quite widespread in epistemology (Berker: 2013a) it faces rather serious problems (see Kelly: 2003 and 2007, and Berker: 2013a and 2013b) and – it seems to me – the motivation for wanting to hold on to a view of this kind can be (better or equally well) accommodated by the other analyses of epistemic reasons I discuss. I do not have the space to discuss this matter any further here.

the claim that there is a *lack of correlation* between epistemic-normative beliefs and epistemic-normative facts, and I will use these forms of words interchangeably in what follows. And I will also consider a causal influence CI on a set of beliefs C to be *off-track* iff the beliefs which result from the operation of CI – i.e. the beliefs which are *caused* by CI and which CI explains – are not correlated with the sort of facts which make beliefs of this kind true.

When I illustrated – by means of an analogy – the general form of the epistemological challenge which DEA relied on, I used the example of a drug which affected our mathematical or colour beliefs in a way which made them unreliable. Given the workings of the drug, beliefs of the kind the drug targeted could *easily have been* false. Because of this, even in the case in which some of these beliefs were true, they would only be *accidentally* so. Another way to put this – I added – was by saying that given the influence of the drug, we would have formed the C-beliefs we formed, regardless of whether they had been true (‘whether they had been true or not’).

Causal influences on a given set of beliefs can determine the lack of correlation between those beliefs and the sort of facts which make those beliefs true in somewhat different ways. This is reflected in the fact that when we believe that the causal origin of a set of beliefs undermines the reliability of those beliefs, there may be different ways of motivating our sceptical worries ensuing from appreciating this fact (White; 2010: 579-80).¹⁵⁰

One such way will emphasise the *contingency* or *chanciness* of our beliefs in a given situation. The other will emphasise the *unavoidability* or *predictability* of us believing as we do in the circumstances envisaged.

Suppose we know that the drug affects the content of our colour-beliefs *in a random way*. Then the thought which motivates our sceptical worries is the thought that we could have easily formed beliefs which are different from the ones we have. In this case, whatever the colour-facts

¹⁵⁰ The terminology I adopt below (‘chanciness’ vs. ‘predictability’, etc.) is from White (2010), although my use of these notions and my understanding of how they are related is slightly different from his.

had been, it would be a matter of luck if colour beliefs matched the facts which would make them true.

Suppose, on the other hand, that the drug affects our colour-beliefs *in a systematic way*. We know that, under the influence of the drug, we are bound to perceive all objects as red, perhaps because we can provide an explanation (e.g. a ‘chemical one’) for why this is the case. In this case, whatever the colour-facts had been, we were bound to believe as we do. And, provided that our environment is not such as to consist mostly or exclusively of red objects, it would be a matter of luck if the colour beliefs we hold matched the facts which would make them true.

Each of the scenarios above is compatible with our describing it by means of expressions such as ‘we could easily have been mistaken’ about our colour-beliefs or ‘we would have had the colour-beliefs we had whether or not they had been true’. Yet the explanation for why it is appropriate to use those expressions would be rather different in each of the two cases.

3.2. ‘Off-track’ Practical-Normative Beliefs

The distinction I have just drawn attention to is important. For, it seems to me, it reveals a significant asymmetry between the way evolutionary-epistemological arguments can be – and have been – put to work in the epistemic and practical domain respectively. Both kinds of intuitions – chanciness and predictability – can in principle be appealed to in the context of a Darwinian-epistemological argument which targeted realism about practical normativity. The case is different when it comes to epistemic normativity. Here, I will suggest, it is ‘predictability’ considerations alone which, ultimately, would have to carry the weight of the argument.

When Street argues against normative realism in the practical domain (Street; 2006), she appears to draw on *both* of what I have called *chanciness* and *predictability* considerations, in turn, in order to articulate her sceptical challenge.¹⁵¹

¹⁵¹ See also White (2010: 585-95) and Clarke-Doane (2012: 318-20) for a similar reading of Street.

Drawing on the former, she writes:

[On a realist theory of practical normativity] allowing our evaluative judgments to be shaped by evolutionary influences is analogous to setting out for Bermuda and letting the course of your boat be determined by the wind and tides: just as the push of the wind and tides on your boat has nothing to do with where you want to go, so the historical push of natural selection on the content of our [practical normative] judgments has nothing to do with evaluative truths. (2006: 117)

And again:

Imagine we had evolved more along the lines of our close primate relatives the bonobos, so that we experienced sexual relations with all kinds of different partners as ‘called for’ in all kinds of different circumstances [...or...] on the model of the social insects [...]. Presumably in these and other such cases our system of full-fledged, reflective [practical-normative] judgments would have looked very different as well, and in ways that loosely reflected the basic evaluative tendencies in question. (2006: 310-20)

The context of the second passage is one in which Street illustrates the sense in which – as she puts it – the ‘content of human evaluative judgments has been tremendously influenced [...] by the forces of natural selection’. Yet it is perhaps quite natural to read the passage as leading to a sceptical worry animated by the kind of ‘chanciness’ considerations I am illustrating here. For one way to read Street’s words in this passage echoes Darwin’s own thought that had our lineage evolved to have a different social system, we would have come to have different moral beliefs. As he puts it ‘had men [been] reared under precisely the same conditions as hive-bees’ we would have come to believe that it is permissible for ‘unmarried females [...] to kill their brothers’ (Darwin; 1871: 70).

The sceptical worry originates as follows. Grant that we could (easily) have evolved differently, and in such a way as to display the kind of evaluative attitudes that bonobos and hive-bees display. Grant also that the kind of evaluative attitudes we would have had in this case, would have affected the content of our normative beliefs. If that had been the case, the content of our normative beliefs would have been different from what it is now, perhaps to the extent that we would have believed ‘fratricide’ permissible. Suppose – as the realist on Street’s

understanding of realism believes – that normative facts *do not depend on evaluative attitudes*, that if our evaluative attitudes had been different, *the normative facts would have been the same*. Then, given the influence of evolution on our evaluative attitudes and (consequently), our normative beliefs, we ought to be sceptical about us justifiably retaining the normative beliefs we do. Our evolutionary history could have been different. Hence, our evaluative attitudes – and our normative beliefs – could have been different. Had this been the case, the normative facts would have still been what they are. Or, at least, they would not have been – as Street puts it – ‘correspondingly different’, since they are independent from those attitudes.

On this reading, the claim that normative beliefs are off-track amounts to the truth of the following counterfactual:

CF: had our normative beliefs been different, the normative facts would not have been correspondingly different.

together with the claim that:

D: our normative beliefs could (easily) have been different.

On a realist understanding of the truth conditions of our normative beliefs, then, it would be a rather surprising fact if our normative beliefs were true. Again, it would be like setting out for Bermuda and letting the course of your journey be affected by whatever winds happen to blow along your way. Call this the ‘chanciness’ reading of the lack of correlation between normative beliefs and normative facts.

But in articulating her sceptical challenge against normative realism in the practical domain, Street can also often be read as relying on considerations of what I have called the ‘unavoidability’ kind. In her ‘Reply to Copp’ (2008) she writes:

As a purely conceptual matter¹⁵² [...] *normative truths might be anything*. In other words, survival might be bad, our children's lives might be worthless, and the fact that someone has helped us might be a reason to hurt that person in return. [...]. Noting the sense in which *the normative truths might be anything*, and noting the role of evolutionary forces in shaping the content of our basic evaluative tendencies, we may wonder whether there is any reason to think these forces would have led us to be capable of grasping the independent normative truths posited by the realist. (2008: 208, italics mine)

Notice how the emphasis on the fact that 'normative truths might be anything' in the quote above is compatible with – and perhaps suggests – the following 'predictability' interpretation of the claim that normative beliefs are off-track.

Suppose that we *could not* (easily) have evolved in such a way as to have (very) different normative beliefs from the ones we have. Indeed, given the evolutionary pressures on our ancestors and their forebears, we believe that it is quite unsurprising that we have come to hold the normative beliefs we do hold. In this case, the claim that there is a 'lack of correlation' between normative beliefs and normative facts will have to be supported by considerations of a different kind.

One way to do this is by defending the belief that the following counterfactual is true:

CF?: had the normative truths been different, our normative beliefs would have been the same

together with the belief that – to paraphrase Street – 'the normative facts might be anything'. In other words:

D?: the normative facts could (easily) have been different.

¹⁵² Street's reliance on conceptual (rather than metaphysical) possibility is an important element of a Darwinian-epistemological argumentative strategy against normative realism which relies on unavoidability considerations. I will return to elaborate on this point in the next section.

The thought behind what is expressed by these counterfactuals can be put informally by saying that we were *bound* to believe as we do – regardless of what the facts had been like. Call this the ‘predictability reading’ of the lack of correlation between normative-beliefs and normative facts.

It is possible to disagree on how promising, when thinking about normative-practical beliefs, each of these ways of framing the sceptical challenge is, the first (the ‘chanciness reading’) relying as it does on a counterfactual scenario in which the normative beliefs we hold are different and the second (the ‘predictability’ reading), on one in which it is the normative facts which are different. Neither need these strategies be thought of as being *exclusive* of one another or *exhaustive* of the spectrum of sceptical possibilities.¹⁵³ I do not have the space to elaborate on these matters any further here. The point of the preceding discussion was merely to bring to view two rather different interpretations of the lack of correlation thesis at work in Darwinian-epistemological style of arguments. In what follows, I turn to look at how a thesis of this kind is best understood when it is put to work in an argument like DEA.

3.3. ‘Off-track’ Epistemic-Normative Beliefs

Whatever one thinks about the practical-normative case, there are strong reasons for believing that a Darwinian-epistemological argument like DEA ought rather to rely on considerations of what I have called the ‘predictability’ kind. This can be seen to follow from my discussion of the kind of constraints which an evolutionary explanation of epistemic-normative beliefs must answer to. In that discussion, I defended the claim that the evolutionary explanations of our epistemic-normative beliefs which the Darwinian-epistemological argument relies on must be *compatible* with our retaining the belief that many of the beliefs which epistemic-normative beliefs cause us to have are true and that we are justified in holding them.

¹⁵³ The ‘chanciness’ and ‘predictability’ considerations need not be exclusive: one might think that it is a *purely contingent* matter that evolution took a certain course, but that once it did, it was *bound* to go in a certain direction. See White (2010: 279-80). They must not be taken to be exhaustive of the kind of considerations one can appeal to in order to argue for the lack of correlation between beliefs and facts of a certain kind either. For the purposes of my discussion here a comprehensive engagement with these other possibilities is not required.

For this to be the case, these explanations ought *not* to be such as to support the counterfactual that our epistemic-normative beliefs *could easily have been very different from what they are*. For if our epistemic- normative beliefs had been very different from what they are, then the beliefs which epistemically rational agents form or retain *because of them*, would *also* have been very different from what they are. Among those latter beliefs, as we saw, are beliefs about mid-sized objects in our most immediate environment and beliefs about core parts of modern science. Call these CS-beliefs. If any of our beliefs are true, it is plausible to think that *a large portion* of our CS-beliefs are true.

Assume now that when CS-beliefs are true, they are made true by corresponding facts. Assume also that those facts are ‘objective’ in the sense that they do not depend¹⁵⁴ on the beliefs (or any other mental states at that) of epistemically rational beings. Then if CS-beliefs had been different, the facts which make them true would not have been correspondingly different. Hence, if those beliefs could easily have been different, even if they were true, they would have been only accidentally true. Appreciation of this fact ought to make us doubt that many of our CS-beliefs are true. And (our awareness of) it ought to lead us to believe that our CS-beliefs are unjustified. If this is the case, they do not amount to knowledge. And if a large portion of those beliefs does not amount to knowledge, it is hard to see what does.

If what I have said so far is correct, the assumption that the epistemic-normative beliefs we hold could easily have been different entails something like the worrying form of scepticism which – I have argued – Darwinian-epistemological arguments ought to avoid. In order to avoid this result, proponents of the Darwinian-epistemological argument will have to understand ‘off-track’ and ‘lack of correlation’ in the epistemological premise of that argument along the lines of

¹⁵⁴ Dependence/Independence of the kind at issue, of course, will have to be qualified. Commonsense facts about my mental states – such as the fact that I am in pain – will clearly depend on what these mental-states are like. The sense of dependence/independence I take to be at work here would count as trivial or irrelevant cases such as the one I have just pointed out. How the relevance and non-triviality at issue are to be specified in detail is a difficult and controversial question. I cannot address this issue here. Whether a certain category of facts is mind-independent *in the relevant sense* will often be clear in practice.

what I have called the ‘predictability’ interpretation of this claim. That premise will then amount to the following counterfactual:

CF³-E: Had the epistemic-normative facts been different, our epistemic- normative beliefs would have been the same.

together with the claim that:

D³-E: the epistemic-normative facts might easily have been different.

Notice how believing that CF³-E is true does not require us to believe that we could easily have had different epistemic-normative beliefs from the ones we have. Thus this interpretation of the epistemological premise of the argument avoids the problem I have discussed above. Indeed, counterfactual claims of this form are quite often put forward in contexts in which we *expect* or *predict* that someone will form beliefs of a given kind, when we also think that forming beliefs in this to-be-expected way is ‘insensitive’ to the kind of facts which make those beliefs true.

One crucial issue to discuss with regards to this interpretation has to do with how we must read the counterfactual claim at CF-E’ and, hence D-E’ itself. We must ask, particular, what kind of modality we must take to be at work in those claims. The general worry here can be put as follows. Epistemic-normative realists (as we saw) typically claim that truths that link epistemic-reason facts to evidence-facts are metaphysically necessary.¹⁵⁵ If so, then it is not metaphysically possible for the epistemic-reasons truths to be different, while the evidence-truths remain unchanged. But if the evidence-truths are not held fixed, how can one be confident that had the

¹⁵⁵ This is true of both epistemic-normative realism of the ‘evidentialist’ kind (according to which epistemic-reason facts *strongly supervene on* evidence-facts) and of epistemic-normative realism of the ‘reasons-as-evidence’ kind (according to which epistemic-reason facts are *identical with* evidence-facts).

epistemic-normative truths been very different, our epistemic-normative beliefs would have been the same?¹⁵⁶

We can illustrate this point as follows. Take epistemic realists of what I have called the ‘evidentialist’ kind. Philosophers who hold this view believe that facts about epistemic-reasons are distinct from, but strongly supervenient on, facts about evidence. They believe that, for instance, the fact that E is evidence for P necessarily entails that there is an epistemic-reason to believe that P. Given this account of the relation between normative-facts and evidence facts then, if the epistemic-reason facts were different, the evidence-facts would also be different. As I have argued in Chapter 2, it is quite safe to assume that what beliefs about epistemic-reasons we have depends on facts about evidence. If this is the case, then had facts about evidence been different, our epistemic-normative beliefs would *themselves* have been different. As a consequence, the counterfactual at CF’-E would be false. If the kind of modality invoked by this counterfactual is *metaphysical* modality, then it would not be possible for the epistemic-normative truths to be different while the epistemic-normative beliefs remained the same.¹⁵⁷ What this suggests is a reading of the possibility invoked by the counterfactual in question along the lines of *conceptual* possibility (i.e. conceivability) – rather than *metaphysical* possibility.

If we read CF’-E in this way, what it claims is that – *for all we can intelligibly imagine* – had the epistemic-normative facts been very different, our epistemic-normative beliefs would have been the same.¹⁵⁸ Clarke-Doane (2012) illustrates how counterfactuals that rely on conceptual possibility can have epistemic significance with the following example:

Imagine a young child guessing the truth of recondite mathematical conjectures at random. Imagine, moreover, that her guesses turn out to be systematically correct. There is an obvious sense in which the forces generating the child’s beliefs in this scenario were “non-truth-tracking,” even given that her guesses were correct and that the relevant truths were

¹⁵⁶ For a similar point, see Clarke-Doane (2012: 320-1).

¹⁵⁷ For a relevantly similar point, see Clarke-Doane (2012: 320-1) and Nicholas Sturgeon (1985). For a discussion of related matters see also Joyce (2014).

¹⁵⁸ In her ‘Reply to Copp’ (2008: 208), Street explicitly formulates her sceptical challenge in terms of ‘conceivability’: ‘as a *purely conceptual matter*, the normative truths could be anything [italics mine]’. See also the passage from Street (2008) I quoted in the main text.

metaphysically necessary. Even if it is not metaphysically possible that the relevant conjectures are false, it is intelligible to imagine that they are. And if—for all that we can intelligibly imagine—the relevant conjectures had been false, the child’s beliefs would have been the same. This raises a *prima facie* puzzle: what is the explanation for the child’s having so many true corresponding beliefs? It would be no answer to this puzzle to simply observe that the relevant conjectures are metaphysically necessary. (2012: 321)

Now, there are two ways in which *for all we can intelligibly imagine* the epistemic-normative facts could have been different:

(1) The evidence-facts they depend on could have been different; (2) The evidence-facts could have been the same, but the normative facts could have been different. For instance, there could have been no normative facts whatsoever (over and above evidence-facts)¹⁵⁹. If this were the case, normative facts would be different. And they would neither be identical to, nor strongly supervenient on, facts of the former kind.¹⁶⁰ For reasons I have already explained above, only the second of those scenarios would be such as to make the counterfactual we are discussing true.

Let us then focus on this latter and let us ask: what reasons are there to believe that it is *conceivable* that the epistemic-normative facts might have been different in this sense? And, whatever those reasons are, are they also such as to support the claim that the epistemic-normative facts might *easily* have been different in the specified way?

One obvious reason to believe that it is *conceivable* that epistemic-normative facts might have been different in the sense under discussion is that it is possible for competent users of epistemic-normative concepts to *intelligibly disagree* over what those facts are like or about whether they exist (Clarke-Doane; 2012: 335).

For instance, both anti-realists and ‘evidentialist’ realists alike will deny, if for different reasons, that epistemic-reasons facts are identical to evidence-facts. Thus they will both disagree with the ‘evidence-as-reason’ realists on this matter. Similarly, both anti-realists and ‘evidence-as-reasons’ realists alike will deny, again for different reasons, that epistemic-reasons facts are

¹⁵⁹ Where we do not take it to be a *conceptual truth* that evidence-facts are irreducibly normative facts.

