

Introduction

Change in Moral View: Higher-Order Evidence and Moral

Epistemology

Preprint

forthcoming in Higher-Order Evidence and Moral Epistemology, edited by

Michael Klenk (Routledge).

Author information

Michael Klenk
m.b.o.t.klenk@tudelft.nl
Delft University of Technology
ORCID 0000-0002-1483-0799

Abstract

Most epistemologists maintain that we are rationally required to believe what our evidence supports. Generally speaking, any factor that makes it more probable that a given state of affairs obtains (or does not obtain) is evidence (for that state of affairs). In line with this view, many metaethicists believe that we are rationally required to believe what's morally right and wrong based on what our moral evidence (e.g. our moral intuitions, along with descriptive information about the world) supports. However, sometimes we get information about our evidence, such as a theory that explains that all moral intuitions are ultimately caused by evolutionary forces. Such genealogical claims like this take form as a puzzle about how to rationally respond to higher-order evidence in moral epistemology. How should we change our moral views in response to genealogical claims about the evolutionary origin of our moral beliefs or about widespread moral disagreement? This introductory chapter first explains the issue about how to change our moral views based on an easily accessible example. Then it shows how recent debates about the puzzle of higher-order evidence bears on recent debates in moral epistemology, notably the debates about evolutionary debunking arguments in metaethics, the epistemic significance of moral peer disagreement, moral testimony, and collective moral knowledge before it introduces the chapters of this book.

1 Introduction

In this introductory chapter, I first explain the issues addressed in this book based on an example of everyday moral thought: how should we change our moral views, such as the conviction that eating meat is impermissible, in response to higher-order evidence – that is, evidence about our evidence? My aim is to make vivid how moral life is permeated with the issues discussed in this book. I then introduce the two philosophical debates that inform this question in Sections 3 and 4 and point out their connection in Section 5. My aim in these sections is to introduce and connect the central positions in the philosophical debate and to show how approaching moral epistemology from the higher-order-evidence perspective is expedient and fruitful. I then present the purpose and content of the volume in Section 6.1

2 Change in Moral View

A good friend of mine is a vegetarian. He believes that it is morally wrong to eat meat because it is wrong to harm sentient beings such as fish, chicken, pigs, and cows. He holds this moral conviction based on sound evidence: he knows that animals can feel pain, he has the intuition that harming sentient being is wrong, and the arguments in Peter Singer's *Animal Liberation* (1975) seem sound to him.²

Recently, my friend read in a newspaper that animal agriculture has a significant impact on climate change.³ He also thinks that we have a moral obligation to protect our planet, so the newspaper article is more evidence for him that eating meat is morally wrong: according to the article, eating meat perpetuates anthropogenic climate change, thereby harming our planet, and since

we have an obligation to protect our planet, we should not eat meat. Intuitively, my friend should therefore be more confident in his moral view about eating meat.

Many epistemologists embrace the view that we are rationally required to believe what our evidence supports. Generally speaking, any factor that makes it more probable that certain states of affairs obtain (or do not obtain), such as the impermissibility of eating meat, is evidence (that those states of affairs obtain). Insofar as there is a sense in which moral beliefs (or moral views – I use these terms interchangeably) can be true, we can thus speak of evidence in support of a moral view.⁴ What counts as evidence in support of a moral view, of course, depends on what makes the particular moral view likely to be true. Descriptive information (e.g. about the fact that animals feel pain) as well as (considered) moral intuitions and moral emotions are often taken as evidence in this context.⁵ Through the newspaper article, my friend encountered more evidence in favour of the wrongness of eating meat, and consequently, it would be rational for him to be more convinced of his moral view against eating meat. Evidence can thus provide more support and therefore make it rationally required to endorse one's moral views more strongly.

Evidence may also take away support for one's moral views, in which case one is rationally required to lower one's confidence in one's moral view. Suppose you do not agree with my friend and instead believe that eating meat is morally permissible. My friend might explain to you in gruesome detail how some animals die in factory farms, describing their shrieks and their anguish, and you might also then have the moral intuition that harming animals is wrong. Moral intuitions are an example of first-order evidence for our moral views.⁶ First-order

evidence bears directly on the truth of a proposition at the object level; that is, in our example, first-order evidence bears on how probable it is that eating meat is morally permissible. You would now have evidence against your moral view. You should weigh it against the relevant evidence you already possess, and then you should endorse what your evidence most strongly supports. Plausibly, your overall confidence that eating meat is permissible now has less evidential support, and consequently, you should now be less confident that eating meat is permissible. This process is simple.

However, the relationship between evidence and rational moral belief becomes more complicated when you consider new information about your evidence. For example, suppose my friend tells you that you just believe that eating animals is permissible because you grew up in a culture where almost everyone thinks this, so you were never confronted with good evidence to the contrary⁷ or that you just believe it because your naturally evolved moral intuitions have resulted in an inhibition against harming human animals, but not other animals.⁸ Alternatively, my friend might point out that you enjoy eating steaks and that therefore your subconsciousness manufactures an appropriate rationalisation so that you can keep doing what you enjoy.⁹ Or perhaps you are simply too callous and ignorant to grasp the impermissibility of eating meat. Or he might also emphasise that he, someone who has considered the question whether eating animals is morally permissible has carefully arrived at a different conclusion. Should you change your moral view?