¹⁶⁰ One way in which facts of a given kind might be different is by there not being any such facts at all.

distinct from but strongly supervenient upon evidence-facts. They will thus both disagree with ‘evidentialist’ realists about this. In general, the very context in which an argument like DEA is put forward and discussed would appear to presuppose the kind of ‘conceivability’ that the counterfactual claim expressed by CF’-E requires. And so does the possibility of drawing a taxonomy along the lines of the one I outlined in Chapter 3.

3.4 Predictability, Epistemic-Normative beliefs and Two Senses of ‘Lack of Correlation/Off-track’

But how does this interpretation of the claim that the epistemic-normative facts *might have been different* sit with the claim that they might *easily* have been so (i.e. D’-E)? It is not at all clear how one is to make sense of expressions such as ‘P could *easily* have been the case’, when P is the proposition that there are no normative facts (over and above evidence-facts), and when the modality at work in those claims is understood along the lines of *conceptual* possibility. One could look for help to the ways in which epistemologists have attempted to formulate luck (/anti-luck) intuitions similar to those which are expressed by the counterfactuals I am discussing. And understand, for instance, CF’-E and D’-E in terms of the idea that epistemic-normative beliefs would be such as to be ‘insensitive’ or ‘unsafe’ (Pritchard: 2008). But the most common way to spell out these notions relies on the idea of *nearby possible worlds*. And if this idea is hard enough to spell out in convincing ways when the modality at work is metaphysical modality,¹⁶¹ its problems will be magnified when the modality at issue is conceptual modality. For if the notion of a nearby possible world is hard enough to get a precise grip on, the notion of conceivability itself is notoriously problematic. What would it mean, for instance, for there to be a nearby *conceivable* world in which there are no epistemic-normative facts (and in which we nonetheless hold the epistemic-normative beliefs we hold)?

¹⁶¹ Much of the strength of important objections to ‘sensitivity’/‘safety’ accounts of the modal conditions which our beliefs must satisfy for them to amount to knowledge would appear to turn on disagreements with regard to which possible worlds are *nearby possible worlds*. See, for instance, Sosa (1999) and Pritchard (2008).

Given the difficulties I have just raised, one would seem to be faced with two options. The first is to drop D'-E as unintelligible or too problematic, and to rest one's interpretation of off-track/lack of correlation in the epistemic case on CF'-E alone, where we take it to be (conceivably) the case that the epistemic-normative facts could indeed have been different in the way I have specified. Call this first way to spell out the sense in which epistemic-normative beliefs would be off-track: *Off-Track – Sense 1*. The second option is to keep D'-E, and try to provide some intuitive way of assessing the truth of this thesis. Call this second way to spell out the sense in which epistemic-normative beliefs would be off-track: *Off-Track – Sense 2*.

With regard to this second option, it seems to me that a plausible way to understand *how easily* it could have been the case that there are no epistemic-normative facts would have this depend on *how likely* it is for realism about epistemic normativity to be true. And it seems also that how likely it is for epistemic-realism to be true will depend on *how good/bad are the reasons* for holding this view.

Clearly epistemic-realists and anti-realists will disagree here. For on this understanding of how to assess D'-E, epistemic realists of the kinds I am considering are likely to reject this claim. After all, they believe that they have good *reasons* to believe that there are epistemic-normative facts, and that these facts have the 'structure' that they take them to have. Thus, if they will surely grant that it is possible (i.e. conceivable) that their view is mistaken, they will nonetheless resist the claim that this could *easily* have been the case.

Moreover, the proponent of the Darwinian-epistemological argument I am discussing here holds that coming to believe that 'we could easily have been mistaken' about our epistemic-normative beliefs (on the assumption that epistemic realism is true) ought to be a consequence of appreciating the fact that evolutionary forces can be appealed to in order to explain why we hold the epistemic-normative beliefs we hold. On the understanding of the claim that 'we could easily have been mistaken about our epistemic-normative beliefs' I have just discussed, however, the truth of this claim requires that epistemic-normative facts might easily have been different (in

the sense that I have specified). But it is not at all clear how evolutionary considerations can be drawn upon in order to argue for the claim that *that there might easily have been no epistemic-normative facts*. On the understanding of D'-E I am discussing here, the truth of this claim will depend on how good the reasons we have for epistemic-realism are. And the epistemic-realist's belief in the existence of epistemic-normative facts is normally defended on the basis of considerations which are *completely independent* of 'the facts of evolution'.

If what I have said above is correct, then there are many reasons for resisting an interpretation of the claim that epistemic-normative beliefs are *off-track* along the lines of claims CF'-E and D'-E *combined* (Off-Track – Sense 2). The problem with this interpretation – I have argued – is D'-E.

But wouldn't DEA still work with an interpretation of off track/lack of correlation that did without this latter claim? Couldn't we understand the thesis that epistemic-normative beliefs are *off-track* as just amounting to the claim that we would have held the epistemic-normative beliefs we hold *even if there had not been facts of the epistemic-normative kind* (Off-Track - Sense 1)?

There would seem to be some clear advantages to this latter reading. Firstly, if true, it would still perhaps carry enough epistemological force for the argument that relies on it to raise a problem for the epistemic realist.¹⁶² Secondly, it does not require the truth of any controversial claim about what might *easily* have been the case. Thirdly, it is indeed a thesis that has better chances of being defended on the basis of evolutionary considerations alone. This interpretation of the claim that epistemic-normative beliefs are 'off-track' would then seem to have much to recommend it. Thus it would be most natural to read the thesis that epistemic-normative beliefs are off-track (if evolution explains them and if realism is true) at work in premise (4) of DEA along the lines of *Off-Track Sense 1*.

¹⁶² The argument would amount to something like: It is 'predictable' that we would have had the N-beliefs we have (evolution can explain them); and we would have had those beliefs whether or not there had been N-facts to make them true. Because of this, once this fact is brought to our attention, we ought to believe that our N-beliefs are not justified.

The distinction between the two senses of ‘*off-track*’ which I have drawn attention to in this section will play a crucial role when, in the next chapter, I will argue that the sceptical conclusion of DEA can be resisted and, hence, that DEA itself cannot be used to as part of a challenge against normative realist views in epistemology. Before doing so, I must turn to examine a central feature of the defence – by proponents of DEA – of the thesis that if evolution explains epistemic-normative beliefs, then those beliefs are off-track. In the discussion which follows, unless I make it explicit, the sense of ‘off-track’ which I will take to be at work must be understood as being either of those that I have distinguished in this section. .

4. When Do Evolutionary Explanations of Beliefs Support the Claim that those beliefs are ‘Off-Track’?

Why do proponents of DEA believe that if evolution explains why we hold the epistemic-normative beliefs we hold, then those beliefs are not correlated with the kind of facts the epistemic realist posits (premise (4) of DEA)? In the previous sections, I have briefly referred – if indirectly – to the kind of reasons evolutionary thinking might provide for the claim that there is a lack of correlation between normative beliefs and normative facts (of a certain kind) in the *practical* and *moral* domains.

In this section and the next, I draw attention to a central feature of the ‘evolutionary’ considerations which lead proponents of DEA to deny the correlation between normative beliefs and normative facts in the *epistemic domain*. Although the terminology I will employ in what follows is not theirs, considerations of the kind I will elucidate are at work in both Street’s (2009) and Gibbard’s (1990) discussions of the Darwinian-epistemological challenges to epistemic realism.

Ultimately, I do not believe that evolutionary considerations of the kind I am about to discuss will lead to the sceptical conclusion that DEA seeks to establish. Thus, to the extent to which these considerations are put to work to establish this conclusion, I take them to be defective in

important ways. But my argument for this will have to wait till the next chapter, where I will engage rather closely with what Street has to say on these topics. My aim in what follows is to continue to clarify the nature of DEA, while discussing yet another important aspect of one of its crucial premises.

I suggest we begin thinking about the topic at issue from the following passage by Walter Sinnott-Armstrong:

The evolutionary explanations [of our having the moral beliefs that we have] work even if there are no moral facts at all. The same point could not be made about mathematical beliefs. People evolved to believe that $2+3=5$, because they would not have survived if they had believed that $2+3=4$, but the reason they would not have survived then is that it is true that $2+3=5$. (2006:46)

In the passage I have quoted, Sinnott-Armstrong is comparing (justification for) two distinct kinds of beliefs we hold. Beliefs of the first kind (moral beliefs) are such that we can provide an evolutionary explanation for them. And the evolutionary explanation for those beliefs is such that it ought to undermine our confidence in those beliefs, when they are ‘realistically’ construed. Beliefs of the second kind (mathematical beliefs) are also such that we can provide an evolutionary explanation for them. But the evolutionary explanation for those beliefs is such that it does not undermine our confidence in those beliefs when realistically construed. Indeed, one could add, this explanation would rather appear to *vindicate* the justification for those beliefs which we thought we already had.

In the context of our discussion, this passage is particularly apt for a series of reasons. For one, it would appear to rely on the interpretation of ‘lack of correlation’ which – I have suggested – is the most plausible in the case of epistemic-normative beliefs. This is the one that can be summarised by the slogan: ‘we would have had the epistemic-normative beliefs we have, even if there were no normative facts to make them true’. The passage is also representative of a very common way of understanding what evolutionary explanations ought to be like in order for them to be deployed *selectively*, that is, in such a way as to call into question the justification for

one category of beliefs, but not for another. Lastly, it directly resonates with the kind of considerations that both Gibbard and Street would appear to rely on in building their case against *epistemic* realism (or for scepticism, on the assumption that epistemic realism is true). I will return to discuss those – Street’s in particular – in some detail in the next chapter.

What the passage I quoted suggests is that in order for evolutionary explanations to call into question the correlation between certain kinds of beliefs (R-beliefs) and certain kind of facts (R-facts), it must be the case that R-facts are not part of the evolutionary explanation for R-beliefs. Conversely, for evolutionary explanations not to call into question (or perhaps to vindicate) the correlation between R-beliefs and R-facts, the R-facts must be part of (or appealed to, or cited in, etc.) the correct evolutionary explanation for R-beliefs. However, this way of understanding how evolutionary considerations bear on epistemological issues – such as the justification for our beliefs of a certain kind – is too crude, at least if we are to take it too literally. One way to see this is by considering the fact that the truth makers for the vast majority of beliefs we hold true do not figure in the evolutionary explanations for those beliefs. This is so even in cases in which we believe evolutionary explanations for beliefs of the kind in question to be extremely plausible and to be such as to vindicate our justification for holding those beliefs true.

As I have already pointed out, philosophers who endorse DEA typically believe that evolutionary considerations do not call into doubt our justification for commonsense beliefs about mid-sized objects in our most immediate environment, or our beliefs about core parts of modern science, or about mathematics. Yet, it is hard to imagine that truths about microwave ovens, electromagnetic fields, or complex mathematical theorems were easy to come by in the environments in which our ancestors evolved their cognitive faculties. Thus it is hard to take seriously the view that these kinds of things and facts must be appealed to in the evolutionary explanation of why we hold the beliefs about them we do. If there is an evolutionary explanation for why we hold many of the beliefs of those kinds, and of why many of them are true, it will have to be – at best – a rather *indirect* one. But then the following question comes to the fore.

Can't whatever evolved faculties or mechanisms underlie our ability to form (many) true beliefs in those other domains also be trusted to deliver (many) truths about, say, epistemic-normative facts? After all, we take our cognitive abilities to be reliable in many and very diverse domains, and the domain of epistemic-normative facts might just well be yet another one of them.

The reason for doubting that epistemic-normative beliefs correlate with epistemic-normative facts must originate in the thought that epistemic-normative beliefs are somewhat different from those others categories of beliefs. Which, again, leads us to the question: how must we understand this difference? I have just argued that one common way to do this – in terms of R-facts figuring in evolutionary explanation for R-beliefs – fails. In the next section I explore a better way to think about these questions.

4.1. Of Milvian Bridges

Let us keep in mind what I have said so far, and go back to some of the terminology I have introduced in earlier sections.

It seems to me a better way of understanding the kind of evolutionary considerations which proponents of DEA draw on is to take proponents of this argument as defending the view that each of the following is the case:

(1) The natural function N of the mechanisms M which give rise to beliefs of kind R (R-beliefs) *is not* that of correctly detecting and representing states of affairs in the world which correspond to R-facts.

(2) Whatever N is, performing N *does not require* the ability to correctly detect and represent states of affairs in the world which correspond to R-facts.

(3) The ability to detect and represent states of affairs in the world which correspond to R-facts *is not a by-product* of N.

and to conclude from this that:

(4) R-beliefs are not correlated to R-facts.

Let us say that there is a ‘Milvian Bridge’ (Griffiths and Wilkins: 2010) which connects R-beliefs and R-facts when the following is true:

Milvian Bridge: R-facts are related to the natural function N of the mechanisms which give rise to R-beliefs in such a way that it is reasonable to believe that R-beliefs are correlated to R-facts.

I take it that believing that each of (1) to (3) above is the case is tantamount to defending the view that there is ‘no Milvian Bridge’ between R-beliefs and R-facts. Thus we can rely on (1) to (3) to specify the ways in which R-beliefs and R-facts ought to be related in order for there to be a ‘Milvian Bridge’ to link them. R-beliefs must be such that they can be traced back to mechanisms M the natural function N of which is such that either (1) it coincides with the ability to represent R-facts, or (2) it entails the ability to represent R-facts, or (3) it has this latter ability as a by-product.

Introducing the concept of a *Milvian Bridge* allows us to make progress on the two (related) questions I have raised above. The first question was after a plausible way of formulating the considerations on which proponents of DEA rely when they defend the thesis that there is no correlation between epistemic-normative and epistemic-normative facts. The answer to this question, I am suggesting, is that proponents of this argument believe that there is no Milvian bridge connecting beliefs and facts of the kind in question. Notice how this answer does not require that we must (directly) appeal to R-facts in order to provide an evolutionary explanation for R-beliefs. On this model the ability to correctly represent R-facts might be – for instance – a by-product of evolved abilities to correctly represent facts of other kinds.

The second question was after a plausible way of understanding (again, on behalf of the proponent of DEA) the difference between beliefs the justification for which is called into question by evolutionary considerations and beliefs which are immune from this challenge. It seems to me there are two obvious ways of drawing the distinction at issue on the basis of the conceptual tools I have adopted here. The first is by saying that beliefs the justification for *which is challenged* by evolutionary considerations are such that *we have reasons* to believe that there is no Milvian bridge connecting those beliefs to the facts which would make them true. On the other hand, beliefs of a kind *which is not challenged* by evolutionary considerations are such that we have *no reason* to believe that there is no Milvian bridge connecting those beliefs to the facts which make them true.

An alternative way of drawing this distinction would be to define the *challenged* beliefs in the same way. But to say that the categories of beliefs which are *not challenged* are such that there are *positive reasons* to think that there is a Milvian bridge connecting them to the facts which would make them true. A view of roughly this kind is what seems to be at work in the passage by Sinnott-Armstrong I have quoted above. Deciding which of those alternatives best captures the distinction at hand raises interesting and difficult issues¹⁶³ which I cannot address here with the attention they would deserve.

Whether this latter question is settled in either of the ways I have explored, will not affect the nature of DEA. For in either case, it will still be the case that proponents of this argument will argue that the *absence* of a Milvian bridge suffices to provide a defence for the lack of correlation

¹⁶³ In most ordinary contexts – it seems to me – when we think we have good reasons to hold beliefs of a certain kind, we don’t need *to also prove* (i.e. to produce strong positive reasons to believe) that there is a Milvian Bridge between beliefs of this category and the facts which would make them true in order to be able to safely retain justification for them. One might think this is because we can assume – by default, so to speak – that *some evolutionary explanation or other* for them being correlated to the facts which they represent must be easily available. It seems to me this ‘assumption’ is problematic. I have already drawn attention to some of the reasons for believing that we ought not to take questions such as these for granted (see Section 4.3.1 of this chapter). Whether the optimism behind the assumption at issue is warranted and, if not, how this would affect that status of Darwinian-epistemological arguments such as the one I have been discussing are topics which I do not have the space to discuss here.

thesis DEA itself draws upon. This is then the thesis we ought to focus on here; for it is on this thesis that premise (4) of the argument draws.

5. Conclusion

In this chapter, I have argued that the reason why proponents of DEA believe that epistemic-normative beliefs are off-track (given that evolution can explain them) is that they believe there to be no Milvian Bridge between epistemic-normative beliefs and epistemic-normative facts. But is it true that we can appeal to evolutionary considerations in order to defend the thesis that there is no Milvian bridge between epistemic-normative beliefs and epistemic-normative facts? And even if this were true, would this entail that epistemic-normative beliefs are *not correlated* with epistemic-normative facts? If it does, what sense of correlation must be at work for this to be the case? And is establishing that epistemic-normative beliefs are off-track *in this sense* sufficient to warrant the sceptical conclusion of DEA?

I will address these questions in the next chapter, where I consider Street's Darwinian-epistemological challenge against epistemic realism and I defend epistemic realism against it.

Part III

Chapter 7

Street's Darwinian-Epistemological Argument: An Epistemic-Realist Reply

1. Introduction

At the end of the previous chapter I discussed an important aspect of premise (4) of DEA. That premise, we will remember, had the form of a conditional. It said that if (a) evolution explains why we hold the epistemic-normative beliefs we hold then (b) there is no correlation between epistemic-normative beliefs and epistemic-normative facts. I have defended the view that proponents of DEA are as best interpreted as believing that (b) is true because they take evolutionary explanations of our epistemic-normative beliefs to entail the claim that *there is no Milvian bridge* between epistemic-normative beliefs and epistemic-normative facts – call this thesis (c), and because they believe that the absence of a Milvian bridge, in turn, entails the lack of correlation between epistemic-normative beliefs and epistemic-normative facts. We can summarise all this as follows. Proponents of DEA believe that:

STEP 1 – thesis: If (a) evolution explains our epistemic-normative beliefs, then (c) there is *no Milvian bridge* between epistemic-normative beliefs and epistemic-normative facts.

and:

STEP 2 – thesis: If (c) there is *no Milvian bridge* between epistemic-normative beliefs and epistemic-normative facts, then (b) there is *no correlation* between epistemic-normative beliefs and epistemic-normative facts.