These are claims about what we might call genealogies of your moral beliefs: they concern the circumstances in which you formed your moral beliefs (e.g. the

culture in which you grew up), the grounds on which you formed them (e.g. your desires or your ignorance), your way of interpreting the grounds (e.g. differently from someone else, such as my friend), or their causes (e.g. your naturally evolved intuitions).¹⁰ Genealogical claims may provide higher-order evidence – that is, evidence about the character of the evidence itself – or about your capacities and dispositions for responding rationally to your evidence (Feldman 2005; Christensen 2010; Lasonen-Aarnio 2014; Kelly 2005, 2010, 2016).¹¹ Moral philosophers have been concerned with such genealogical claims. In particular, many have sought to debunk moral beliefs – that is, to show that they are less justified, or not justified at all, in light of the genealogical information, as I will show in more detail in Section 4.

Although it is not immediately apparent how genealogical claims have a bearing on whether it is permissible to eat meat, claims of that kind strike many philosophers as relevant to your moral outlook (e.g. Street 2006; Joyce 2016b; see Wielenberg 2016 for an overview), and therefore, they are relevant to whether you should maintain your conviction about eating meat. But how precisely should you change your moral view in response to such claims? How does higher-order evidence matter in moral epistemology? The issue quickly gets complicated here.

First, we have what I will term the classification problem. The classification problem is the problem of finding out whether a given piece of information counts as evidence or not (and thus makes some states of affairs more or less probable). Some genealogical claims are relevant to what one is rationally required to believe about moral matters, and thus, they are higher-order evidence for moral truths. Such claims seem to affect the rationality or justification of one's beliefs at least

indirectly.¹² After all, it is evidence about how good your evidence is or how well you interpreted it, which is relevant to whether or not your belief is true, even though it does not directly concern whether some moral claim is true.¹³ But clearly not all genealogical claims are relevant to what one is required to believe. Unsupported confabulations do not matter. One must at least have sufficient reason to think that genealogical claims about a (set of) belief(s) could be true for them to affect what one is rationally required to change about the affected (set of) belief(s). For example, the claim that you never looked at good evidence about the permissibility of eating meat in the first place strikes me as relevant for what you are rationally required to believe about eating meat, but only insofar as the claim is (likely to be) accurate. If you are reasonably sure that the claim is false, then it should not play a role in how you change your moral view about eating meat. Moreover, even among genealogical claims that are likely to be true, only some of them should affect what we are rationally required to believe. For example, consider the claim that you formed your beliefs about the permissibility of eating meat on a Tuesday. That is a genealogical claim, and it may likely be true, but I imagine that you would not change your view on eating meat because of it – and you should not, because weekdays do not affect the truth of moral beliefs about the permissibility of eating meat. So when can we count genealogical information as higher-order evidence instead of as mere higher-order information? Intuitively, there needs to be some “connection” between the genealogical claim and the truth of the belief that it is about for the genealogical claim to count as evidence. The question, of course, is what that “connection” should be. These observations verge on the trivial. However, as I will show in Section 4, attempts to explicate the

nature of that connection in the moral case are controversial, and yet they can be understood as (inexplicit) attempts to solve the classification problem.

Second, the classification problem is closely related to what I will term the accounting problem, which concerns a more fundamental puzzle about the proper relation between higher- and lower-order evidence. As with the classification problem, it helps to illustrate the accounting problem with a non-moral case. Let us consider Watson and Holmes. Suppose Watson correctly reasoned from his available first-order evidence (e.g. a bloody knife) that the butler committed the murder. But Holmes, the best criminal investigator of all time, tells Watson that he reasoned irrationally. Watson now has both first-order and higher-order evidence, and he must take account of what his total evidence supports. So even though he may have solved the classification problem, the accounting problem still looms. Likewise in the moral case. Higher- and lower-order evidence add up to one's total evidence (e.g. your original evidence in support of the permissibility of eating meat plus the genealogical information provided by my friend). What view does one's total evidence support? That is, in forming a moral view, how should one account for one's total evidence? One option is to ignore the higher-order evidence. For example, you could maintain that eating meat is permissible and continue to believe that your view is based on good evidence. But in that situation, one would not only irrationally ignore a part of one's evidence (i.e. one's higher-order evidence), but this stance would also seem uncomfortably dogmatic, because one would not seem open to changing one's moral view. Alternatively, one could let one's higher-order evidence override one's first-order evidence. For example, you could be less confident that eating meat is permissible, and you could also be

less confident that your view is based on good evidence. But then you would just have irrationally ignored a different part of your evidence, and if your higher-order evidence is misleading, then your altered moral view would make you epistemically and morally worse off (because you gave up a rational belief in response to misleading higher-order evidence). Misleading higher-order evidence suggests that one's (moral) view fails to be rational, even though one's (moral) view is, in fact, rational. Finally, you could somehow try to weigh up your first-order and your higher-order evidence. If you do that, however, it would seem that you failed to fully "respect" both types of evidence (cf. Feldman 2005). So it seems as if there is no fully satisfactory way to change one's moral view in response to higher-order evidence and thus, in other words, no fully satisfactory way to solve the accounting problem.

This book looks at questions of moral epistemology through the lens of recent research on higher-order evidence. It thereby aims to break new ground in debates concerning debunking arguments in ethics (cf. Sauer 2018) and the epistemic significance of moral disagreement (cf. Tersman 2006) and to show how research on higher-order evidence has a bearing on debates about moral testimony, collective moral knowledge, and the truth of substantive moral theories such as consequentialism.

In the rest of the Introduction, I will first explain the accounting problem in more detail and will then reveal the prominent, though hitherto mostly unacknowledged, role of questions about the epistemic significance of higher-order evidence for debates in moral epistemology, before I summarise the insights provided by the 12 contributions to this volume.