In this chapter, I examine an argument which – I will argue – is best viewed as an instance of DEA. This is the Darwinian-epistemological challenge to epistemic realism that Street defends (Street; 2009). In particular, I will discuss in some detail the kind of considerations Street relies on in defending the thesis that if evolution explains our epistemic-normative beliefs, then those beliefs are *off-track* (Premise (4) of DEA).

This will require looking at why Street believes that *there is no Milvian Bridge* between epistemic-normative beliefs and epistemic-normative facts, and at why she believes the absence of a Milvian bridge can motivate a sceptical conclusion of the kind arguments like DEA seek to establish. My discussion of Street will thus provide a direct insight into the kind of evolutionary considerations which proponents of arguments like DEA draw on in order to establish that if epistemic-realism is true, then epistemic-normative beliefs are off-track. The discussion will thus flesh out – and provide an instance of – the general strategy of defence of premise (4) of DEA which I discussed, in an abstract form, at the end of the previous chapter.

Against Street I will argue that, on the assumption that evolutionary explanations of epistemic-normative beliefs of the kind she discusses are correct, the thesis that those beliefs are not justified does not follow.

On some versions of epistemic-realism which fall within the target of DEA – those according to which epistemic-reasons facts are identical with facts about evidence – there would seem to be a straightforward Milvian bridge between epistemic-normative beliefs and epistemic-normative facts. Thus STEP 1 – thesis would be false. And this would secure the claim that, contrary to what premise (4) of DEA claims, those beliefs *are indeed* correlated with the facts which make them true. Thus the inference at STEP 2 – thesis would also be false. This will be sufficient to block DEA's sceptical conclusion.

On different versions of epistemic-normative realism which fall within the target of DEA – those according to which epistemic-reasons facts are *not identical* with facts about evidence – the situation is be more complex. On one sense of lack of correlation (*off-track – sense 1*), it will

indeed appear to be true that if evolution explains our epistemic-normative beliefs, then there will be *no Milvian Bridge* between epistemic-normative beliefs and epistemic-normative facts. Thus STEP 1 – thesis would be true. And it will also be true that the absence of a Milvian bridge will entail the lack of correlation between epistemic-normative beliefs and epistemic-normative facts. Thus STEP 2 – thesis will also be true. Yet this – I will argue – does not suffice to warrant the sceptical conclusion of the argument.

On a different sense of lack of correlation (*off-track – sense 2*), it will be true that the evolutionary influences on our epistemic-normative beliefs make it so that there is *no Milvian bridge* between those beliefs and epistemic-normative facts. Thus STEP 1 – thesis will be true. Yet it would not follow from this that epistemic-normative beliefs are *not correlated* to epistemic-normative facts. Thus STEP 2 – thesis will be false. Establishing the lack of correlation between epistemic-normative beliefs and epistemic-normative facts in the sense at issue would indeed suffice to conclude that epistemic-normative beliefs are not justified. However – given the failure of STEP 2 – thesis – we cannot infer lack of correlation in this sense on the basis of evolutionary considerations. Thus, once again, the sceptical conclusion of DEA will not follow.

2. Street on the Darwinian Epistemological Problems of Epistemic Realism

In her paper ‘Evolution and the normativity of epistemic reasons’ (2009) Sharon Street defends the view that evolutionary considerations can be drawn upon in order to show that realism about epistemic normativity faces an ‘unacceptable epistemological problem’. In what follows, I argue that the kind of epistemological problem Street raises for epistemic realism are best read as the result of Street endorsing an argument with the form of DEA. I look in some detail at the evolutionary considerations that Street relies on in order to defend the thesis that assuming the truth of epistemic-realism would lead to our epistemic-normative beliefs being unjustified. I then provide a response to Street’s arguments on behalf of the realist.

Although – as I will return to make clear – Street does not *explicitly* endorse an argument such as DEA, there are strong reasons for focusing our discussion on her work. The first is that – as I have already pointed out – the considerations she puts forward are best read as providing an instance of such an argument. The second is that her epistemological challenge to epistemic realism is sensitive to the fact that normative realism in the epistemic-domain can take importantly different forms. The third is that Street’s account constitutes one of the best worked out accounts in the literature as to how evolution can be relied upon in order to formulate an epistemological challenge for versions of normative realism *in the epistemic domain*.¹⁶⁴

Let us then begin to look at why Street believes that ‘the [epistemic] realist has *no* [...] *acceptable explanation* of the striking coincidence’ (2009: 234) between epistemic-normative beliefs of the kind we hold and the epistemic-normative facts that realism posits. Street distinguishes two varieties of epistemic-normative realism, what she calls ‘Realism about Cliffordian Normativity’ and ‘Normative Realism about Likelihood’ (Street; 2009: 213-227).¹⁶⁵ These roughly correspond to the varieties of epistemic realism which I have called ‘evidentialist realism’ and ‘evidence-as-reasons’ realism. Street argues against each of those views separately. I will look at each of her arguments in turn.

¹⁶⁴ In each of these respects Street’s discussion is a better focus for our investigations here than, for instance, Gibbard’s Darwinian argument(s) against normative realism in Gibbard (1990).

¹⁶⁵ Street’s terminology – as well as the way in which she distinguishes the varieties of epistemic-realism she discusses – is, in many respects, misleading. On the one hand, it leads to a conflation between what I have called Non-Humean instrumentalist and Evidentialist interpretations of epistemic-reasons. On the other, it fails to distinguish claims about evidence from claims about probability. While the claim that epistemic reasons depend on evidence – in the minimal sense I have specified in Chapter 2 – is, I take it, fairly uncontroversial, the claim that epistemic-reasons depend on probability facts is on far less secure conceptual grounds. For many philosophers think that facts about evidence do not reduce to facts about probability. Whether facts about evidence are to be understood in this way is a vexed question in the philosophy of science and in epistemology. See, for instance, Achinstein (1978). Street’s failure to draw the distinctions I have just indicated affects both the strength and the breath of her arguments against epistemic realism. This will become apparent from my discussion in the rest of this chapter.

3. Street's Epistemological Argument against Evidentialist Realism

The general form of Street's argument should by now be familiar. Street argues that given that evolution can be called upon to explain why we hold the epistemic-normative beliefs we hold, if we were to assume the truth of epistemic-realism, we could not provide an explanation for the correlation between normative-beliefs and normative facts. Why does Street believe this to be the case? In short: because she believes that the only explanation available to the epistemic realist here is 'scientifically unacceptable' (2009: 234).

As we will remember, evidentialist realists believe that epistemic-reason facts are distinct from but strongly supervenient upon facts about evidence. The evidentialist realist takes what I have called the evidentialist interpretation of epistemic-reason claims at face value. According to this version of epistemic realism, epistemic-normative expressions of the form '*E is an epistemic reason for P*' are equivalent in meaning to expressions of the form 'the fact that E is evidence for P provides a reason to believe that P'. And expressions of this latter form are such as to 'mirror' and 'reveal' the structure of the complex facts which make true epistemic-normative beliefs true.

Street grants evidentialist realists that epistemic-normative beliefs are best understood along the lines that their 'evidentialist' conceptual analysis suggests. She takes it to follow from this analysis that creatures who believe that E is a reason for P take considerations regarding what is evidence for what as providing normative reasons for or against holding beliefs of various kinds. And they do so 'both as a fairly routine, unconscious matter, and also when they deliberate about what to believe in full consciousness' (2009:227). Thus they believe that F provides a reason to believe that Q, G a reason to believe that R, etc. on the basis of the fact that (they believe that) F, G, etc. are *evidence for*, respectively, Q, R, etc.

She also assumes that the thesis that there is an evolutionary explanation for why we hold many specific epistemic-normative beliefs of this kind has 'obvious plausibility': 'it tends to promote survival and reproduction to have roughly true beliefs about one's environment' (2009: 235). Taking considerations of evidence to determine what one ought to believe is an obvious

way to secure the acquisition of a large number of true beliefs. This explains why creatures like ourselves have a tendency to form many epistemic-normative beliefs of this kind.

Given that our having those normative beliefs can be explained by evolution, Street claims, unless the epistemic realist can draw upon this fact to *also* explain the correlation between those beliefs and epistemic-normative facts, then those normative beliefs would have to be considered ‘off-track’(2009: 234). As I have anticipated, Street believes the explanation that the realist can provide here is ‘unacceptable’. She thinks this explanation can only be of the following kind – what she calls the ‘tracking account’.¹⁶⁶

Creatures who took truth-related considerations [i.e. considerations about what is evidence for what] to be reasons for belief did better in the struggle to survive and reproduce because the former creatures grasped an independent normative truth that the others did not. In other words, on this account, it is an independent normative truth that truth-related considerations are reasons for or against belief, and creatures who accepted the corresponding judgment (or better ‘proto’ versions of it) did better because they were grasping an independent normative truth. (2009: 234)

And she compares this explanation with her preferred ‘competitor’ to the above account. She calls this latter the ‘adaptive-link account’. According to this account, creatures like ourselves routinely take considerations of evidence to provide normative reasons for or against beliefs and, hence, hold many beliefs such as that E is evidence for P is a reason to believe that P, not because creatures of this kind:

Grasped an independent normative truth that the other types of creatures missed; it is rather, much more simply, that [this type of] creatures tended to arrive at beliefs about their environments that were true [...] and this, in turn, was advantageous for the purposes of survival and reproduction. (2009: 235)

Street goes on to say that:

The adaptive-link account [...] defeats the tracking account hands down as a scientific explanation of why we ended up [with epistemic-normative beliefs of the kind we hold]. (2009: 236)

¹⁶⁶ Street’s vocabulary of choice here clearly reflects the one she adopted in her paper *Darwinian Dilemma* (2006). This is to be expected, given that she here is (allegedly) framing an argument against epistemic realism similar to the one she outlines in that paper.

The reasons Street believes this to be the case is that the adaptive-link account explanation is much better in terms of familiar criteria we use to evaluate scientific theories, such as ‘parsimony, clarity, and explanatory power’ (2009: 241). I will not comment on this argument at any length here. However, it is worth bringing out the structure of the argument I have just outlined. This – it seems to me – runs something like the following:

- (1) The *only* (evolutionary/scientific) explanation that the epistemic-realist can provide *of the correlation* between epistemic-normative beliefs and epistemic-normative facts is the tracking account (TA).
- (2) There is a much better (evolutionary/scientific) explanation – the adaptive-link account (AL) – *of why we hold the epistemic-normative beliefs* we hold.
- (3) The adaptive link account (AL) does not at any point require mention of epistemic-normative facts.

Therefore:

- (4) The tracking account (TR) is false.

Therefore:

- (5) The epistemic-realist can’t provide *any* (acceptable) explanation of the correlation between epistemic-normative beliefs and epistemic-normative facts.¹⁶⁷

I think there are many rather serious problems with Street’s argument as I have presented it here. Some of them are as follows. For one, as I will return to argue, there are good reasons to believe that its first premise is false. It is not clear what ought to count as a scientific/evolutionary explanation here. On a certain understanding of what scientific/evolutionary explanations

¹⁶⁷ Street seems to believe that the ‘tracking account’ would be no better than saying that ‘evolution landed us on the independent normative truths about such matters purely as a matter of chance’ (2009: 234).

require¹⁶⁸ it is perhaps the case that epistemic-realists *cannot* provide an explanation of this kind. But they can nonetheless provide *another explanation* which is, I will claim, both ‘acceptable’ and ‘compatible’ with the existence of the scientific/evolutionary explanation in question. And this explanation is – it seems to me – *sufficient* to deny that there is a lack of correlation (between epistemic-normative beliefs and epistemic-normative facts) of the kind that Street’s argument threatens to establish.¹⁶⁹

Secondly, notice that, as things stand the tracking and the adaptive-link accounts are not even – at this stage – *competing* explanations. The first is (an attempt at) an evolutionary/scientific explanation of the *correlation* between epistemic-normative beliefs and epistemic-normative facts; the second an (alleged) scientific/evolutionary explanation of *why we hold the epistemic-normative beliefs we hold*. If the two accounts have different *explananda*, they then can hardly *compete* with one another. And if they are not competitors, one can hardly conclude that the latter is much better than the former. It is true that Street believes that if we were to adopt her own meta-normative view of epistemic normativity: (1) the correlation between epistemic-normative beliefs and epistemic-normative facts will be explained, (2) it will be explained in a way which is compatible with endorsing the adaptive-link account, and (3) this explanation of the correlation in question will be much better than that provided by the combination of evidentialist-realism and the tracking account. Setting aside the fact that it is far from obvious that (3) is the case,¹⁷⁰ it is (3) that Street would need to discuss in order to compare two rival explanations of the same phenomenon. But Street takes the conclusion of her argument to be established before she ever discusses this claim.¹⁷¹

¹⁶⁸ Is an explanation of the kind: ‘evolution explains why we track evidence-facts, epistemic-reasons strongly supervene on evidence-facts, therefore: evolution explains why we track evidence-facts’ an evolutionary explanation? Parfit (2011, Ch: 33) and Enoch (2010) seem to believe so, although in a qualified way. If evolutionary explanations of why we ‘track’ facts of some kind require causal interaction with those facts, then a non-naturalist realist cannot help himself to them. If they don’t (perhaps in the way I have sketched above), then perhaps they can. The matter – it seems to me – is to a large extent a *terminological* one.

¹⁶⁹ I will return to elaborate on this point in later sections of this chapter.

¹⁷⁰ A defence of this claim will take more space than I have at my disposal here. My reply to Street’s arguments will not, ultimately, depend on arguing for this claim.

¹⁷¹ Indeed she never directly compares the two rival explanations I have just drawn attention to in her paper.

Thirdly, notice that, in any case, it does not follow from (2) the adaptive link account is much better than the tracking account that (4) the tracking account is unacceptable; and to the point of being *no account* at all as (5) implies. The tracking account could be a worse account in many respects, but still count *as an 'acceptable' account* nonetheless.¹⁷²

3.1 A Better Argument

I think there is a quick way to turn Street's argument against epistemic realism here into a much better argument. The reading I am about to suggest is also more in tune with Street's claim that the kind of epistemological problems that epistemic realism faces are – in significant respects – analogous to the ones that normative realism face in the practical case (Street: 2006). Street herself explicitly says that her Darwinian challenge to epistemic-normative realism is to be read along the lines of the Darwinian challenge she articulates against normative realism in the practical domain. Thus the reading of Street's argument I am about to suggest would seem to have much to recommend it.

Here is my reconstruction of Street's argument:

- (1) Evolution explains our epistemic-normative beliefs.
- (2) The best evolutionary account of why we hold the epistemic-normative beliefs we hold – i.e. the adaptive-link account (AL) – explains those beliefs in terms of mechanisms M the natural function – N – of which *is not* that of correctly representing epistemic-normative facts (of the Evidentialist kind). Neither is the ability to correctly represent normative facts *required by* or a *by-product* of those mechanisms M performing N.

¹⁷² One can appeal here to the fact that there are some situations in which an explanation is *so bad* in comparison with its alternatives, that we can safely say that it counts as *no explanation* at all. And that Street has such a context in mind when she claims what she does here. However, whether this is the situation we are in fact in when it comes to the case under consideration is far from obvious. And settling whether this is the case would require discussion of the issues I mention in the previous paragraph. Prior to doing so, we cannot assume that we can safely infer from *bad explanation* to *no explanation* as Street does.

Therefore:

(3) There is *no Milvian bridge* between epistemic-normative beliefs and epistemic-normative facts.

Add that:

(4) If there is no Milvian bridge between epistemic-normative beliefs and epistemic-normative facts, then epistemic-normative beliefs are off-track with respect to epistemic-normative facts.

We can then conclude that:

(5) Epistemic-normative beliefs are off-track with respect to the kind of facts that – according to Evidentialist realism – would make them true.

And if we assume that:

(6) Evidentialist realism is true

We can conclude that:

(7) Epistemic-normative beliefs are not justified.

It is easy enough to recognise in the argument I have just outlined a slightly more complicated version (a variation) of an argument of the kind of DEA. If we leave these complications aside, we can count the argument above as an *instance* of DEA.

But what are the reasons for thinking that (2) is supported by Street's 'adaptive-link' account? They are as follows. Street seems to believe that the natural function of those mechanisms is to get creatures like us to form beliefs about our environment which are *true* (2009: 236). Presumably, she means that the natural function of those mechanisms was to allow our ancestors to form true beliefs (or accurate 'belief-like' representations) about *their most immediate environments*, and that we can trust those mechanisms to yield – as a by-product of their operations – true beliefs in very different environments; such as the one creatures like us currently move in.¹⁷³ In any case, all that is required – in terms of correct representation – for those mechanisms to

¹⁷³ As I have already argued, something like this claim is required by the 'strong' Darwinian thesis the DEA requires. This interpretation is also independently supported by what Street says elsewhere in this paper and in her 'Darwinian Dilemma' paper (2006).

perform their function is that they accurately represent *evidence-facts*. And if those mechanisms are still to yield true beliefs in very ‘distant’ environments from the one our ancestors evolved in, they must still – as a by-product of their original natural function – be accurately representing *evidence-facts* in those environments.

If all this is correct, then Street can say that we would have held many of the epistemic-normative beliefs we hold *whether there had been epistemic-normative facts or not*. And this – one might think – lends support to the claim that epistemic-normative beliefs are *off-track* with respect to epistemic-normative facts construed along the lines of evidentialist realism. This fact can then in turn be drawn upon in order to defend the thesis that – on the assumption that evidentialist-realism is true – our epistemic-normative beliefs are unjustified.

3.2 An Evidentialist-Realist Reply to Street’s Argument

How compelling is this argument? Does it establish *the lack of correlation* between epistemic-normative beliefs and epistemic-normative facts that it purports to establish? And does it provide warrant for the claim that, if evidentialism is true, our epistemic-normative beliefs are not justified? Notice how Street’s argument relies on the two steps/theses I have discussed at the beginning of this chapter. If we update those on the basis of my preceding discussion of Street, we have:

*STEP 1 – thesis**: If (a) *adaptive-link account (AL)* explains our epistemic-normative beliefs, then (c) there is *no Milvian bridge* between epistemic-normative beliefs and *evidentialist*¹⁷⁴ epistemic-normative facts.