3 The Accounting Problem: Puzzles About Higher-Order Evidence

The philosophical debate about higher-order evidence focuses predominantly on two questions, which will be relevant in the volume.¹⁴ First, there is the fundamental question of determining what one's total evidence supports when one receives higher-order evidence. Second, there is a closely related question about the level of impact that higher-order evidence has. In this section, I will briefly introduce the debate concerning both questions so that we have a good grasp of the current debate about higher-order evidence. It will then become clear in the next section how these debates bear on metaethics.

Most epistemologists accept that higher-order evidence is normatively significant – that is, it has a bearing on what a rational thinker should believe. All chapters in this volume presuppose this claim, and for good reasons.¹⁵

First, it would appear dogmatic to ignore higher-order evidence. Suppose that you get higher-order evidence that your reasoning from evidence E (e.g. your moral intuition that eating meat is morally wrong) to judgement B is erroneous. You would have to think as follows to be dogmatic. Given E, B. But given the higher-order evidence H, your reasoning from E to B might be erroneous. If you simply assert that E really supports B and thus conclude that H cannot be true and that therefore B, you would be dogmatic. In effect, you have reasoned your way to B by relying on B.¹⁶

Second, we are fallible creatures, and it is plausible that even beliefs about necessary truths are prone to error, and thus, there must be situations in which we have reason to revise such judgements, even when we are correct. For example, even after a competent deduction, there will be cases where we have reason to

change our views. Regular undercutting defeat is often characterised as the phenomenon of showing that an evidential relation that holds between E and H in circumstances C does not hold now insofar as C does not obtain (cf. Pollock and Cruz 1999). For example, suppose you are on a factory visit, looking at seemingly red wedges on a conveyor belt. In normal circumstances, you can trust your sight, so the fact that the wedges seem red to you is evidence that they are. But the factory foreperson's claim that the wedges are illuminated by red light undercuts your justification for thinking so. So in typical cases of undercutting defeat, an evidential relation that holds in normal circumstances is revealed to be missing in the actual circumstances. But if one has competently worked out a proof the evidential relation between E and H necessarily holds, undercutting defeat may not be able to explain a change in a moral view about necessary truths.¹⁷ Higher-order defeat, in contrast to regular undercutting defeat, suggests that one has taken an evidential relation to obtain that never obtained; it is "revisionary" in this sense (Lasonen-Aarnio 2014). So the normative significance of higher-order evidence may help to explain why we ought to change our views even about necessary truths.

Thus, higher-order evidence about morality should, somehow, affect our moral views. But how, exactly? The two views that have most commonly been defended in the literature are the conciliatory view and the steadfast view.¹⁸ Both will play a role in what follows.

According to the conciliatory view, one should respond to (negative) higher-order evidence by becoming less confident in one's original view and less confident that one's first-order evidence supports one's view. The position developed from

considerations of the significance of peer disagreement (Christensen 2007, 2009; Matheson 2009). However, its scope can be broadened to encompass appropriate reactions to higher-order evidence in general. Conciliatory views are partly motivated by the arguments supporting the normative significance of higher-order evidence introduced earlier. They also accord with the intuition that ignoring higher-order evidence is dogmatic and therefore inappropriate. This intuition seems particularly strong when one considers dogmatic reactions to peer disagreement.

According to the steadfast view, one should respond to (negative) higher-order evidence by maintaining one's confidence in one's original view, and one should also remain confident that one's first-order evidence supports one's view. Several views have been proposed to deal with this problem, and as several contributions to this volume show, adopting or strengthening one of these views can have implications for the debunking arguments mentioned earlier.

Kelly (2010) suggests that both the steadfast view and the conciliatory view are onto something. The conciliatory view is correct in suggesting that what is rational to believe about our evidence is not wholly independent of what is rational to believe about the world. Learning that someone like Peter Singer believes that eating meat is wrong may be a reason to think that you assessed the evidence incorrectly. The steadfast view correctly shows, argues Kelly, that if something is genuinely good evidence for a given conviction, then that fact itself will contribute to the epistemic justification for believing that conviction (Kelly 2010, 159). Kelly goes on to defend what he terms the total evidence view. According to this view, the evaluation of one's total evidence (e.g. your moral intuition about harming

animals and the information about, say, the meat-eating culture in which you grew up) will depend on the strength of the higher-order and the lower-order evidence.¹⁹

As Tiozzo (2019, ch. 2) points out, Kelly's view appears to have the advantages of both conciliatory and steadfast views, but it also has their downsides. By aggregating the evidence to find a suitable compromise, it seems that one will give neither the higher-order evidence nor the first-order evidence its full due.

Another way to look at the significance of higher-order evidence is in terms of the interaction of levels of belief. The conciliatory, steadfast, and total evidence views all agree that levels of belief are connected; that is, they agree that there is a close connection between what is rational to believe about some proposition *p* and what is rational to believe about *p*. In contrast, the proponents of level-splitting views maintain that one's total evidence sometimes rationally requires divergent beliefs at different levels (Coates 2012; Hazlett 2012; Lasonen-Aarnio 2014; Weatherson n.d.).