And:

¹⁷⁴ By ‘evidentialist epistemic-normative facts’ I mean epistemic-normative facts of the kind the evidentialist-realist believes in/posits.

*STEP 2 – thesis**: If (c) there is *no Milvian bridge* between epistemic-normative beliefs and *evidentialist* epistemic-normative facts, then (b) there is *no correlation* between epistemic-normative beliefs and *evidentialist* epistemic-normative facts.

As I have anticipated in the introduction to this chapter, I believe there are two kinds of reply available to the evidentialist realist here. Each of them is compatible with assuming that the evolutionary explanations that Street relies on in her argument are correct.

On one of the two senses of off-track which I have already discussed (off-track – sense 1), the evidentialist can reply that although it is true that Street’s evolutionary considerations lend support to the claim that epistemic-normative beliefs are off-track in this sense, this will not warrant the sceptical conclusion of her argument (and hence of DEA).

On a different sense of off-track (off-track – sense 2) the evidentialist-realist will say that Street’s evolutionary considerations do not suffice to establish that epistemic-normative beliefs are off-track in this sense. Hence that the sceptical conclusion of her argument (and hence of DEA) does not follow.

I will return to articulate those replies in some detail soon. Before doing so, however, I want to explore a different line of reply to Street’s argument; one that does not assume either of the senses of ‘off-track’ I have relied on in the replies to her argument which I have sketched above.

3.2.1 Street on ‘Lack of Correlation’

To be fair to Street’s original discussion, we must acknowledge that she relies on a rather different – and much more intuitive - understanding of ‘lack of correlation’ than the one(s) I have arrived at in my previous discussion of how this notion is best understood in arguments like DEA (Chapter 6, Section 3.4). She thinks that if we grant the truth of evidentialism and if we acknowledge the fact that evolution explains why we hold the epistemic-normative beliefs we

hold, the evidentialist won't be able to explain why: something like the following bi-conditional holds:

LC: for a large number of epistemic-normative beliefs which creatures like us hold, we hold those beliefs iff those beliefs are true.

Short of appealing to an 'implausible coincidence', Street believes that the evidentialist will not be able to account for the 'striking' line up between epistemic-normative beliefs and epistemic-normative facts which this bi-conditional expresses. It seems to me that – if this were the challenge – the evidentialist will have a rather non-mysterious reply to it, although not, perhaps, the kind of reply Street thinks the realist must resort to.

Indeed, the first thing to notice is that 'the tracking account' which Street believes would constitute *the only explanation* (2009: 234) available to the realist here is, in fact, a kind of explanation that the realist is rather unlikely to want to help himself to. A paradigmatic evidentialist is Derek Parfit.¹⁷⁵ Parfit believes that epistemic-reason facts strongly supervene, but are not identical to what he calls 'alethic facts' (what I have been calling evidence-facts). He also believes that epistemic-reason facts are *not causally efficacious*.

As we saw, Street presents the 'tracking account' as a competing scientific/evolutionary explanation to the adaptive link account. But if evolutionary/scientific explanations are a kind of causal explanation, then it will presumably have to be the case that the categories of facts which figure in them must be of the *causally efficacious* kind. But paradigmatic evidentialists do not take normative facts to be causally efficacious (Parfit; 2011). Hence the 'tracking account' alternative

¹⁷⁵ David Enoch (2010) does not discuss epistemic-normativity separately. But it is likely that his *Robust Metanormative Realism* about normativity, when it comes to normativity *in the epistemic domain*, will take a form rather like that of Parfit's evidentialist realism. This interpretation would appear to be somewhat supported by the fact that Enoch's reply to the Darwinian-epistemological challenge against practical-normative realism in Enoch (2010) is structurally similar to Parfit's reply to a version of that challenge directed against normative realism in the epistemic domain. See Parfit (2011: Ch. 33).

that Street discusses is unlikely to be a live option for them.¹⁷⁶ What is instead a rather more plausible option for the evidentialist here is something like the following.¹⁷⁷

The evidentialist can accept something along the lines of Street's adaptive link account, but then go on to say that this is compatible with thinking that (B) is true, and that it is not so in virtue of 'a fortuitous accident'. This is because if evidentialism is true, epistemic-normative facts will strongly supervene on evidence facts. But as we saw above, Street's adaptive-link account must have 'built into it', so to speak, the fact that the adaptive mechanisms which explain why creatures like us hold the epistemic-normative beliefs they hold, must be sensitive to (i.e. 'correctly represent') evidence-facts. This is, after all, part of what explained why creatures who exhibited the belief-forming tendencies which those mechanisms are responsible for are such as to form beliefs about their – immediate and less immediate – environments which are largely true.

As we saw, according to Street, those adaptive-link mechanisms are also such as to ultimately give rise to the epistemic-normative beliefs we hold. If this is the case, then the adaptive link account would appear to entail that there is a Milvian Bridge between *epistemic-normative beliefs* and *evidence-facts*. But if there is a Milvian Bridge between epistemic-normative beliefs and *evidence facts*, and if epistemic-reason facts strongly supervene on evidence-facts, then it will be the case – of many of the epistemic-normative beliefs we hold – that we hold those beliefs iff those beliefs are true. Hence the truth of (B) will be accounted for, and without having to appeal to any miracle or fortuitous coincidence.

To be sure, the correlation in question will not have been explained *directly* in terms of the adaptive-link mechanisms correctly representing *epistemic-reason facts*. On this account, we would have the epistemic-normative beliefs we have not *because it is reproductively advantageous to correctly*

¹⁷⁶ To be fair to Street, Parfit does at times write as if he were subscribing to something like the 'tracking account'. However, it is clear from his extended discussion of these matters (2011; Ch. 33) that his actual explanation of the correlation between epistemic-normative beliefs and epistemic normative facts is rather one along the lines of that which I discuss below.

¹⁷⁷ For analogous strategies in the practical-normative and moral domains, see Parfit (2011: Ch.33), Enoch (2010), Copp (2008) and Wielenberg (2010), among others.

represent epistemic-normative facts. Neither will the ability to correctly represent epistemic-normative facts be *required* by the natural function realised by the mechanisms which ultimately give rise to (and ‘explain’) epistemic-normative beliefs. Lastly, the ability to correctly represent normative facts will also not appear to be *a by-product* of the operations of those mechanisms. In this sense, there will be *no Milvian Bridge* between epistemic-normative beliefs and epistemic normative facts of the evidentialist kind.

And if the presence of a Milvian Bridge is required for the explanation that Street is demanding of the realist here to count *as a scientific/evolutionary explanation*, then Street will indeed be right in saying that the realist can’t provide any such *scientific* explanation.

Nonetheless, the realist account I have just sketched is (1) compatible with there being a scientific/evolutionary explanation of why we hold the epistemic-normative beliefs we hold; (2) compatible with believing that the best evolutionary-explanation in question will proceed along the lines sketched by Street’s adaptive-link account; and (3) such as to provide an explanation – although not perhaps a *scientific* one – of the ‘striking’ correlation between (many of) the epistemic-normative beliefs we hold true and the epistemic-normative facts the evidentialist realist posits. In other words, given the truth of evidentialism, and given the possibility to explain epistemic-normative beliefs in evolutionary terms, the ‘striking line up’ between normative beliefs and normative truths would not appear to be all that ‘striking’ after all.

Notice how the evidentialist-realist explanation at issue need not be the ‘best possible (evolutionary) explanation around’ for it to be such as to forestall the implausible ‘miraculous coincidence’ scenario that Street believes threatens if we assume that evidentialism is true. It will suffice that it is an *acceptable* explanation. And neither what Street says in her (original) argument nor what she can be taken to say in the way I have reconstructed her argument (i.e. as an instance of DEA) would appear to provide reasons for the claim that the explanation I have just discussed is *unacceptable*.

3.2.2 The Evidentialist Reply: Different Senses of ‘Off-track’

Can an argument of the kind I have attributed to Street be made more compelling, if we read the claim that epistemic-normative beliefs are *off-track* along the lines I have discussed in the previous chapter? There, we will remember, I discussed how the notion of ‘off-track/lack of correlation’ at work in arguments such as DEA was best to be understood. I suggested that we narrow the interpretations of this notion down to two distinct senses (sense 1 and sense 2), and I briefly discussed some of the advantages which one of those two interpretations (sense 1) would seem to present (Chapter 6; Section 3.4).

Let us then begin with this sense, what I called: *Off-track – sense 1*. This amounted to the claim that:

Off-track – sense 1: Had the epistemic-normative facts been different, our epistemic-normative beliefs would have been the same.¹⁷⁸

Here the modality invoked by the counterfactual is *conceptual* and the *conceivability* of the epistemic-normative facts being very different is understood as it being conceivable that there are no epistemic-normative facts of the kind the realist posits. In the specific case under discussion – that of evidentialist-realism – I also suggested that we understand the conceivability of there being no epistemic-normative facts of the kind the evidentialist realist posits in terms of *the conceivability of there being no epistemic-reason facts over and above evidence-facts*;¹⁷⁹ in other words, in terms of it being conceivable that *there are no epistemic-reason facts strongly supervening upon evidence-facts*.

¹⁷⁸ This was the thesis I called CF’-E in Section 3.4 of the previous chapter.

¹⁷⁹ One might think it ought to be possible to read what Street says in her paper as relying on the *conceivability* that epistemic-normative facts are very different in a different sense than the one I am taking into consideration here. That is, Street could be taken to say that it is conceivable that *the kind of epistemic-reason-providing facts* could have been very different. For instance, reasons for beliefs could have been provided by considerations regarding what beliefs would be pleasant to have, or fulfil an agent’s wish, etc. If this had been the case – because of the evolutionary influences on our epistemic-normative beliefs – we would still have had the epistemic-normative beliefs we have. Thus there would be a conceivable world in which the epistemic-normative facts are very different, but in which we would still hold the epistemic-normative beliefs we currently hold. In this world, epistemic-normative beliefs would not be correlated with epistemic-normative facts, and they would be ‘off-track’ with respect to those facts. The obvious problem with this scenario is that *the concept* of an epistemic reason Street is explicitly operating with is that of a reason for belief which is appropriately sensitive to considerations which bear on the truth of a belief. In other

An epistemological challenge formulated with this understanding of lack of correlation in mind would appear *prima facie* to have more traction than one formulated with Street's more generic and intuitive understanding of lack of correlation. For it is true that the evidentialist-realist can explain the correlation between epistemic-normative beliefs and epistemic-normative facts on the interpretation of 'correlation'/'lack of correlation' that seems to be work at work in Street (2009). Yet, even if that explanation stands, there remains a sense in which it would still have been true that we would have had the epistemic-normative beliefs we have *whether or not they had been true*; and hence that had the epistemic-normative facts been very different, our epistemic-normative beliefs would have been the same. This will be the case if we *adopt off-track – sense 1* as our preferred way of specifying how the lack of correlation between epistemic-normative beliefs and epistemic-normative facts ought to be understood.

Given that epistemic-normative facts of the evidentialist kind need not enter into the explanation of our epistemic-normative beliefs (if Street's evolutionary account is correct), then had there not been facts of this latter kind, we would still have come to hold the epistemic-normative beliefs we hold. In the conceivable world in which there are no evidentialist epistemic-normative facts (in the sense I have specified) and in which we still hold the epistemic-normative beliefs we hold, *there would be no Milvian bridge* between epistemic-normative beliefs and epistemic-normative facts¹⁸⁰ of the kind the evidentialist posits.

Thus *STEP 1 – thesis** would be true. On the sense of lack of correlation/off track I am considering here, *STEP 2* – thesis* would also be true. For its antecedent is true (there is no Milvian bridge), and so is its consequent, for it is true that, on this interpretation, had the

words, the concept of epistemic-reason Street is interested in discussing in her paper is that of (what I have called in Chapter 2 of this thesis) *an evidence-based reason*. But if this is the case, then *it cannot be conceivable* that epistemic-reason could have been based on considerations other than evidence.

¹⁸⁰ This follows – trivially, so to speak – from the non-existence of epistemic-normative facts of the kind at issue. But – as we saw from my discussion in the two preceding sections – the absence of a Milvian bridge would be established regardless, merely by relying on the kind of evolutionary explanation for our epistemic-normative beliefs which Street favours.

epistemic-normative facts been different, our epistemic-normative beliefs would have been the same).

Is establishing the lack of correlation between epistemic-normative beliefs and epistemic-normative facts *in this sense* sufficient to warrant the claim that our epistemic-normative beliefs are, then, not justified? In the previous chapter (Section 3.4), I have illustrated how the appeal to *conceptually possible* scenarios can be put to the service of defending counterfactuals which have, so to speak, ‘epistemological consequences’. There I gave the example of the child who guesses the truths of remote mathematical conjectures at random and in which her guesses turned out to be systematically correct (Clarke-Doane; 2012: 321). And I quoted Clarke-Doane commenting on this example as follows:

There is an obvious sense in which the forces generating the child’s beliefs in this scenario were “non-truth-tracking,” even given that her guesses were correct and that the relevant truths were metaphysically necessary. Even if it is not metaphysically possible that the relevant conjectures are false, it is intelligible to imagine that they are. And *if—for all that we can intelligibly imagine—the relevant conjectures had been false, the child’s beliefs would have been the same*. This raises a *prima facie* puzzle: what is the explanation for the child’s having so many true corresponding beliefs? (2012: 321, italics mine)

I have also indicated (Chapter 6, Section 3.2) how Street herself does explicitly resort to an appeal to what could *conceivably* have been the case, in articulating her Darwinian-epistemological argument against normative realism in the *practical* domain. In Street (2008) she rests her epistemological challenge on the claim that ‘as a purely *conceptual* matter [...] *normative truths might be anything*’ (2008: 208).

Perhaps then, we ought to read the dialectic between Street and the evidentialist-realist here along the lines of the situation illustrated by the example of the child I have discussed above. After all, we might think, it is true that *we can intelligibly imagine the epistemic-normative facts being different (in the sense I have indicated)* and that, in this case, *our epistemic-normative beliefs would have been the same*. And if we believe that the child’s beliefs in

that scenario are not justified, perhaps we ought to believe that epistemic-normative beliefs on an evidentialist-interpretation of their truth conditions are not justified either.

I think this latter claim can be resisted. That is to say, establishing that epistemic-normative beliefs are not correlated to epistemic-normative facts in this sense (Off-track – sense 1) does not suffice to secure the sceptical conclusion that our epistemic-normative beliefs are not justified. This is because there are important relevant differences between the child-scenario above and the situation we would be in if evidentialist-realism were true. Rather than discussing those, I will illustrate why the sceptical conclusion does not follow by means of another example.

It seems to me the dialectic between Street and the evidentialist-realist here can be exemplified by means of an analogy with the following situation. Suppose we are given a drug which affects the colour-beliefs we form. Suppose also that we are aware of this fact and we are told that the drug is such that – whatever the colour of the objects in the environment we are in may be – we will have a visual appearance as of those objects being red. As a consequence, we will form the beliefs that the objects in this environment are red. Thus *we would believe those objects are red whether or not they were red.*

Ought we to think that whatever colour-beliefs we form in these circumstances are not justified? Well, only if we can reasonably assume – as it is plausible to do in normal circumstances – that we do not move in an environment in which all of the objects we are likely to encounter are indeed red. For, suppose we moved in an environment in which all those objects are red. Then, despite the fact that the content of our colour-beliefs has been determined by the drug, the content of our colour-beliefs would still ‘match’ what is the case colour-wise. It would still be true that we would believe that objects are red whether or not they were red. Yet, if we know or have strong reasons to believe (say we are on the set of a movie...) that we move in an all-red environment, then the claim that our colour-beliefs are not justified would not seem to follow.

In the case I have just considered, *the mere possibility* that objects in our environment are not red does not suffice to establish the claim that our colour-beliefs are not justified. For this conclusion to follow, it must be the case that we know/have reason to believe that the environment we move in does not only (or even mainly) consist of red objects. Or we must have no good reason to believe that that we move in an environment in which all (or most) of the objects are red. In other words, the conditions we are in must be not only such that (1) we would have believed that objects are red whether or not they were red, but also such (that we can safely assume that) that (2) our colour beliefs could have easily been false. Because (we can safely assume that) it is the case that the colour of the objects in the environment we move in *might have been anything*. Thus our colour-beliefs could have easily been false. And knowing/justifiably believing all this would indeed seem to undermine justification for our colour-beliefs.

I believe that the situation with evidentialist-realism and a Darwinian-epistemological challenge that relied on *off-track - sense 1* is relevantly similar to that which I have illustrated above. That is to say, the mere *conceivability* of there being no epistemic-normative facts over and above evidence-facts does not suffice to establish the thesis that epistemic-normative beliefs (on an evidentialist-construal of their truth conditions) are not justified. For this sceptical conclusion to follow, it must not only be the case that *had the epistemic-normative facts been different, our epistemic-normative beliefs would have been the same*, in the sense that it is conceivable to imagine that there are no evidentialist-realist normative facts and that we still hold the epistemic-normative beliefs we do.

For the conceivability of this scenario to have *epistemological import* – i.e. to bear on what is epistemically justified/unjustified – it seems to me, it would also have to be the case that we have no good reasons to think that epistemic-normative facts are such as the evidentialist-realist takes them to be; or that we have positive reasons to think that they are not as the evidentialist-realist takes them to be; or that we are in no position to determine the likelihood of there being facts of the kind at issue.

For suppose that epistemic-normative facts were as the realist takes them to be (or that we had reason to believe that they are so). Then, even if it were true that had there not been facts of this kind, we would still have had the epistemic-normative beliefs we have,¹⁸¹ the sceptical conclusion would not follow. This is because we would be in a scenario which is relevantly similar to that in which, despite the fact that we are bound to believe that the objects we see are red due to the effect of a drug, we also have (drug- independent) reasons to believe that we are in an environment in which all the objects are red.

For that conclusion to follow, it would have to be true that *it could easily have been the case that the epistemic-normative facts are not like what the evidentialist-realist takes them to be*. In the case under consideration, this would presumably have to amount to us either having good reason to believe that there are no epistemic-normative facts of the kind the evidentialist posits, or having no good reason to think that there are epistemic-normative facts of that kind. But this just amounts to saying that epistemic-normative beliefs are off-track with respect to epistemic-normative facts of the evidentialist kind *on a different interpretation of off-track/ lack of correlation*, i.e. what I have called *off-track – sense 2*.

In Section 3.4 of the previous chapter, I defined this sense of off-track/lack of correlation in terms of the combination of two theses. There I said that *Off-track – sense 2* amounted to the conjunction of:

CF²-E: Had the epistemic-normative facts been different, our epistemic- normative beliefs would have been the same.

and:

D'-E: the epistemic-normative facts might easily have been different.

¹⁸¹ In the sense I have been assuming all along in this discussion, i.e. that it is conceivable that there are no facts of this kind, etc.

In that section, I also discussed how, if we are to make any sense of D'-E at all, we ought to understand its truth to depend on what reasons we have for believing that epistemic-normative facts are like the epistemic-realist takes them to be. And this is also the result I derived from my discussion of the analogy with the colour-example above. In order for us to believe that our epistemic-normative beliefs are not justified (on an evidentialist understanding of their truth-conditions) it must not only be the case that we would still have held the epistemic-normative beliefs we hold, even if there were no epistemic-normative facts of the evidentialist kind. But we must also have positive reasons to believe *that there are no facts of this kind*.¹⁸²

If what I have said so far is correct, it shows that on *off-track – sense 1*, the fact that STEP 1 – thesis* and STEP 2 – thesis* are true, does not suffice to support the conclusion that our epistemic-normative beliefs are not justified (on an evidentialist-realist construal of the truth condition of those beliefs). It also illustrates that, in order for that conclusion to follow, we ought to establish that epistemic-normative beliefs and 'evidentialist' epistemic-normative facts are not correlated in the sense of 'lack of correlation' that *off-track – sense 2* expresses. Moreover, if an argument like Street's is to be successful, establishing that epistemic-normative beliefs are off-track in this sense ought to be the result of *evolutionary* considerations of the kind she appeals to. In other words, it ought to follow from the truth of her adaptive link account (AL).

However, Street's adaptive-link account can hardly be taken to support the lack of correlation between epistemic-normative beliefs and 'evidentialist' epistemic-normative facts *in this latter sense*. For establishing lack of correlation in this sense required establishing that we have good reasons to think that evidentialist-realism is false, or no good reasons to think it is true, or perhaps that we are in no position to determine whether or not we have reasons to believe that evidentialist-realism is true. But notice that (1) whether or not we have good reasons to think that evidentialist-realism is true is *independent* of evolutionary considerations of the kind Street

¹⁸² Or no reasons to believe that there are, or we must be in no position to determine whether there are epistemic-normative facts of this kind or not.

relies on; and that (2) surely, the evidentialist-realist will resist that we have good reasons to think that evidentialist-realism is false, or no reasons to believe it is true, or that we are in no position to adjudicate whether there are epistemic-normative fact of the evidentialist kind.

The evidentialist-realist will grant that it is surely *conceivable* that his view is mistaken. Yet he will also maintain that he has good reasons for holding this view. Those reasons will, for instance, have to do with the fact (that the realist believes) that alternative meta-normative accounts will fail to capture certain central aspects of the ‘epistemic-normative phenomena’; and that the view he holds can successfully deal with the most serious objections that have been directed against it in the philosophical literature. Can Street’s evolutionary considerations in any way be brought to bear on the ‘goodness’ of those reasons?

Notice how in order for *evolutionary* considerations to bear on the ‘goodness’ of those reasons, it would at least have to be the case that evolution can explain why realists hold the beliefs which comprise their *meta-normative* theory; or why they are inclined to take certain considerations as good reasons for holding this theory. But evolution in general and Street’s evolutionary considerations in particular can hardly be thought to provide such explanations. The goodness of the reasons for holding evidentialism will then seem to be largely *independent* of evolutionary considerations. And the ‘score’ between realist and anti-realists will have to be settled with more traditional means and by going back to more familiar debates.

To conclude: Street’s Darwinian-epistemological argument against evidentialist-realism does not succeed. In the sense (off-track – sense 1) in which Street’s considerations can be taken to establish that epistemic-normative beliefs are not correlated with ‘evidentialist’ epistemic-normative facts, these considerations do not suffice to establish the thesis that epistemic-normative beliefs are unjustified (on an evidentialist-realist construal of their truth conditions). And in the sense of lack of correlation which we would need in order to establish that sceptical conclusion (off-track – sense 2) Street’s evolutionary considerations cannot establish that

epistemic-normative beliefs are off-track. Whether this is the case would seem to depend on the settling of more traditional – and non-evolutionary – meta-normative debates and disagreements.

4. Street's Epistemological Argument against Evidence-as-Reasons Realism.

Street believes that what I have called 'evidence-as-reasons' realism is unacceptable 'for the same reason that realism about Cliffordian normativity [i.e. 'evidentialist' realism] is unacceptable' (2009: 238).

A view of this kind, as I have already discussed, holds that epistemic-reason facts are identical with evidence-facts. In Chapter 3, we saw how evidence-as-reason realism is compatible with very different accounts of what epistemic-reason facts and, hence, evidence-facts are like; for instance with the view that evidence facts (and hence epistemic-reason facts) reduce to probability facts of some specific kind; or with the view that evidence-facts are irreducibly normative facts.

As I have also already argued in Chapter 3, it seems to me that evidence-as-reasons realism is compatible with an evidentialist conceptual analysis of epistemic-reason claims. Street, however, does not consider this possibility. The variety of evidence-as-reasons realism she focuses on is also committed to an evidence-as-reasons conceptual analysis of epistemic reason claims.¹⁸³ Thus Street takes evidence-as-reasons realists to think that claims about what is evidence for what – such as *that E is evidence for P* – are 'roughly equivalent' (i.e. equivalent in meaning) to claims about reasons – such as *that E is a reason for believing that P* (2009: 237). And they believe that, when those claims are true, they are made true by the same kind of facts. We can call those equally *evidence-facts* or *epistemic-reason facts*.

¹⁸³ I will not consider this option myself in what follows for the sake of simplicity.

At a very early stage of her argument Street says something which appears to be very puzzling. Unless we get clear about what she says here, it will be hard to appreciate in full the nature of her challenge to realism.

Here is what she says:

It's worth emphasising that I have no quarrel with a normative interpretation of [evidence] claims *per se*, just as I have no quarrel with realism about [evidence] *per se*. My quarrel is with the *combination* of the two positions – in other words, normative realism about [evidence] – which not only interpret claims such as [E is evidence for P] along the lines of [E is a reason for P], but then goes on to insist that [claims such as E is a reason for P] hold [i.e. are true] in a way that is robustly independent of an agent's attitudes. (2009: 238)

It is this latter view – the one she has 'quarrels with' – that Street says 'is unacceptable'. But how must we make sense of what Street says here? The 'puzzle' I referred to above can be expressed as follows. If evidence-claims and epistemic-reasons claims are (roughly) equivalent in meaning then, presumably, if they refer at all, they refer *to the same facts*. Street has no quarrel with a normative interpretation of evidence claims. Neither has she any quarrel with 'realism about evidence' – i.e., as she puts it, the view that evidence-facts are 'robustly independent of an agent's attitudes'.

But then how can she have quarrels with the *combination* of these views? Clearly, there is an obvious candidate for the reference of these claims. Both epistemic-reason claims and evidence-claims (which, remember, are *roughly equivalent in meaning*) can be taken to refer to evidence-facts understood in the robust realist (i.e. attitude-independent) way Street has no problems with. If this is the case, then epistemic-reason claims can easily be taken to refer to 'attitude-independent facts' (i.e. evidence-facts). But isn't this precisely what the kind of epistemic-realists she is taking issues with believe?

I think what Street *really* means to take issue with here is the thesis that evidence-facts are *irreducibly normative* facts.¹⁸⁴ Indeed, what she goes on to say later in her argument supports this interpretation (2009: 239-241). Thus the ‘target’ of Street’s argument here is best viewed as being evidentialist realist views that are also committed to the thesis that evidence-facts are *irreducibly normative*, and that our epistemic-normative beliefs – when true – are made true by *irreducibly normative facts*.

But now remember that Street’s explicit target in the paper I am discussing is a family of views – epistemic realist views – which she defines in terms of their being committed to the existence of *attitude-independent* normative facts. Thus even if she were right to think that evidence-as-reasons realist views which are committed to what I have called the ‘irreducibility thesis’ face unacceptable epistemological problems, her argument would still fall short of establishing that (all evidence-as-reasons) *epistemic realist* views are unacceptable. For not all versions of epistemic-realism – included those of the evidence-as-reasons kind – need be committed to the *irreducibility thesis*.¹⁸⁵ In the best case scenario (i.e. that in which her argument is indeed a compelling one), Street’s case against epistemic-normative realism here would then still be incomplete.

With these important qualifications in mind, we are now ready to look at the details of Street’s argument (2009: 237-242). Street starts from the observation that most creatures like us we observe, ourselves included, have the widespread tendency to form very similar epistemic-normative beliefs and, as a result, hold many epistemic-normative beliefs with the same content.

On the conceptual interpretation of those beliefs under consideration here (i.e. the evidence-as-reason interpretation), these are beliefs of the form: ‘That E *is evidence* for P’. And they are equivalent to beliefs of the form: ‘That E *is a reason for believing* that P’. Specific instances of

¹⁸⁴ See Chapter 3 for a more detailed discussion of what a position of this kind amounts to in the specific case under discussion here.

¹⁸⁵ Again, see Chapter 3 for discussion of what this *irreducibility thesis* amounts to.

beliefs of this kind would have to include beliefs such as (some of these are Street's own examples).¹⁸⁶

- a) The fact that all previously encountered tigers were carnivorous *provides evidence for the truth of the claim/is a reason for believing that* the next one will be too.
- b) The fact that the past ten tosses of a fair coin came up heads *provides evidence for the truth of the claim/is a reason for believing that* the next toss will come up tails.
- c) The fact that n/m of the observed Xs have been Y *is evidence for the truth of the claim/is a reason to believe that* n/m of the unobserved Xs are Y.
- d) The fact that all the lights are off and the car is not parked outside *is evidence for the truth of the claim/is a reason for believing that* Jane is not at home.
- e) The fact that she is a Pisces *provides no evidence for the truth of the claim/is no reason to believe that* she will hit it off with him.

Once again, Street asks why 'as a matter of sheer scientific explanation' it is the case that creatures like ourselves come to hold (and share) many of the specific normative beliefs of this kind we hold, rather than, for instance, beliefs such as:

- f) The fact that all previously observed tigers were carnivorous is evidence for the truth of the claim that/*is a reason to believe that* the next tiger will *not* be carnivorous.

or

- g) The fact that n/m of the observed Xs have been Y is evidence for the truth of the claim/*is a reason to believe that* $1 \text{ minus } n/m$ of the unobserved Xs are Ys.¹⁸⁷

¹⁸⁶ With the difference that I have replaced Street's talk of 'likelihood' with talk of 'evidence'.

¹⁸⁷ These examples are Street's.

Again, Street thinks it is very plausible to believe that the scientific explanation in question will be an *evolutionary explanation*, presumably¹⁸⁸ because she believes that creatures who exhibited the mechanisms and sets of dispositions that gave origin to evidence/normative beliefs such as those in the first list had much better chances at surviving and reproducing than creatures who did not. Her thinking seems to be that on the assumptions that beliefs about evidence/reasons play the causal role they normally do play with respect to other beliefs (i.e. the belief that E is strong or decisive evidence/a reason for P causes creatures to believe that P), creatures who exhibited the mechanisms and set of dispositions which resulted in beliefs such as (a) – (e), etc. were more likely to end up with true beliefs than creatures who exhibited the mechanisms and set of dispositions that resulted in beliefs such as (f) or (g), etc. And ‘it tends to promote survival and reproduction to have roughly true beliefs about one’s environment’ (2009: 235).

Street explicitly indicates that her argument against ‘evidence-as-reasons realism’ is to be read along the lines of the argument she articulates against *evidentialist* forms of realism. This latter argument was an argument for the claim that evidentialist realism has ‘unacceptable epistemological consequences’. For, according to Street, it cannot provide an (acceptable) explanation of why – given the evolutionary explanations available for our epistemic-normative beliefs – many of our epistemic-normative beliefs are correlated with the facts which would make them true. Keeping this in mind and drawing on what Street actually says in the sections in which she discusses these matters, she once again seems to be putting forward an argument with the following structure:

(1) The *only* (evolutionary/scientific) explanation that the epistemic-realist can provide *of the correlation* between epistemic-normative beliefs and epistemic-normative facts is a version of the tracking account (TA*).

¹⁸⁸ It seems to me there are a lot of crucial and relevant details missing from the very sketchy evolutionary explanations Street is helping herself to here. However, I do not have the space to discuss this question here.

(2) There is a much better (evolutionary/scientific) explanation – i.e. the adaptive link account (AL*) – for *why we hold the epistemic-normative beliefs* we hold.

(3) The adaptive link account (AL*) does not at any point require mention of epistemic-normative facts.

Therefore:

(4) Tracking account (TR*) is false.

Therefore:

(5) The epistemic-realist can't provide *any* (acceptable) explanation of the correlation between epistemic-normative beliefs and epistemic-normative facts.¹⁸⁹

The difference here is that what epistemic realists are asked to explain – i.e. the *explanandum* of the evolutionary explanation which is requested of them – is different. In the case under consideration, they are asked to explain why epistemic-normative beliefs *understood along the lines of the evidence-as-reasons interpretation of epistemic-reasons claims* would 'line up' with epistemic normative facts *understood as being facts of the kind evidence-as-reason realism posits*. If we add the qualification about *irreducibility* which – as I have argued – we must add, then the facts with which epistemic-normative beliefs will have to 'line up' must be taken to be *irreducibly normative* facts.

Insofar as this argument has the same structure of the argument I have discussed in the previous section, it seems to me it will inherit exactly the same structural problems that I have already discussed there. But what is of specific interest here is what Street has to say about what the tracking account (TA*) and the adaptive-link account (AL*) come down to – and how they differ – in this case.

Here is what she says about the tracking account (AL*):

According to the *tracking account* as it arises in this context [...], creatures who (in some 'proto' fashion) accepted normative judgment (c) did better

¹⁸⁹ Street seems to believe that the 'tracking account' would be no better than saying that 'evolution landed us on the independent normative truths about such matters purely as a matter of chance' (2009: 234).

than creatures who accepted normative judgment (g) because the former judgment is true and it pays to grasp the normative truth. In other words, on this account, it is *true* that ‘the fact that n/m of the observed Xs has been Y is a normative reason to believe that n/m of the unobserved Xs are Y’, and creatures who made this judgment did better because they were grasping (if only at some ‘proto’ level) a normative fact that the others were missing. (2009: 240)

And here is what she contrasts the ‘tracking account’ with:

According to [the *adaptive link*] account, creatures who made the normative judgment expressing the inductive rule did better than creatures who made the normative judgments expressing the rule of counter-induction, not because members of the former group grasped an independent normative truth that others missed, but rather, more simply, because they tended to form true beliefs where the others didn’t (which in turn promoted reproductive success). Here it is useful to consider a specific example. Why did creatures who took the fact that *all previously encountered tigers were carnivorous* to be a reason to believe that *the next tiger would be carnivorous* do better than those who took the same fact to be a reason to believe that the next tiger would not be carnivorous? According to the adaptive-link account, the answer is not that the one group was sensitive to a normative fact that the other group missed; it is rather that creatures who made the former normative judgment tended to believe that the next tiger would be carnivorous, and, as it turned out, the next tiger *was* carnivorous. (2009: 240)

Street thinks that it is ‘absolutely clear’ that the adaptive link account wins out over the tracking account: it provides a much more ‘parsimonious’ explanation and it does not need to posit facts (i.e., presumably, *irreducibly normative facts*) which do no work in the explanation at issue and are, because of this, ‘unnecessary’ (2009: 241). And she concludes that, once the ‘tracking account’ has been so defeated, the realist will be left with no good explanation of the ‘line up’ between normative beliefs and normative truths (2009: 242). The only option for the realist, she thinks, would be to insist that it ‘all happened as a matter of sheer luck’.

As I said, Street’s argument here inherits the problems of her argument against evidentialist realism. For instance, it simply does not follow from the fact that the adaptive link account (AL*) provides *a better (evolutionary) explanation* of why we hold the epistemic-normative beliefs we hold that the epistemic realist can provide *no (evolutionary) explanation* of the correlation between

epistemic-normative beliefs and epistemic-normative facts (short of appealing to a ‘lucky coincidence’).

For one, the two explanations have *different explananda*. Thus it is, at best, premature to compare them. Before we can do so, we must also ask: how is the adaptive link account going to help us make sense of the thesis that ‘many of our epistemic-normative beliefs are (non-accidentally) true?’ After all, this is what Street originally challenges the epistemic realist to explain. But most importantly, why or how exactly are anti-realist theories of epistemic normativity such as the one Street herself favours going to be superior to the realist alternative here?¹⁹⁰

In any case, even if the two accounts – the realist/tracking account and the anti-realist/adaptive-link account – turned out to be so that the latter is (much) better than the former, Street’s conclusion would *still* not follow. For from A is a (*much*) *better explanation* than B, it does not follow that B is *no explanation*. All that Street could conclude on this basis is that the realist does not have *as good an explanation* of the correlation between epistemic-normative beliefs and epistemic-normative facts as the one the anti-realist can provide.

4.1. A Better Argument

Once again, it seems to me that there is a quick way to turn Street’s argument into a better argument here, in particular into an argument that – if successful – would indeed seem to have better chances to yield the conclusion that she seeks to establish. Understanding Street’s argument this way would bring her argumentative strategy against epistemic realism once again more in line with the evolutionary-epistemological strategy she adopts against realism about practical normativity (Street: 2006).