Level-splitting views offer a way of dealing with misleading higher-order evidence. Misleading higher-order evidence misleads you about your first-order beliefs. For example, suppose you correctly believe that eating meat is permissible because of my friend's testimony. When a seeming source then tells you that my friend is a liar, you would be rationally required to be less confident about the impermissibility of eating meat. But, clearly, by connecting the impact of your higher-order beliefs about what your evidence supports to your first-order beliefs, you have put yourself in a worse position epistemically and morally. Level splitting

can avoid this problem. You could correctly maintain that eating meat is permissible while also maintaining that you don't have good evidence in support of this view. As we will see later, the level-splitting view is of interest in the metaethical debate because it offers a way for new information about the bad grounds for our moral beliefs to leave intact our justification for our first-order moral beliefs.²⁰

The main problem for level-splitting views, which, at the same time, is the main motivation for level-connecting views, is a concern with akratic beliefs. On a level-splitting view, one ends up believing things like “p but my evidence does not support p” or “p but I am not sure whether it is rational to believe that p.” Such akratic beliefs, or Moorean propositions, are widely considered to be irrational (cf. Horowitz 2014; Adler 2002). Moreover, the level-splitting view also entails that bootstrapping is permissible (Sliwa and Horowitz 2015, 2848).²¹

Thus, there is a fundamental question about what attitudes or beliefs one's total evidence supports when higher-order evidence is in the picture. The accounting problem is perfectly general; it arises in all areas in which one strives for rational belief. But there is special reason to attend to the accounting problem in the case of moral epistemology, as the next section shows.

4 The Classification Problem: Genealogy and Moral Philosophy

Several prominent metaethical debates turn on the question of how to react to genealogical information. In this section, I aim at showing that these debates can usefully be thought of as being (implicitly) concerned with the classification problem.²²

Let us first consider some of the sources of genealogical claims of relevance for moral epistemology. According to a long tradition in the humanities, which rose to prominence with Freudian psychoanalysis, subconscious processes influence and control our thoughts and behaviour (cf. Leiter 2004).²³ Of course, we would need to be much more precise about this claim to properly assess its truth. The basic thought behind it, however, is clear and sufficient for the purposes of this chapter: the introspective seeming that we are “in control” of our thoughts and behaviour in rational ways is a chimera and wrong in so many ways. An anecdote by the philosopher G.A. Cohen (2001) illustrates this thought nicely. Cohen relates how he chose Oxford over Harvard for graduate school and later realised that that his fellow Oxonians generally accepted the analytic/synthetic distinction, while Harvard students generally rejected it. But both groups were confronted with the same arguments. Apparently, Cohen surmised, the respective environment explained the philosophical views that his peers and he adopted. Cohen’s anecdote illustrates how even our considered philosophical views might be significantly influenced by seemingly irrelevant factors beyond rational control (see Vavova 2018; Sher 2001; White 2010).

From a different perspective, evolutionary biologists and evolutionary psychologists argue that the capacity for moral thought and behaviour (e.g. Tomasello 2016; Baumard 2016) as well as the content of some of our moral beliefs have evolutionary origins (cf. Barkow, Cosmides, and Tooby 1995). For example, the philosopher Sharon Street (2006) took up the latter position and argued that some of our fundamental moral intuitions, such as if doing something would promote the interests of a family member then that is a reason to do it, are best

explained by the content of moral intuitions being ultimately determined by natural selection (2006, 115). Organisms that promoted the interests of their family members had comparative advantages of organisms that evaluated the world differently, as explained by the theory of kind selection, so evolution explains well the content of some of our moral beliefs (Street 2006; see also Joyce 2006; see Buchanan and Powell 2015 for criticism)

With a view to the proximate causes of moral beliefs, neuroscientists (cf. Liao 2016), experimental moral psychologists (cf. Doris 2010), and social psychologists (cf. Forgas, Jussim, and van Lange 2016), among others, have revealed a host of seemingly morally irrelevant situational influences on our moral judgements.²⁴ For example, the philosopher and psychologist Joshua Greene argued that fMRI evidence suggests that moral judgements of a particular type, namely characteristically deontological moral judgements, have their origin in brain regions typically associated with emotional processing that get triggered by situations that are “up close and personal” (Greene et al. 2001, 2106; Greene 2008; see Kahane et al. 2015 for criticism).

The list of descriptive, genealogical accounts of moral judgement could be continued almost indefinitely. The pertinent metaethical question, however, is what can we learn philosophically from such genealogical findings?²⁵

The answers to that question given by moral philosophers can be grouped into three categories: some suggested that genealogical accounts of our moral beliefs have implications for what we ought to do (that is, they have first-order, normative implications); others claimed that they have implications for what there is (that is, they have metaphysical implications about the nature of morality); and

others have argued that genealogical claims have implications for what we ought to believe (that is, they have epistemological implications).

However, the former two claims can be subsumed under the latter, at least if interpreted in a plausible way. If genealogical claims have implications for what we ought to do, then they must do so by “trickling down” to our first-order moral beliefs from our higher-order beliefs about our moral beliefs.²⁶ For example, they might imply that we are rationally required to believe that we do not know that eating meat is permissible, which might require us to abstain from eating meat.²⁷ If genealogical claims have metaphysical implications, then their metaphysical implications are plausibly only indirect, in the sense that they determine what kinds of existential claims we are justified in making.²⁸ For example, genealogical claims might imply that some notable fact, such as that some moral beliefs are extremely widespread or that others are extremely disputed, is best explained by a theory that does not appeal to moral facts, so that we might lose our justification for postulating the existence of moral facts. The important point is that, in each case, genealogical claims about our moral beliefs are taken to have implications for what we ought to believe. Thus, such arguments can be subsumed under the following genealogical schema:

Genealogical Schema

Empirical premise: Beliefs of a set, M, have genealogy G (they are caused by process x, formed in environment y, or influenced by factor z).

Epistemological premise: If M-beliefs have genealogy G, then we are rationally required to change our M-beliefs in a way w.

Conclusion: So we are rationally required to change our M-beliefs in a way w.

The genealogical schema provides a generic answer to our question about the philosophical relevance of genealogical claims about moral beliefs. It says that we are rationally required to somehow change (a subset of) our moral beliefs because of their genealogy. How precisely should we change them? The most common interpretation of “change in a way w” has been to suggest that some genealogy “should reduce or nullify” our confidence in the respective moral beliefs, though an increase in confidence is equally possible.

Arguments that fall within the genealogical schema have been tremendously influential in moral philosophy, as the following sketches illustrate. Montaigne (1877) and more recently and more sophisticatedly Wong (2006) argue that most of our moral beliefs are heavily influenced by culture and that we therefore cannot rationally maintain that these beliefs are true independently of culture. Ruse and Wilson (1986) argued that all our moral beliefs are ultimately caused by evolutionary pressures and that we therefore cannot rationally maintain that these beliefs are true independent of our human nature. In recent metaethics, Greene (2008) argues that characteristically deontological moral beliefs are caused by processes that do not work in our current environments and that therefore deontological beliefs are unreliable. Street (2006) and Joyce (2006) argue that the evolutionary origin of our moral beliefs implies that we cannot have moral

knowledge, at least if robust versions of moral realism are true. McGrath (2008), among others, has wondered whether widespread moral disagreement implies that at least some of us formed our moral beliefs in a problematic way such that we might lose our justification for maintaining them.

If it is correct that these different debates fit the genealogical schema, then we can make an interesting observation. The genealogical schema raises the classification problem when we consider the soundness of the epistemic premise: why should a given genealogy, *G*, have a bearing on whether we are rationally required to change our moral beliefs in some particular way? Thus, several central debates in recent moral epistemology rely on an answer to the question of when new information provides evidence concerning our higher-order moral beliefs (e.g. the belief that our moral beliefs are based on good evidence).

Interpreting the debate about genealogical explanations in moral philosophy as asking for a solution to the classification problem will help us to see how these different scientific approaches to morality are ultimately connected on the epistemological level. Thus, why do any of these claims bear on the rationality of our higher-order moral beliefs (if they do)?

First, as we have seen in Section 2, we should be reasonably confident that such genealogical claims are true, so the ongoing debates about their truth is crucial to moral epistemologists interested in the actual truth and justification of our moral beliefs. Nonetheless, even a true genealogical claim does not per se bear on the rationality of our moral beliefs: true genealogical claims are not sufficient for a warranted change in moral view. There needs to be some “connection” between the higher-order information and moral beliefs. We need to ascertain

whether these claims, if true, bear on whether our higher-order views (e.g. “my moral beliefs are based on good evidence”) are warranted.

Therefore, to assess the soundness of any argument that is an instance of the genealogical schema, we need a solution to the classification problem. Once we have solved the classification problem, we should be able to tell which genealogical accounts classify as higher-order evidence, as having a bearing on the rationality of our beliefs about our moral evidence. Once we have solved the accounting problem, we should be able to tell how that bears on the rationality of our moral beliefs themselves.²⁹ This is what this book aims to contribute to.³⁰

5 Connecting the Classification Problem With the Accounting Problem

The classification problem and the accounting problem are connected in moral epistemology. I will focus on five issues in recent moral epistemology that connect genealogical claims about morality with the classification problem and the accounting problem. The metaethical debate on genealogical information, in the debates about evolutionary debunking arguments³¹ and moral disagreement³² in particular, has increasingly “gone epistemological” in the sense that its participants have more explicitly begun to interpret these debates as concerning questions about what we should take as higher-order evidence concerning our moral beliefs (and why).

First, can we classify genealogical information about our moral beliefs as higher-order evidence? That is, do genealogical claims have a bearing on whether our moral beliefs are rationally formed? That is the classification problem. According to at least one interpretation, debunking arguments in moral philosophy provides us with evidence of error. In the characteristic language used

by some proponents of debunking arguments, some influences seem to take our moral beliefs “off track” with regard to moral truth (cf. Street 2006). If genealogical claims are evidence that our moral beliefs are unlikely to be true, then they classify as bona fide higher-order evidence. However, several scholars have noted that targeting all moral beliefs by way of such an argument runs the risk of self-defeat (e.g. Klenk 2017; Vavova 2018; Rini 2016). That is because, presumably, we can know that all our moral beliefs are unlikely to be true only by presupposing some moral truths, which is illegitimate in an attempt to defeat all moral beliefs. However, that is the case only if one accepts a conciliatory view on the accounting problem (as Wittwer discusses in this volume). Apart from aiming to defeat all moral beliefs, there is an open question regarding whether genealogical claims might be evidence of error for a subset of our moral beliefs, which leaves open the possibilities for answering the classification problem at this level (Sauer 2018; May 2018).

A second, closely related debate depends on solving the allocation problem. Even if some genealogical claim is legitimate evidence of error concerning (a subset of) our moral beliefs, there is the possibility to provide additional evidence that our moral beliefs are likely to be true. Providing such evidence requires assumptions about the moral truth. However, for such a strategy to work, one must violate a variant of what is known as the independence principle in the higher-order evidence debate (cf. Christensen 2010, forthcoming). The thought behind the independence principle is as follows. Suppose that you receive evidence that your reasoning in support of your belief that *p* was faulty. According to the independence principle, you ought to reason about whether *p* without taking into

account the evidence that led you to believe that p in the first place. If violating the independence principle is permissible, then debunking arguments can probably be disproved (cf. Klenk 2018c). At the same time, if the independence principle is false, then conciliatory solutions to the accounting problem lose some of their support. Thus, these questions about the classification problem in metaethics bear on solutions to the accounting problem.