¹⁹⁰ Street writes in many points as if she was arguing for an ‘error-theoretic’ account of epistemic-normativity, rather than for the account she explicitly favours, i.e. a form of ‘Humean-constructivism’ about epistemic reasons. This, I believe, renders the question: ‘how is her favoured account going to be a better fit with the adaptive-link account she defends?’ much more urgent for Street.

This argument will have the following form:

(1) Evolution explains our epistemic-normative beliefs.

(2) The best evolutionary account of why we hold the epistemic-normative beliefs we hold (i.e. the ‘adaptive-link account’ AL*) explains those beliefs in terms of mechanisms M* the natural function – N* – of which *is not* that of correctly representing epistemic-normative facts (of the evidence-as-reasons kind). Neither is the ability to correctly represent epistemic-normative facts *required by* or a *by-product of* those mechanisms M* performing N*.

Therefore:

(3) There is *no Milvian bridge* between epistemic-normative beliefs and epistemic-normative facts.

Add that:

(4) If there is no Milvian bridge between epistemic-normative beliefs and epistemic-normative facts, then epistemic-normative beliefs are off-track with respect to epistemic-normative facts.

We can then conclude that:

(5) Epistemic-normative beliefs are off-track with respect to the kind of facts that – according to Evidence-as-reason realism – would make them true.

And if we assume that:

(6) Evidence-as-reason realism is true.

We can conclude that:

(7) Epistemic-normative beliefs are not justified.

Again, it is easy enough to recognise in the argument I have just outlined a slightly more complicated version of DEA. And again leaving these complications aside, we can loosely speak of the argument above as an instance of DEA.

The obvious question we must ask here is in what way – if in any way at all – what Street says about the adaptive link account (AL*) can be thought to support (2). Let us then examine this question.

Street's adaptive link account would seem to entail that mechanisms M^* do need to be such as to correctly represent (or at least be sensitive to) facts of *some* kind in order for those mechanisms to be able give rise to (or realise) N^* .¹⁹¹ This can be illustrated as follows. Let us think – as it is quite natural to do – of mechanisms M^* as realising a function which pairs inputs and outputs in such a way that the outputs are, in a large enough number of cases, true beliefs (or accurate belief-like representations). The 'adaptive-link' mechanisms Street has in mind must be, for instance, such as to prompt or produce the belief – or belief-like representation – *that the next tiger will be carnivorous* (as output) upon registering the fact *that all of the tigers which have hitherto been observed have been carnivorous* (as input).

The explanation of why mechanisms which rely on input-output pairings of such kind are likely to result in a large number of true beliefs must be that they – in one way or another – encode information to the extent that the facts they register (as input) and the content of the beliefs they produce (as output) are related in such a way that the first *evidentially support* the (truth of the) second. In the same way as the fact that all of the previously observed tigers have been carnivorous *evidentially supports* the truth of the proposition that the next tiger will be carnivorous. In other words, those mechanisms must be 'sensitive to' what I have called *evidence-facts*. Another way to put this is by saying that – in one way or another – those mechanisms must 'correctly represent' evidence-facts.

¹⁹¹ Let us assume that N^* here is, once again, the disposition to form true beliefs about one's environment. It is in virtue of giving rise to this ability that mechanisms M^* are passed on to the next generations; and it is in virtue of this fact that their frequency in a population of creatures like ourselves is to be explained. I am once again working under the assumption that the ability to form largely true beliefs about one's environment enhances the survival and reproduction prospects of the creatures who exhibit this ability.

It is *in virtue of* their ‘sensitivity’ to facts of this kind that they can produce largely true beliefs and, thus, perform their function N^* . And it is because of this that creatures who exhibited the mechanisms in question ‘tended to form true beliefs where [other creatures] didn’t’.

Street believes¹⁹² that the adaptive-link account can be invoked to explain why we hold (many, or a large number of) the epistemic-normative beliefs we hold. And many of the epistemic-normative beliefs we hold have been formed in environments which are – in many obvious respects – very different and very distant from the environments in which our ancestors evolved. Call these environments E' . Such beliefs will include beliefs about having reasons to believe that the microwave is broken, or beliefs about the kinds of reasons we have to believe in complex mathematical theorems, or about what reasons there are in favour or against complex scientific theories, etc.

Assume then that epistemic-normative beliefs of those kinds are still to be explained in terms of mechanisms M^* . Let us also assume that we must trust M^* to still result in the production of largely true beliefs as outputs in the environments E' in which we come to hold epistemic-normative beliefs of those kinds. If this is so, then we must also think that the disposition to correctly represent evidence-facts in what are, presumably, very different environments from the ones in which they originally evolved is *a by-product* of the operations of M^* .

If what I have said so far is correct, then we can say that:

The best evolutionary account of why we hold the epistemic-normative beliefs we hold (i.e. the ‘adaptive-link account’ AL^*) explains those beliefs in terms of mechanisms M^* the natural function – N^* – of which both requires (i.e. entails) the ability to correctly represent evidence-facts (in environments E) and has as a by-product the ability to correctly represent evidence-facts (in environments E').

¹⁹² This follows from what Street says in her paper, but also from the abstract form that evolutionary-epistemological challenges such as the one I am discussing here must take. See my discussion (in the previous chapter) of the ‘constraints’ the evolutionary explanations at work in arguments like DEA must fulfil.

And thus that:

(MB) If AL* is true, there is a Milvian Bridge between *epistemic-reason beliefs* and *evidence-facts*.

The interpretation of the consequences of the adaptive-link account I have just provided, is corroborated by what Street herself says:

In making judgments about epistemic reasons, no doubt we have a strong sense that we are tracking facts that are robustly independent of us. And indeed we are, on my view too. (2009: 243)

A question we have already encountered now arises again. As we saw, evidence-as-reasons realists believe that claims about evidence (E is evidence for P) and claims about epistemic reasons (E is a reason to believe that P) are roughly equivalent in meaning. We also saw that Street has no qualms with the thesis that evidence-facts are attitude-independent facts. Indeed, she herself takes them to be so. We also just now saw that – if Street is right – then the best evolutionary account of why we hold the epistemic-normative beliefs we hold is such as to ensure that there is (what I have called) a Milvian Bridge between epistemic-reason beliefs and evidence-facts.

But then there would then seem to be a perfectly respectable form of evidence-as-reasons realism that can provide a perfectly respectable evolutionary explanation for the correlation between epistemic-normative beliefs and epistemic-normative facts. This is the form of evidence-as-reasons realism which believes that epistemic-normative facts are identical with the (attitude-independent) evidence-facts which mechanisms M* enable us to ‘track’.

How can we then still believe that there is a challenge for epistemic realism to be raised here on the basis of what Street says? As I have already pointed out, Street’s argument against epistemic realism is best read as an argument against versions of this view which are also committed to the *irreducibility* of epistemic-normative facts.

Street believes that the *evidence-facts* which mechanisms M^* allow us to be able to track are not *irreducibly normative facts*. Indeed, she seems to believe that what I have been calling evidence-facts are ultimately to be understood in terms of ‘probability-facts’ of some form or other, and that facts about probability, in turn, are to be understood in terms of facts about ‘patterns in the world’ (2009: 241) or about ‘frequencies of events’ (2009: 243).

Let us assume – for the time being - that Street is right here and that some *non-irreducibly normative* account of evidence-facts can be provided. If this is so, then it would appear to be the case that the truth of (2) will be restored. For mechanisms M^* will enable us to track evidence-facts, but we would have no ‘evolutionary reason’ to think that they will also enable us to track *irreducibly normative facts*. Indeed, we can explain why we hold the epistemic-normative beliefs we hold in terms of dispositions and mechanisms which at no point entail that those beliefs must, ultimately, be correlated to irreducibly normative facts of any kind. And if (2) is true, then it will be true that:

There is no Milvian bridge between epistemic-normative beliefs and (irreducibly normative) epistemic-normative facts.

An argument which relied on considerations of this kind would appear to have (some) intuitive grip. For we would now seem to have a case for the claim that we would have held the epistemic-normative beliefs we hold *whether or not there had been irreducibly normative facts to make them true*. And this, we might think, lends support to the thesis that epistemic-normative beliefs are *not correlated* with¹⁹³ the kind of facts which – according to evidence-as-reason realists of the variety

¹⁹³ All I am relying on at this stage is whatever intuitive sense of ‘lack of correlation’, if there is any, one might take to be such that (1) it is supported by evolutionary considerations of the kind I have just reviewed and (2) it provides strong reasons in support of the thesis that our epistemic-normative beliefs are unjustified (on the construal of their truth conditions which will be provided by an evidence-as-reason realist who also subscribes to the irreducibility of epistemic-normative facts). Ultimately, as I will go on to argue, I don’t believe that Street’s evolutionary considerations are successful. Thus, I do not believe that there is a sense of lack of correlation which fulfils (1) and

we are considering – would make them true. We might then go on to think that this fact, in turn, is sufficient to establish that – if epistemic-realism of the variety under consideration is true – epistemic-normative beliefs are unjustified.

4.2. An Evidence-As-Reasons Realist Reply to the Argument

4.2.1. The Point So Far

How compelling is Street's Darwinian-epistemological argument against evidence-as-reasons epistemic realism? I have already pointed out how there are versions of evidence-as-reasons realism – those according to which evidence-as-reasons facts are not *irreducibly normative* - which far from being called into question by Street's evolutionary considerations, would rather seem to be *vindicated* by them.

In the sense that, on the assumption that her adaptive link account (AL*) is correct, it would itself appear to provide a perfectly respectable explanation of the correlation between epistemic-normative beliefs and evidence-as-reasons normative facts of the kind in question. And it would be so in a way that ought to satisfy Street's own standards for believing that evolutionary considerations can secure a correlation between beliefs of a certain category and the facts which make them true.

For if AL* is correct, then *there is* a rather straightforward Milvian bridge between epistemic-normative beliefs and *evidence-facts* of the kind Street believes we must be able to represent if we are to hold the epistemic-normative beliefs we hold. And on Street's preferred account of what those evidence-facts are like, those facts are (1) not irreducibly normative facts and (2) attitude-independent (on a sense of attitude-independence that includes the sense in which the evidence-as-reason realist believes the epistemic-normative facts are attitude-independent).

(2) above. At this stage, I am just attempting to provide as charitable a reading of Street's Darwinian-epistemological argument as possible.

But as I have also already pointed out, there is a perfectly respectable variety of evidence-as-reasons realism which says that epistemic-normative facts are *identical with* evidence-facts construed along the lines of Street's preferred account of what those facts amount to. And if establishing the presence of a Milvian bridge is sufficient to secure the correlation between a category of beliefs and the facts which make them true,¹⁹⁴ then we ought to believe that our epistemic-normative beliefs are correlated to epistemic-normative facts (on the version of evidence-as-reasons realism at issue). And if this is true, then at least some versions of evidence-as-reasons realism will clearly survive Street's Darwinian-epistemological challenge. Thus Street's argument against evidence-as-reasons realism is, at best, incomplete.

But what about versions of evidence-as-reasons realism which also subscribed to the thesis that epistemic-normative facts are *irreducibly normative*? At the end of the previous section, I have suggested that Street's argument was best read as a challenge for views of this kind. I have also said that when Street's argument is read with this target in mind, it would (perhaps) appear to have some intuitive force. This argument relied on the thesis that *the only kind of facts* that the evolved mechanisms M^* at work in Street's adaptive link account (AL^*) need to correctly represent are evidence-facts, where those latter facts are not taken to be irreducibly normative. To the extent that our epistemic-normative beliefs are to be ultimately explained in terms of M^* then, they would seem to be explained in terms of mechanisms which are not sensitive to – i.e. they do not correctly represent – irreducible normative facts. Thus one might be tempted to conclude that we would have had those beliefs *regardless of what the irreducible normative facts had been or whether there had been irreducible normative facts or not*. And this may, in turn, provide reasons for believing that on a construal of the truth conditions for those beliefs that required that they correctly represent irreducible normative facts, we ought to believe that those beliefs are not justified. I turn to look at this argument in the next section.

¹⁹⁴ That Street believes this is the case is apparent from what she says in Street (2009) but also in Street (2006) and Street (2008).

4.2.2. Evidence-as-Reasons Realism and Irreducibly Normative Facts

What are we to make of this latter argument? Let us look at it in more detail. As did her previous argument against evidentialist realism, this argument also relies on the two the crucial steps I have already illustrated. Let us call *irreducibly-normative evidence facts* the kind of facts the existence of which the version of evidence-as-reasons realism we are focusing on posits.¹⁹⁵ Taking into account my preceding discussion, the crucial steps of Street's argument can be formulated as follows:

*STEP 1 – thesis***: If (a) *adaptive-link account* (*AL**) explains our epistemic-normative beliefs, then (c) there is *no Milvian bridge* between epistemic-normative beliefs and *irreducibly-normative evidence facts*.

*STEP 2 – thesis***: If (c) there is *no Milvian bridge* between epistemic-normative beliefs and *irreducibly normative evidence facts*, then (b) there is *no correlation* between epistemic-normative beliefs and *irreducibly-normative evidence facts*.

Can Street establish the truth of each of those theses on the basis of her preferred evolutionary account of why we hold the epistemic-normative beliefs we hold? Does it matter which sense of 'off-track' we must take to be at work in them? And would establishing the truth of both of those theses suffice to warrant the claim that our epistemic-normative beliefs are not justified, if we took them to refer to irreducible normative facts?

I believe that we need not go past addressing the first of the questions I have just asked in order to see that Street's argument here cannot succeed. In other words, and to be more precise, I believe the problem with Street's argument here will have to do with STEP 1 – thesis**.

¹⁹⁵ This is because this brand of realism believes that (1) facts about evidence are identical to facts about epistemic reasons and (2) facts about evidence are irreducibly normative facts. See Chapter 3 for discussion of what this irreducibility claim amounts to.

Street's adaptive link account (AL*) must rely, quite crucially for her argument here, on the claim that evidence-facts are not *irreducibly normative* facts. As we saw, Street takes evidence-facts to be identical with – and ultimately to be understood in terms of – probability facts of some sort or other. And she seemed to favour either some 'propensity' or some 'frequency' account of probability. Street does not go into details as to which theory of probability she is relying on here. But what is clear is that she believes that whatever evidence-facts will turn out to be, they will not turn out to be *irreducibly normative facts*.

But one might wonder whether this is something that Street can just *assume* – rather than having to *argue for* – in the context of an argument such as the one she is putting forward. For that evidence-facts are to be understood in this way is something that the brand of evidence-as-reasons realists she is arguing against here will hardly concede.

Proponents of this version of epistemic realism – as Street herself characterises it – believe that evidence-reasons facts are identical with evidence-facts, and that these facts are *irreducibly normative* facts. And surely, epistemic-realists of this variety will believe they there are *strong arguments* for this latter thesis. Among the main motivations for holding a version of normative realism of the kind under discussion here would be – for instance – an argument to the extent that any account of evidence-facts which does not ultimately appeal to irreducible normative properties is bound to be incomplete and, because of this, an inadequate account of evidence facts themselves.

Thus it is open to an epistemic realist of this variety to agree with Street on the claim that evolution explains why we hold the epistemic-normative beliefs we hold. And he can also agree with the claim that – in the way Street's adaptive-link accounts describes – the epistemic-normative beliefs we hold originate in mechanisms which are such as to enable us to correctly represent evidence-facts. But – as I have pointed out above – he will most likely *disagree* on what the correct analysis of those evidence-facts ought to be. On this interpretation of Street's argument then, the bone of contention between Street and the evidentialist-realist will rather

come down to the correct understanding of evidence-facts. And settling this question one way or another *will not itself depend* on an appeal to evolutionary considerations of the kind Street builds on in her argument.

Let us now go back to the first crucial step of Street's argument. This was that:

*STEP 1 – thesis***: If (a) *adaptive-link account* (AL^*) explains our epistemic-normative beliefs, then (c) there is *no Milvian bridge* between epistemic-normative beliefs and *irreducibly-normative evidence facts*.

It seems to me that evidence-as reason realists of the kind under discussion here can reply to Street by resisting this step in either of two ways. We can put this reply in the form of a dilemma. Remember that this brand of epistemic realism believes that the best account (indeed, the *correct* account) of evidence-facts will have to be provided in terms of *irreducibly normative* facts. Let us assume that this is true. As we saw in my discussion in the previous section, it followed from the details of Street's discussion of her adaptive link account (AL^*) that:

(MB): If AL^* is true, there is a Milvian Bridge between *epistemic-reason beliefs* and *evidence-facts*.

Now, either we take the evidence-facts which the adaptive mechanisms M^* (appealed to in Street's adaptive link account AL^*) enable us to correctly represent to be *irreducible normative facts* or we don't.

If we do, then the realist can say that we have thereby established a straightforward Milvian bridge between epistemic-normative beliefs and irreducibly normative facts of the kind he posits. Hence Step 1 – thesis** will be false. And because this is a crucial step in Street's argument, that argument cannot ultimately succeed.

If we don't then, on the assumption that – as the realist believes – evidence-facts are *in fact* irreducibly normative, we will have a problem. For whatever Milvian bridge AL* will be taken to establish, it cannot be taken to be a Milvian bridge between epistemic-normative beliefs and evidence-facts at all. In other words, (MB) above will be false. One problem with this is that AL* was meant to provide an evolutionary explanation of our epistemic-normative beliefs on an evidence-as-reasons interpretation of those beliefs – that is to say on an interpretation of epistemic-normative beliefs according to which those beliefs were equivalent in meaning to beliefs about evidence. But if those beliefs are equivalent to beliefs about evidence and if the evolutionary explanation for them does not secure a Milvian bridge between those beliefs and (irreducibly normative) evidence-facts, then this explanation can hardly be considered by the realist a good (or the best) explanation for those beliefs. Thus it is open to the realist here to say that, because he has strong reasons to believe that evidence facts are irreducibly normative, and because Street's adaptive link account cannot explain how we can have (true) beliefs about those facts, then this account is a defective (evolutionary) account and it ought to be rejected. Hence the realist will say that proposition (a)¹⁹⁶ in Step 1 – thesis** is false. And because Street's strategy relies on the truth of this proposition, this strategy cannot ultimately be successful.