Third, genealogical claims may be evidence that moral beliefs fail some other significant epistemic requirement. They need not be evidence of error, but may be, for example, evidence about the reliability, safety, or sensitivity of one's moral beliefs. In evaluating whether that is the case, we face another instance of the classification problem. Are available genealogies evidence in support of any of these claims? And, if so, in what ways should we change our moral beliefs in response to them? Here, a sometimes neglected aspect of the classification problem (what is higher-order information evidence of?) and the accounting problem (what are we rationally required to believe in response, given all our evidence?) combine. Metaethicists have begun to question the principles in virtue of which we ought to change our moral views (Klenk 2019; Clarke-Doane and Baras 2019). Scholars have defended several epistemic principles, some arguing that none of them are violated by any currently available genealogical information.

Fourth, once we have classified some genealogical claim as higher-order evidence, we face the accounting problem of deciding what our total evidence supports. That issue has been most pronounced in debates about moral peer disagreement. Moral disagreement, at least among peers that are roughly equally likely to get moral matters right, is widely thought to raise a pernicious challenge

for the view that moral judgements are justified.³³ Disagreement challenges presuppose an answer to the classification problem; the heart of the debate concerns the allocation problem. Given that one possesses good evidence for one's moral views, how should one account for the higher-order evidence provided by the information about the disagreement? Interestingly, the allocation problem should in equal measure arise in the debunking debate. Here as well, we should suppose that one possesses good evidence for one's moral views and thus ask how any higher-order evidence can be accounted for. Some connections between both debates have already been established (Klenk 2018a; Bogardus 2016; Mogensen 2016). In principle, the same questions that arise in the debate about peer disagreement regarding the accounting problem should be applicable in other debates concerning genealogical information.

Finally, the genealogical schema, as noted earlier, makes room also for increasing confidence in one's moral views in response to genealogical information. On this point, there have been two debates in recent moral epistemology, concerning the epistemic significance of moral testimony (Hills 2009; McGrath 2009) and the requirements for collective moral knowledge (e.g. Anderson 2016). The classification problem looms for the first debate: can testimony provide us with evidence in moral epistemology? The questions we discussed earlier, in turn, should equally apply on the collective level: how should higher-order evidence be accounted for on the collective level?

In reaching the current stage of the debate, metaethicists who have considered genealogical claims have increasingly "gone epistemological" by recognising the importance of some epistemological claims when thinking about

the metaethical implications of genealogical information. This volume aims to make further progress in this area by explicitly addressing the assumptions about the normative significance of higher-order evidence that underlie much of the recent metaethics debate. By exploring the implications of different views about the epistemic significance of higher-order evidence in moral epistemology, the book should also provide insights helpful in theory choice for scholars working on higher-order evidence.

6 The Structure of the Volume

This volume is the first to explicitly address the implications of higher-order evidence for topics in moral epistemology. It contains 12 previously unpublished essays, and it brings together leading international philosophers working in moral epistemology and the debate about higher-order evidence and several promising scholars who are at an earlier stage in their careers.

The volume is divided into four parts: the first is about the higher-order defeat of morality; the second argues against the higher-order defeat of morality; the third is about the wider implications of higher-order evidence in moral epistemology; and the fourth part is about permissible epistemic attitudes in response to higher-order evidence in moral epistemology. The first two parts build on the debates discussed earlier, namely the debate concerning the implications of genealogical explanations of moral beliefs and the epistemic significance of moral disagreement. The third part stays with the method of looking at moral epistemological topics through the higher-order evidence lens but widens the perspective to topics beyond the recent debunking debate. The chapters in this part address the role of moral testimony as higher-order evidence, higher-order

defeat in collective moral epistemology, and the rationality of fanatic beliefs in the light of research on higher-order evidence. Finally, the two essays in the fourth part take on a more fundamental perspective, asking how we can rationally respond to higher-order evidence concerning our moral views. In each part of the book, I have sought to include chapters that display a broad range of views, including established positions and new insights.

In Part I of the book, which is about the higher-order defeat of morality, **Silvan Wittwer** argues (in his contribution titled “Evolutionary Debunking, Self-Defeat and All the Evidence”) that evidence about moral peer disagreement could defeat all moral beliefs about which there is peer disagreement if the total evidence view is true (as opposed to a variant of conciliationism). Wittwer discusses how the total evidence view implies the falsity of the independence principle, according to which evidence of error regarding beliefs of some set, B, should be assessed independently of the contents of B. If the independence principle is false, then debunkers can show that widespread moral disagreement provides us with evidence of error. Wittwer explicitly discusses the self-defeat challenge raised by Vavova (2015, 2018) and shows that adopting the total evidence view is a way for debunkers to circumvent the challenge. However, to defeat all moral beliefs, debunkers must show that our total evidence implies that our moral beliefs are likely to be false. He suggests that an argument to this effect is unlikely to succeed in terms of showing the unreliability of moral beliefs but that it might succeed in terms of possible moral peer disagreement.

Looking beyond the epistemic relevance of disagreement, **Norbert Paulo**, in his contribution, “Moral Intuitions Between Higher-Order Evidence and Wishful

Thinking,” argues that recent moral psychology provides higher-order evidence against the reliability of moral intuitions as evidence for moral beliefs. In contrast to Wittwer, Paulo accepts the independence principle for the sake of argument. He claims, however, that the available higher-order evidence supports negative implications for the justification of our moral views. Paulo attests to a lack of evidence for the reliability of our moral intuitions and examines the evidence against the reliability of some moral intuitions that is discussed in recent situationist moral psychology. Paulo argues that continuing to rely on (considered) moral intuitions in the face of this evidence amounts to wishful thinking. Thus, according to Paulo, those who are sceptical about moral intuitions need not reject the independence principle to erect debunking arguments, as long as they can show that there are systematic influences on our moral beliefs that make them unreliable.