Lastly, and quite apart from the considerations I have just reviewed, notice also the following problematic aspect of Street's argument here. Remember that Street's overall Darwinian-epistemological strategy consisted in (1) *granting* the epistemic-realist – for the sake of argument – that facts of the kind they posit exist, and that they are the truth-makers of epistemic-normative beliefs and (2) showing how *assuming* the truth of epistemic-realism and acknowledging the influence of evolution on our epistemic-normative beliefs would ultimately lead to the conclusion that our epistemic-normative beliefs are not justified.

But if this is to be the overall strategy, then – on the specific instance of Street's argument which I am discussing here – it is not at all clear whether she is *granting* the epistemic realists what

¹⁹⁶ This was the proposition that (a) adaptive-link account (AL*) explains our epistemic-normative beliefs.

she ought to be granting them. For if Street is to grant the epistemic-realist under attack here anything at all, she ought to be *granting* that evidence-facts are *irreducibly normative facts*. Yet far from doing this, Street is rather *assuming* that evidence-facts are not of this kind.

The reply to Street's argument against evidence-as-reason realism which I have discussed so far can be summarised as follows. Either the evidence-as-reason realist has a rather straightforward way to answer the Darwinian-epistemological challenge that Street levels against this view. Or Street's argument must rely on substantive assumptions about the nature of evidence-facts that proponents of the version of realism under attack are likely to resist. In any case, settling the debate on the correct interpretation of evidence-facts one way or another is not something that can plausibly be the outcome of evolutionary considerations of the kind Street appeals to in her paper.

5. Conclusion

In this chapter I looked at a specific instance of DEA. This has allowed me to illustrate in some detail the kind of evolutionary explanations for epistemic-normative beliefs which proponents of DEA must draw on in order to argue for the argument's sceptical conclusion. I have focused on Street's version of this argument because it is one of the most detailed in the philosophical literature. It specifically targets epistemic-normative realism, it is sensitive to the diverse forms which a view of this kind can take, and it exemplifies rather clearly the general structure which an appeal to evolutionary consideration by proponents of DEA must take. I have then argued that Street's argument is not successful. I have done so in a way which showed how, despite the fact that all versions of epistemic realism I have considered (and which fall within the target of the argument) can address the challenge successfully, they will do so in significantly different ways.

In the previous chapters, I have drawn attention to some central features which arguments such as DEA must display. They had to do with the specific constraints on the kind of

evolutionary explanations which these arguments must rely on, and with the specific constraints on the kind of interpretation of off-track/lack of correlation which must be taken to be at work in those arguments.

Although I have not defended this claim directly here, it seems to me safe to think that – because of those general constraints on DEA – the epistemic realist strategies of reply I have explored in this chapter will generalise. That is to say, they will still be available to epistemic-realists who are faced with versions of DEA which may differ – in some of their details – from Street’s own; instance versions that rely on accounts of evidence which are different from that which Street relies on, or versions which ascribe to the evolved mechanisms which give rise to our epistemic-normative beliefs a ‘natural function’ different from that which Street ascribes to them (Gibbard: 1990).

Conclusion

In this thesis I explored one way in which the ‘facts of evolution’ have been brought to bear on the evaluation of competing meta-normative views in epistemology. More specifically, I investigated whether the truth of epistemic-realism is compatible with us retaining justification for our epistemic-normative beliefs, on the assumption that those beliefs can be explained in evolutionary terms.

I began by noticing how – despite there being general agreement on the claim that many of the central concepts in epistemology are *normative* – locating normativity in epistemology is far from an easy task. *Which* epistemic concepts are normative and *what* the normativity of epistemic concepts consists in are questions the answers to which can hardly be taken for granted. Yet we cannot have a sense of what normative realism in epistemology amounts to prior to settling those questions. In turn, we can hardly evaluate arguments against normative realism in epistemology – whether they be evolutionary or not – if we don’t know what this latter view amounts to. I suggested that we begin making progress on those issues by starting from the concept of an epistemic-reason; and I defended the view that a plausible characterization of reasons of this kind could be given in terms of the notion of an *evidence-based* reason. I then suggested that we go about identifying which epistemological concepts are normative on the basis of the relation they bear to the concept of an evidence-based reason. I argued that the thesis that epistemic-reasons are evidence-based reasons can be understood in importantly different ways. In order to elucidate this point, I distinguished three concepts of epistemic reasons by drawing attention to three different ways to interpret the relation between epistemic-reasons and evidence.

I then moved on to discuss *meta-normative* issues in epistemology and I focused my discussion on claims about epistemic-reasons. I showed how the different interpretations of epistemic

reasons that I distinguished can give rise to very different forms of epistemic realism. I thereby sketched a taxonomy of normative realist views in epistemology which took this fact into account.

Finally, I turned to discuss evolutionary arguments. I began by looking at the kinds of evolutionary arguments against normative realism which have received much attention in recent *metaethical* literature. I argued that arguments that took what I called an ‘epistemological’ form represent a more interesting object of philosophical investigation than evolutionary arguments of a different ‘metaphysical’ kind. I then outlined a Darwinian-epistemological argument which was meant to raise a challenge for *epistemic* realism (DEA) and discussed some central features of this argument at some length. In particular, I drew attention to: (1) the kind of constraints that the evolutionary explanations at work in this argument must answer to; (2) how the claim that epistemic-normative beliefs are off-track (if epistemic-realism is true) is best understood; and (3) which versions of epistemic-realism are meant to fall within the target of the argument. I then engaged with a specific instance of an argument of this form – the Darwinian-epistemological argument put forward by Sharon Street – and I defended different varieties of epistemic-realism against this argument. I said that the epistemic-realist strategies of response to Street’s argument that I articulated could be expected to generalise to other versions of DEA as well. The structure of my reply to DEA can be summarised as follows:

(a) The *conceptual links* between epistemic-reasons claims and evidence claims ensure that, if there are epistemic-normative facts, those facts are either identical to or strongly supervenient upon evidence-facts.

(b) The requirement that Darwinian-epistemological arguments be *selective* and that they do not entail troublesome forms of (global) scepticism ensure that whatever explanations of epistemic-normative beliefs evolution can be thought to provide, they will have to be such as to explain why – in normal circumstances – the epistemic-normative beliefs we hold true cause us

to have largely true beliefs about our most (and least) immediate environments. In order for those explanations to do so, they will have to say that whatever adaptive mechanisms led to our holding the normative beliefs we hold, these will have to be ‘appropriately sensitive’ to evidential-support relations and what I called ‘evidence facts’.

(c) Given (a) and (b) there is, in an intuitive sense, *an obvious correlation* between epistemic-normative beliefs and epistemic-normative facts of the kind the realist posits. This is because epistemic-normative beliefs are correlated with evidence facts and, epistemic-normative facts – on the versions of realism the argument targets – are either identical to or strongly supervenient upon those facts.

(d) On a less intuitive note, I said that on a (rather) charitable reading of the crucial epistemological premise of this argument (premise (4)), what evolutionary considerations would have to establish for DEA to call into question the justification for our epistemic-normative beliefs is that: (1) had the epistemic-normative facts been different, our epistemic-normative beliefs would have been the same and (2) the epistemic-normative facts could *easily* have been different, where the possibility at work in both claims is *conceptual* possibility (i.e. *conceivability*), and where the claim that epistemic-normative facts *could have been different* amounts to the claim that (it is conceivable that) there are no epistemic-normative facts of the kind realists posit. I argued that evolutionary considerations of the kind that proponents of DEA rely upon can hardly establish those claims. Deciding whether (1) and (2) are true will depend on the outcome of substantive debates in meta-epistemology which will mirror some of the traditional debates we are familiar with from the metaethical literature. And evolutionary considerations would appear to have *no bearing* on settling those debates one way or another.

(e) Given (c) and (d), DEA relies on a premise – the ‘epistemological’ premise (4) – which is false. Establishing this fact provides sufficient grounds for believing that DEA is not a successful argument.

Bibliography

- Achinstein, P. (1978). 'Concepts of Evidence', *Mind* 87/345 (*New Series*), 22–45.
- (ed.) (1983). *The Concept of Evidence*. Oxford: Oxford University.
- Alden Smith, E. (2000). 'Three Styles in the Evolutionary Analysis of Human Behaviour', in L. Cronk, N. Chagnon and W. Irons (eds), *Adaptation and Human Behaviour*. New York: Aldine DeGruyter.
- Alexander, R. (1987). *The Biology of Moral Systems*. New York: Aldine de Gruyter.
- Alston, W.P. (1988). *Epistemic Justification*. Ithaca, NY: Cornell University Press.
- Anscombe, G. E. M. (1963). *Intention*. Cambridge, MA: Harvard University Press.
- (2005). *Human Life, Action and Ethics: Essays by G. E. M. Anscombe*. Ed. by M. Geach and L. Gormally. Exeter, UK: Imprint Academic.
- Audi, R. (2001). 'An Internalist Theory of Normative Grounds', *Philosophical Topics* 29, 31–45.
- Austin, J. L. (1956–57). 'A Plea for Excuses.' *Proceedings of the Aristotelian Society*. *New Series*, 57, 1–30.
- Axelrod, R. (1984). *The Evolution of Cooperation*. New York: Basic Books.
- (2006). 'Biology to Ethics: An Evolutionist's View of Human Nature', in G. Boniolo and G. De Anna (eds) *Evolutionary Ethics and Contemporary Biology*. Cambridge: Cambridge University Press.
- Ayer, A. J. (1946). *Language, Truth and Logic*. 2nd edn (1st edn, 1936), London: Victor Gollancz.
- Bedke, M. (2009). 'Intuitive Non-Naturalism Meets Cosmic Coincidence', *Pacific Philosophical Quarterly* 90, 188–209.
- Berker, S. (2009). 'The Normative Insignificance of Neuroscience', *Philosophy and Public Affairs* 37, 293–329.
- (2013a). 'The Rejection of Epistemic Consequentialism', *Philosophical Issues* 23/1, 363–87.
- (2013b). 'Epistemic Teleology and the Separateness of Propositions', *Philosophical Review* 122/3, 337–93.
- (2014). 'Does Evolutionary Psychology Show that Normativity is Mind-Dependent?', in J. D'Arms and D. Jacobson (eds), *Moral Psychology and Human Agency: Essays on the New Science of Ethics*, Oxford: Oxford University Press.

- Blackburn, S. (1984). *Spreading the Word*. Oxford: Oxford University Press.
- (1993). *Essays in Quasi-Realism*. Oxford: Oxford University Press.
- (1998). *Ruling Passions*. Oxford: Oxford University Press.
- BonJour, L. (1999). 'Foundationalism and the External World', in J. Tomberlin (ed.), *Philosophical Perspectives* 13: Epistemology. Malden, MA: Blackwell.
- Boyd, R. (1988). 'How to Be a Moral Realist', in G. Sayre-McCord (ed.), *Essays on Moral Realism*. Ithaca: Cornell University Press.
- Bradie, M. and Harms, W. (2012). 'Evolutionary Epistemology', *The Stanford Encyclopedia of Philosophy*; <http://plato.stanford.edu/archives/win2012/entries/epistemology-evolutionary/>.
- Brink, D. O. (1984). 'Moral Realism and the Sceptical Arguments from Disagreement and Queerness', *Australasian Journal of Philosophy* 62/2, 111–25.
- (1986). 'Externalist Moral Realism', *Southern Journal of Philosophy* 24 (Supplement), 23–42.
- (1989). *Moral Realism and the Foundations of Ethics*. Cambridge: Cambridge University Press.
- (2001). 'Realism, Naturalism, and Moral Semantics', *Social Philosophy and Policy* 18, 154–76.
- Broome, J. (1999). 'Normative Requirements', *Ratio* 12/3, 398–419.
- (2004). 'Reasons', in R. J. Wallace, P. Pettit, S. Scheffler and M. Smith (eds), *Reasons and Value: Themes from the Moral Philosophy of Joseph Raz*. Oxford: Clarendon Press.
- (2005). 'Does Rationality Give Us Reasons?', *Philosophical Issues* 15, 321–37.
- (2008). 'Is Rationality Normative?', *Disputatio* 11, 153–71.
- Buller, D. (2005). *Evolutionary Psychology and the Persistent Quest for Human Nature*. Cambridge, MA: MIT Press.
- Buss, D.M. (2005). *The Handbook of Evolutionary Psychology*. Hoboken: Wiley.
- Campbell, D. T. (1974). 'Evolutionary Epistemology', in P. A. Schilpp (ed.), *The philosophy of Karl R. Popper*. LaSalle, IL: Open Court.
- Chignell, A. (2013). 'The Ethics of Belief', *The Stanford Encyclopedia of Philosophy*; <http://plato.stanford.edu/archives/spr2013/entries/ethics-belief/>.
- Child, B. (1993). 'Anomalism, Uncodifiability, and Psychophysical Relations', *Philosophical Review* 102/2, 215–45.
- Chrisman, M. (2010). 'Expressivism, Inferentialism and the Theory of Meaning', in M.S. Brady (ed.), *New Waves in Metaethics*. Basingstoke: Palgrave Macmillan.

- Churchland, P. (1987). 'Epistemology in the Age of Neuroscience', *Journal of Philosophy* 84/10, 544–53.
- Clarke-Doane, J. (2012). 'Morality and Mathematics: The Evolutionary Challenge', *Ethics* 122, 313–40.
- Clifford, W.K. (1877/1999). 'The ethics of belief', in T. Madigan (ed.), *The ethics of belief and other essays*, Amherst, MA: Prometheus, 70–96
- Cloos, C.M. (2011). 'On the Relationship between Reasons and Evidence', Unpublished Manuscript. <http://philpapers.org/rec/CLOOTR-2>
- Conee, E. and Feldman, R. (2004). *Evidentialism: Essays in Epistemology*. Oxford: Oxford University Press.
- Copp, D. (2008). 'Darwinian Skepticism about Moral Realism', *Philosophical Issues* 18, 186–206.
- (2011). 'Normativity and Reasons: Five Arguments from Parfit against Normative Naturalism', in S. Nuccetelli and G. Seay (eds), *Ethical Naturalism: Current Debates*. Cambridge: Cambridge University Press.
- (2011a). 'Realist Expressivism: a Neglected Option for Moral Realism', *Social Philosophy and Policy* 18/ 02, 1–43.
- Crisp, R.(2005). 'Value, Reasons and the Structure of Justification: How to Avoid Passing the Buck', *Analysis* 65, 80–5.
- (2006). *Reasons and the Good*. Oxford: Oxford University Press.
- Cuneo, T. (2007). *The Normative Web: An Argument for Moral Realism*. Oxford University Press.
- Dancy, J. (2000a). *Practical Reality*. Oxford: Oxford University Press.
- (2000b). 'Should We Pass the Buck?', in A. O'Hear (ed.), *Philosophy, the True, the Good and the Beautiful*, Cambridge: Cambridge University Press, 159–73
- (2004). *Ethics Without Principles*. Oxford: Clarendon Press.
- (2006). 'Nonnaturalism', in Copp, D. (ed.), *Oxford Handbook of Ethical Theory*. New York: Oxford University Press.
- (2009). 'Reasons and Rationality', in S. Robertson (ed.), *Spheres of Reason: New Essays in the Philosophy of Normativity*. Oxford: Oxford University Press.
- Danielsson, S. and Olson, J. (2007). 'Brentano and Buck-Passers', *Mind* 116, 511–22.
- Darwin, C. (1859). *The Origin of Species*. London: John Murray.
- (1871). *The Descent of Man, and Selection in Relation to Sex*. London: John Murray.
- Dawkins, R. (1982). *The Extended Phenotype*. New York: Oxford University Press.

- (1989). *The Selfish Gene*, 2nd edn, expanded, Oxford: Oxford University Press.
- Delbruck, D. (1978). 'Mind from Matter', in W. Heidcamp (ed.), *The Nature of Life*. Baltimore: University Park Press.
- DeRose, K. (2004). 'Sosa, Safety, Sensitivity, and Skeptical Hypotheses', in J. Greco, (ed.), *Ernest Sosa and His Critics*. Oxford: Blackwell.
- Enoch, D. (2007). 'An Outline of an Argument for Robust Metanormative Realism', *Oxford Studies in Metaethics* 2, 21-50.
- (2010). 'The Epistemological Challenge to Metanormative Realism: How Best to Understand It, and How to Cope with It. *Philosophical Studies* 148 (3): 413 - 438.
- (2011). *Taking Morality Seriously: A Defense of Robust Realism*. Oxford: Oxford University Press.
- FitzPatrick, W. (2008). 'Robust Ethical Realism, Non-Naturalism, and Normativity', in R. Shafer-Landau (ed.), *Oxford Studies in Metaethics* 3, 159-205.
- (2012). 'Ethical Non-Naturalism and Normative Properties', in M. Brady (ed.), *New Waves in Metaethics*. New York: Palgrave MacMillan.
- (2013). 'Debunking evolutionary debunking of ethical realism.', *Philosophical Studies*, 1–22.
- (2014). 'Morality and Evolutionary Biology', *The Stanford Encyclopedia of Philosophy*; <http://plato.stanford.edu/archives/fall2014/entries/morality-biology/>.
- (2014a). 'Why There is No Darwinian Dilemma for Ethical Realism', in M. Bergmann and P. Kain (eds.), *Challenges to Moral and Religious Belief: Disagreement and Evolution*, Oxford: Oxford University Press.
- Fodor, J. (1998). 'The Trouble with Psychological Darwinism', *London Review of Books* 20/2, 11–13.
- (2000). *The Mind Doesn't Work That Way: The Scope and Limits of Computational Psychology*. Cambridge, MA: MIT Press.
- (2002). 'Is Science Biologically Possible?', in J. Beilby (ed.), *Naturalism Defeated? Essays on Plantinga's Evolutionary Argument Against Naturalism*, Ithaca, NY: Cornell University Press.
- Frankish, K. (2007). 'Deciding to Believe Again', *Mind* 116/463, 523–47.
- Gibbard, A. (1990). *Wise Choices, Apt Feelings*. Oxford: Clarendon Press.
- (2003). *Thinking How To Live*. Cambridge, MA: Harvard University Press.
- Gould, S. J. and Lewontin, R.C. (1979). 'The Spandrels of San Marco and the Panglossian Paradigm: A Critique of the Adaptationist Programme', *Proceedings of the Royal Society: Biological Sciences* 205, 581–98.

- Gould, S. J. (1997a). 'Darwinian Fundamentalism', *New York Review of Books*, 44/10, 34–7.
- (1997b). 'Evolutionary Psychology: An Exchange', *New York Review of Books* 44/15, 55–8.
- Greco, J. (2002). 'Virtue and Luck, Epistemic and Otherwise', *Metaphilosophy* 34, 353–66.
- Greene, J. (2008). 'The Secret Joke of Kant's Soul', in W. Sinnott-Armstrong (ed.), *Moral Psychology 3. The Neuroscience of Morality: Emotion, Brain Disorders, and Development*. Cambridge, MA: MIT Press.
- Griffiths, P. and Wilkins, J. (2010). 'When Do Evolutionary Explanations of Belief Debunk Belief?', unpublished manuscript; <http://philsci-archive.pitt.edu/5314/>.
- Harman, G. (1977). *The Nature of Morality*. New York: Oxford University Press.
- (1986). 'Moral Explanations of Natural Facts: Can Moral Claims Be Tested Against Reality?', *Southern Journal of Philosophy* 24 (Supplement), 57–68.
- Hawthorne, J. (2004). *Knowledge and Lotteries*. Oxford: Oxford University Press.
- Hieronymi, P. (2005). 'The Wrong Kind of Reason', *The Journal of Philosophy* 102, 437–57.
- (2009). 'Believing at Will', *Canadian Journal of Philosophy* 35 (Supplement), 149–87.
- Horwich, P. (1982). *Probability and Evidence*. Cambridge: Cambridge University Press.
- Huemer, M. (2005). *Ethical Intuitionism*. Basingstoke: Palgrave Macmillan.
- Hull, D. (1988). *Science as a Process: An Evolutionary Account of the Social and Conceptual Development of Science*. Chicago: The University of Chicago Press.
- Jackson, F. (1999). 'Non-Cognitivism, Normativity, and Belief', *Ratio* XII, 420–35.
- James, W. (1896/1979). 'The Will to Believe', in F. H. Burkhardt, F. Bowers, and I. K. Skrupskelis (eds), *The Will to Believe and Other Essays in Popular Philosophy*. Cambridge: MA, Harvard.
- Joyce, J. (2005). 'How Probabilities Reflect Evidence', *Philosophical Perspectives* 19/1, 153–78.
- Joyce, R. (2002). *The Myth of Morality*. Cambridge: Cambridge University Press.
- (2006). *The Evolution of Morality*. Cambridge, MA: MIT Press.
- (2013). *Irrealism and the Genealogy of Morals*, *Ratio* 26/4, 351–72.
- Kahane, G. (2011). 'Evolutionary Debunking Arguments', *Nous* 45, 102–25.
- Kelly, T. (2003). 'Epistemic Rationality as Instrumental Rationality: A Critique', *Philosophy and Phenomenological Research* 66/3, 612–40.

- (2006). 'Evidence', *The Stanford Encyclopedia of Philosophy*; <http://plato.stanford.edu/entries/evidence/>.
- (2007). 'Evidence and Normativity: Reply to Leite', *Philosophy and Phenomenological Research* 75/2, 465–74.
- (2008). 'Evidence: Fundamental Concepts and the Phenomenal Conception', *Philosophy Compass* 3/5, 933–55.
- Kolodny, N. (2005). 'Why Be Rational?', *Mind* 114, 509–63.
- Koorsgaard, C. (1996). *The Sources of Normativity*. Cambridge: Cambridge University Press.
- (1997). 'The Normativity of Instrumental Reason', in G. Cullity and B. Gaut (eds), *Ethics and Practical Reason*, Oxford: Clarendon Press, 215–54.
- Korcz, K.A. (2010). 'The Epistemic Basing Relation', *The Stanford Encyclopedia of Philosophy*; <http://plato.stanford.edu/archives/spr2010/entries/basing-epistemic/>.
- Kvanvig, J. (2003). *The Value of Knowledge and the Pursuit of Understanding*. Cambridge: Cambridge University Press.
- Leite, A. (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, 456–64.
- Lenman, J. (2008). 'Review of Terence Cuneo: The Normative Web', *Notre Dame Philosophical Reviews*; <http://ndpr.nd.edu/news/23575/>.
- (2014). 'Moral Naturalism', *The Stanford Encyclopedia of Philosophy*; <http://plato.stanford.edu/archives/spr2014/entries/naturalism-moral/>.
- Levy, N. (2006). 'Cognitive Scientific Challenges to Morality', *Philosophical Psychology* 19, 567–87.
- Lillehammer, H. (2003). 'Debunking Morality: Evolutionary Naturalism and Moral Error Theory', *Biology and Philosophy* 18, 567–81.
- Lipton, P. (2004). *Inference to the Best Explanation*. London: Routledge.
- Mackie, J. L. (1977). *Ethics: Inventing Right and Wrong*. Harmondsworth: Penguin.
- McDowell, J. (1982). 'Criteria, Defeasibility, and Knowledge', *Proceedings of the British Academy* 68, 455–479.
- (1995). 'Knowledge and the Internal', *Philosophy and Phenomenological Research* 55, 877–93.
- McNaughton, D. and Rawling, P. (2003). 'Naturalism and Normativity', *Proceedings of the Aristotelian Society* 77/1(Supplement), 23–45.
- Millar, A. (2009). 'How Reasons for Action differ from Reasons for Beliefs', in S. Robertson (ed.), *Spheres of Reason: New Essays in the Philosophy of Normativity*. Oxford: Oxford University Press.

- Nagel, T. (1979). 'Ethics Without Biology', in T. Nagel (ed.), *Mortal Questions*, Cambridge: Cambridge University Press.
- (1997). *The Last Word*. Oxford: Oxford University Press.
- Nozick, R. (1981). *Philosophical Explanations*. Oxford: Oxford University Press.
- (1993). *The Nature of Rationality*. Princeton: Princeton University Press.
- (2001). *Invariances: The Structure of the Objective World* Cambridge, MA: Belknap Press.
- Oddie, G. (2005). *Value, Reality, and Desire*. Oxford: Oxford University Press.
- Olson, Jonas (2004). 'Buck-Passing and the Wrong Kind of Reasons', *Philosophical Quarterly* 54, 295–300.
- (2009). 'Reasons and the New Non-Naturalism', in S. Robertson (ed.), *Spheres of Reason: New Essays in the Philosophy of Normativity*. Oxford: Oxford University Press.
- Papineau, D. (2003). *The Roots of Reason: Philosophical Essays on Rationality, Evolution, and Probability*. Oxford University Press.
- Parfit, D. (2001). 'Reasons and Rationality', in D. Egonsson, J. Josefsson, B. Petersson, and T. Rønnow-Rasmussen (eds), *Exploring Practical Rationality*. Aldershot: Ashgate.
- (2011). *On What Matters*. Oxford: Oxford University Press.
- Pinker, S. (1997a). *How the Mind Works*. New York: Norton.
- (1997b). 'Evolutionary Psychology: An Exchange', *New York Review of Books* 44/15, 55–8.
- Plantinga, A. (1993). *Warrant and Proper Function*. Oxford: Oxford University Press.
- (2002). 'The Evolutionary Argument against Naturalism', in J. Beilby (ed.), *Naturalism defeated?* Ithaca, NY: Cornell.
- Popper, K. R. (1972). *Objective Knowledge: An Evolutionary Approach*. Oxford: Clarendon Press.
- Pritchard, D. H. and Turri, J. (2014). 'The Value of Knowledge', *The Stanford Encyclopedia of Philosophy*; <http://plato.stanford.edu/archives/spr2014/entries/knowledge-value/>.
- Pritchard, D. H. (2002). 'Resurrecting the Moorean Response to Scepticism', *International Journal of Philosophical Studies* 10, 283–307.
- (2005). *Epistemic Luck*. Oxford: Oxford University Press.
- (2008). 'Sensitivity, Safety, and Anti-Luck Epistemology', J. Greco (ed.), *The Oxford Handbook of Skepticism*. Oxford: Oxford University Press.
- Pryor, J. (2001). 'Highlights of Recent Epistemology', *British Journal for the Philosophy of Science* 52, 95–124.

- Putnam, H. (1983). 'Why Reason Can't Be Naturalized', reprinted in H. Putnam, *Realism and Reason: Philosophical Papers* 3. Cambridge: Cambridge University Press.
- (1990). 'Beyond the Fact-Value Dichotomy', in James Conant (ed.) *Realism with a Human Face*. Cambridge: Harvard University Press.
- Quine, W.V. (1951). 'Two Dogmas of Empiricism,' *Philosophical Review* 60, 20–43.
- (1963). *From A Logical Point of View*. New York: Harper & Row, Harper Torchbooks.
- Rabinowicz, W. and Rønnow-Rasmussen, T. (2004). 'The Strike of the Demon: On Fitting Pro-Attitudes and Value', *Ethics* 114, 391–423.
- Railton, P. (1986). 'Moral Realism', *The Philosophical Review* 95, 163–207.
- (1997). 'On the Hypothetical and Non-Hypothetical in Reasoning about Belief and Action', in G. Cullity and B. Gaut (eds), *Ethics and Practical Reason*. Oxford: Clarendon Press.
- Raz, J. (1999). *Practical Reasons and Norms*. Oxford: Oxford University Press.
- (1999a). 'Explaining Normativity: On Rationality and the Justification of Reason', in J. Raz, *Engaging Reason*, Oxford: Oxford University Press.
- Reisner, A. (2009). 'The Possibility of Pragmatic Reasons for Belief and the Wrong Kind of Reasons Problem', *Philosophical Studies* 145, 257–272.
- Ridge, M. (2014). 'Moral Non-Naturalism', *The Stanford Encyclopedia of Philosophy*, <http://plato.stanford.edu/archives/fall2014/entries/moral-non-naturalism/>.
- Robertson, S. (2009). 'Introduction: Normativity, Reasons, Rationality', in S. Robertson (ed.), *Spheres of Reason: New Essays in the Philosophy of Normativity*. Oxford: Oxford University Press.
- Ruse, M. (1986). *Taking Darwin Seriously*. Oxford: Blackwell.
- (2006). 'Is Darwinian Metaethics Possible (And If It Is, Is It Well Taken)?', in G. Boniolo and G. De Anna (eds), *Evolutionary Ethics and Contemporary Biology*.
- Sayre-McCord, G. (1986). 'The Many Moral Realisms', *Southern Journal of Philosophy* 4 (Supplement), 1–22.
- (1988). 'Introduction: The Many Moral Realisms', in Geoffrey Sayre-McCord (ed.), *Essays on Moral Realism*. Ithaca, NY: Cornell UP, 1988.
- Scanlon, T.M. (1998). *What We Owe To Each Other*. Cambridge, MA: Belknap Press of Harvard University Press.
- (2009a). 'Being Realistic about Reasons: Introduction' University of Oxford, First Locke Lecture; http://www.philosophy.ox.ac.uk/lectures/john_locke_lectures/past_lectures.

- (2009b). ‘Being Realistic about Reasons: Metaphysical Objections’ University of Oxford, Second Locke Lecture;
http://www.philosophy.ox.ac.uk/lectures/john_locke_lectures/past_lectures.
- (2009c). ‘Being Realistic about Reasons: Epistemological Worries’. University of Oxford, Fourth Locke Lecture;
http://www.philosophy.ox.ac.uk/lectures/john_locke_lectures/past_lectures.
- Schafer, K. (2010). ‘Evolution and Normative Skepticism’, *Australasian Journal of Philosophy* 88, 471–88.
- Schiffer, S. (2003). *The Things We Mean*. Oxford: Oxford University Press.
- Setiya, K. (2008). ‘Believing at Will’, *Midwest Studies in Philosophy* 32/1, 36–52.
- Shafer-Landau, R. (2003). *Moral Realism: A Defence*. Oxford: Oxford University Press.
- (2012). ‘Evolutionary Debunking, Moral Realism and Moral Knowledge’, *Journal of Ethics and Social Philosophy* 7/1, 1–37.
- Shah, N. and Velleman, D. (2005). ‘Doxastic Deliberation’, *The Philosophical Review* 114/4, 497–534.
- Shah, N. (2002). ‘Clearing Space for Doxastic Voluntarism’, *The Monist* 85, 436–45.
- (2003). ‘How Truth Governs Belief’, *Philosophical Review* 112, 447–82.
- Shroeder, M. (2012). ‘The Ubiquity of State-Given Reasons’, *Ethics* 122/3, 457–88.
- Sinnott-Armstrong, W. (2006). *Moral Skepticism*. Oxford: Oxford University Press.
- Skorupski, J. (1999). ‘Irrealist Cognitivism’, *Ratio* 12, 436–59.
- (2002). ‘The Ontology of Reasons’, *Topoi* 21, 113–24.
- (2009). ‘The Unity and Diversity of Reasons’, in S. Robertson (ed.), *Spheres of Reason: New Essays in the Philosophy of Normativity*. Oxford: Oxford University Press.
- Sober, E. (1981). ‘The Evolution of Rationality’. *Synthese* 46, 95–120.
- (1984). *The Nature of Selection: Evolutionary Theory in Philosophical Focus*. Cambridge, MA: MIT Press.
- (ed.). (1994). *Conceptual Issues in Evolutionary Ethics*, 2nd edn, Cambridge, MA: MIT Press.
- Sosa, E. (1999). ‘How to Defeat Opposition to Moore’, *Philosophical Perspectives* 13, 141–54.
- Southwood, N. and Chuard, P. (2009). ‘Epistemic Norms Without Voluntary Control’, *Nous* 43, 599–632.

- Stich, S. (1990). *The Fragmentation of Reason*. Cambridge, MA: MIT Press.
- Stratton-Lake, P. (2005). 'How to Deal with Evil Demons: Comment on Rabinowicz and Rønnow-Rasmussen', *Ethics* 115, 788–99.
- Street, S. (2006). 'A Darwinian Dilemma for Realist Theories of Value', *Philosophical Studies* 127, 109–66.
- (2008). 'Reply to Copp: Naturalism, Normativity, and the Varieties of Realism Worth Worrying About', *Philosophical Issues* 18, 207–28.
- (2008a). 'Constructivism about Reasons', in R. Shafer-Landau (ed.), *Oxford Studies in Metaethics*, vol. 3. Oxford: Clarendon Press.
- (2009). 'Evolution and the Normativity of Epistemic Reasons', *Canadian Journal of Philosophy* 39 (Supplement 1), 213–248.
- (2010). 'What Is Constructivism in Ethics and Metaethics?', *Philosophy Compass* 5/5, 363–84.
- Sturgeon, N. (1986). 'Harman on Moral Explanations of Natural Facts', *Southern Journal of Philosophy* 24 (Supplement), 69–78.
- (1988). 'Moral Explanations', in G. Sayre-McCord (ed.), *Essays on Moral Realism*, Ithaca: Cornell University Press.
- (1992). 'Nonmoral Explanations', *Philosophical Perspectives* 6, 97–117.
- (2006). 'Ethical Naturalism', in Copp, D. (ed.) *Oxford Handbook of Ethical Theory*. New York: Oxford University Press.
- Tersman, F. (2008). 'The Reliability of Moral Intuitions: A Challenge from Neuroscience', *Australasian Journal of Philosophy* 86, 389–405.
- Toulmin, S. (1967). 'The Evolutionary Development of Natural Science', *American Scientist* 55, 456–71.
- VanFraassen, B. (1980). *The Scientific Image*. Oxford: Oxford University Press.
- (1984). 'Belief and the Will', *Journal of Philosophy* 81, 235–56.
- (1989). *Laws and Symmetry*. Oxford: Clarendon Press.
- Velleman, D. (2000). 'On the Aim of Belief', in D. Velleman, *The Possibility of Practical Reason*. New York: Oxford University Press.
- (2000a). 'The Possibility of Practical Reason', reprinted in D. Velleman *The Possibility of Practical Reason*. Oxford: Oxford University Press.
- Walton, K. (1990). *Mimesis as Make-Believe*. Cambridge, MA: Cambridge University Press.
- (1993). 'Metaphor and Prop-oriented Make-Believe', *European Journal of Philosophy* 1, 39–56.

- Wedgwood, R. (2002). 'The Aim of Belief', *Philosophical Perspectives* 16, 267–97
- (2007). *The Nature of Normativity*. Oxford: Oxford University Press.
- White, R. (2010). 'You just believe that because ...', *Philosophical Perspectives* 24, 573–615.
- Wielenberg, E. (2010), 'On the Evolutionary Debunking of Morality', *Ethics* 120, 441–64.
- Williams, B. (1973). 'Deciding to Believe', in B. Williams, *Problems of the Self: Philosophical Papers, 1956–1972*. New York: Cambridge University Press.
- (1980). 'Internal and External Reasons', reprinted in B. Williams (1981), *Moral Luck*. Cambridge: Cambridge University Press.
- (1981). 'Internal and External Reasons', in B. Williams. *Moral Luck*. Cambridge: Cambridge University Press.
- (2001). 'Postscript. Some Further Notes on Internal and External Reasons', in E. Millgram (ed.), *Varieties of Practical Reasoning*. Cambridge, MA: MIT Press.
- Williamson, T. (1997). 'Knowledge as Evidence', *Mind* 106/424, 717–41.
- (2000). *Knowledge and Its Limits*. Oxford: Oxford University Press.
- Wright, C. (1992). *Truth and Objectivity*. Cambridge, MA: Harvard University Press.
- (1995). 'Truth in Ethics', *Ratio* 8: 209–26.
- Zagzebski, L. (1996). *Virtues of the Mind*. Cambridge: Cambridge University Press.
- (2003). 'The Search for the Source of Epistemic Good', *Metaphilosophy* 34, 12–28.