In the chapter by **Paul Silva Jr** titled “Debunking Objective Consequentialism: The Challenge of Knowledge-Centric Anti-Luck Epistemology,” we see how considerations about higher-order evidence can affect which normative theories we adopt. Silva argues that we lack support for accepting consequentialism as the true normative theory because of higher-order evidence against the view that our moral intuitions are best explained by consequentialism. Silva operates on the assumption that there is a knowledge-centric anti-luck epistemology. Accordingly, for a belief to be knowledge, its truth must not be lucky, and for a belief to be justified, it must qualify as knowledge. Hence, the view implies that propositions that cannot be known cannot be the content of justified beliefs. For example, before a winner is drawn, the belief that one’s lottery ticket

is a losing ticket cannot amount to knowledge, according to most epistemologists, because even in lotteries with minuscule odds of winning, it could be the case that one's ticket is a winner. Consequently, one cannot know that one's ticket is a losing ticket. Silva goes on to show, and this is the main innovation of the chapter, that consequentialist beliefs are relevantly similar to beliefs about lotteries. The net value of an action, it seems, could easily be different. That is so, argues Silva, even for seemingly uncontroversial, deeply held moral beliefs. Since support for consequentialism is often supposed to come from its ability to explain such uncontroversial moral beliefs, Silva's appeal to the higher-order evidence about the knowability of such propositions takes away that support.

A fourth way to rescue debunking arguments is discussed by **Olle Risberg and Folke Tersman** in their chapter "Disagreement, Indirect Defeat, and Higher-Order Evidence." They argue that some kinds of moral disagreement defeat our moral views because these kinds provide us with undercutting defeaters rather than higher-order evidence. They assume for the sake of argument that level-splitting views about higher-order evidence are correct. On the face of it, that assumption makes it difficult to argue that evidence of moral peer disagreement defeats the justification of some or all moral beliefs. After all, level-splitting views imply that higher-order evidence affects only our higher-order beliefs (e.g. that you have good evidence for your belief that eating meat is wrong) and not on our substantive moral beliefs (e.g. that eating meat is wrong). But Risberg and Tersman's strategy is to show that peer disagreement provides us with undercutting defeaters, which sever the link between the truth of a belief and the grounds on which a given agent holds the belief. In contrast, higher-order evidence

is often supposed to leave that link intact and thus have a different epistemic force than undercutting defeaters. Risberg and Tersman dispute this interpretation and they argue that higher-order evidence works like ordinary undermining defeaters. With that argument at hand, they show that moral peer disagreement provides us with undercutting defeaters and thus implies that we are rationally required to reduce confidence in our moral beliefs. Also, by weighing in on the nature of higher-order defeat, assimilating it with undercutting defeat, their chapter makes a contribution to the epistemological debate about the nature of higher-order defeat.

In the second part of the volume, which contains arguments against the higher-order defeat of morality, three chapters assess the prospects of higher-order defeat with a contrary conclusion. **Brian C. Barnett** argues in “Higher-Order Defeat in Realist Moral Epistemology” that given the correct view on higher-order defeat, moral realists can cope with ubiquitous debunking challenges (some of which I introduced earlier, such as the evolutionary debunking challenge and the moral disagreement argument), concluding that a significant count of ordinary first-order moral beliefs are safe from full defeat. Barnett’s argumentative strategy involves two steps. First, he sketches a theory of higher-order defeat in which he pays close attention to the relation between higher-order evidence and lower-order evidence to deduce which types of higher-order evidence defeat lower-order evidence. Second, similar to Risberg and Tersman, he looks at the evidence we gain from moral peer disagreement and evolutionary explanations of morality to see whether the evidence fits any of the defeating types of higher-order evidence that he deduced from his theory of higher-order defeat. The upshot of Barnett’s

work is that neither moral peer disagreement nor evolutionary explanations of morality fully defeat our moral beliefs. The reason for the failure of debunking of all our moral beliefs is that, argues Barnett, it is evidence about our moral beliefs, but the (negative) support it confers is inscrutable and thus unable to be fully defeated all our moral beliefs.

Marco Tiozzo takes a different route to put pressure on sceptical moral disagreement arguments. In his chapter “Moral Peer Disagreement and the Limits of Higher-Order Evidence,” he argues that higher-order evidence gained from moral (peer) disagreement fails to imply widespread moral scepticism. Tiozzo pays close attention to what it takes for a belief to lose its justification, given a body of higher-order evidence, and thereby focuses on the link between higher-order evidence and higher-order defeat. In doing so, he distinguishes the question of what one’s higher-order evidence supports (which Wittwer, Paulo, and Barnett focus on) from what one can rationally believe given that evidence. He argues that higher-order defeat, the lowering or nullifying of the justification of some set of first-order moral beliefs, requires that one believes that one’s belief fails to be rational. If one remains ignorant of this fact, then there is no defeat. Tiozzo then shows that on this interpretation of higher-order defeat, the success of the argument from moral peer disagreement is in jeopardy. He distinguishes between objective and subjective defeaters. The former are facts that affect evidential relations. The latter are beliefs about evidential relations. He argues that subjective defeaters are better suited to explaining higher-order defeat and that therefore higher-order defeat is contingent on what people believe. Since people generally do not take moral peer disagreement to be evidence against the

rationality of their moral beliefs, sceptical implications from moral disagreement are unlikely.

In the final essay of Part II of the book, **Michael Huemer**, in his contribution “Debunking Skepticism,” turns the tables on aspiring debunkers of our moral beliefs. Huemer exploits the fact that higher-order defeat can appear at many levels. If we take genealogical information concerning our moral beliefs to have sceptical implications – that is, to consider that information as constituting higher-order defeaters – then, argues Huemer, we should conclude that those beliefs in sceptical alternatives are based on inadequate evidence too. Find a defeater for them and they lose their force. Mirroring Paulo’s strategy of attending carefully to psychological explanations of moral beliefs, though for an opposing conclusion, Huemer argues that “moral scepticism is the product of an unreliable belief-forming process.” In doing so, Huemer accepts for the sake of argument that we have a good grasp of the defeating power of higher-order defeat, but he applies the same strategy “one level up.” He thereby adds a challenge for debunkers. When debunkers can solve the worry that their arguments are self-defeating, the same solution should allow us to raise a higher-order debunking argument against their view. That is, this can be the case as long as one accepts Huemer’s substantial empirical claims about the causes of sceptical belief, which may also provide an impetus for further moral psychological research.

In Part III of the book, which is about the broader implications of higher-order evidence in moral epistemology, we look beyond particular metaethical positions about the nature of morality (such as moral realism) to a broader range

of issues concerning the beliefs of fanatics, collective moral epistemology, and the epistemic significance of moral testimony.

Marcus Lee, Neil Sinclair, and Jon Robson assess the implications that accepting the normative significance of higher-order evidence has for a view about moral testimony (e.g. when my friend tells you that harming animals is wrong and you believe it on that basis) in their contribution “Moral Testimony as Higher-Order Evidence.” Moral testimony is often viewed with suspicion regarding its ability to confer moral knowledge or justification (Hills 2013, 2009). Does such pessimism about moral testimony imply that moral testimony cannot be a source of higher-order defeat? Lee, Sinclair, and Robson show that pessimism about moral testimony does not imply scepticism about the negative, or debunking, effect of moral testimony on justification. Thus, although moral testimony may not help us to justify moral beliefs, it may defeat our beliefs. When Risberg and Tersman note that peer disagreement can be taken as first-order evidence, they are implicitly relying on the view that moral testimony can work as higher-order evidence, a view that is corroborated by the chapter of Lee, Sinclair, and Robson. Lee, Sinclair, and Robson are thus able to show how peer disagreement may be significant, even if moral testimony from our peers would not help us to justify moral beliefs or moral knowledge.

The epistemology of group beliefs, collective epistemology, has been another topic of recent interest in moral epistemology. **J. Adam Carter and Dario Mortini** examine it through the lens of research on higher-order evidence in their chapter “Higher-Order Defeat in Collective Moral Epistemology.” They assume for the sake of argument that moral knowledge involves cognitive achievement, a

thesis that has recently gained much support (e.g. Pritchard 2012). According to this view, it is not sufficient for a belief to be reliably true to qualify as knowledge, but its truth has to be, in some sense, achieved by the believer in question. Carter and Mortini transport the achievement requirement to the collective level, briefly outlining what the conditions would be for groups to attain moral knowledge. They argue that the combination of an achievement view concerning moral knowledge and current proposals regarding collective epistemology suggest that collective moral knowledge is “extremely fragile.” That is, groups have a hard time satisfying the requirements of the achievement thesis and thus claim that a particular piece of collective moral knowledge would easily be defeated. It remains to be discussed whether this result is a strike against the achievement view or whether an improved account of collective moral knowledge can deliver a more stable moral epistemology.

In the final chapter of Part III, we turn to a practical problem of particularly recent concern: the beliefs of fanatics. Fanatics, who often but not necessarily have a religious background, have what seem like problematic beliefs, and they seem resistant to change. Nonetheless, some argue that fanatics can be rational. On one account, they are sufficiently ignorant of conflicting information; they have a “crippled epistemology” (Sunstein 2009). So sociological facts might explain how the beliefs of fanatics can be rational, even though their beliefs seem far from it from the outside. **Joshua DiPaolo** shows in his contribution, “The Fragile Epistemology of Fanaticism,” however, that fanatics are ignorant of higher-order evidence and that existing explanations of the rationality of fanatic beliefs fail to take this into account. By approaching the epistemology of fanatic beliefs from the

higher-order evidence perspective, DiPaolo refines the explananda of those theories that try to explain the rationality of fanatics. He describes how fanatics, in fact, react to higher-order evidence, claiming that fanatics possess relevant higher-order evidence against their fanatic beliefs while also ignoring this evidence. He finds that fanatics do not treat higher-order evidence as evidence (as they, plausibly, should) but instead as threats to their identity. Hence, if one wants to maintain that fanatics are rational, one has to explain how taking evidence as a threat to one's identity can be rational. The chapter thus also perfectly fits the volume by showing how a particular moral epistemological debate has thus far neglected the normative significance of higher-order evidence.

In Part IV of the book, which is about permissible epistemic attitudes in response to higher-order evidence in moral epistemology, **Margaret Greta Turnbull and Eric Sampson** argue in their chapter "How Rational Level-Splitting Beliefs Can Help You Respond to Moral Disagreement" that one can undogmatically split higher-order evidence so that it has a bearing only on one's higher-order beliefs. They thereby show that one need not adopt a conciliatory view to avoid the charge of dogmatism and irrationality and, nonetheless, it is possible to avoid higher-order defeat of morality (which may come, as several other contributors discuss, through moral disagreement and evolutionary explanations of morality). In doing so, their chapter contributes to the mounting defence against the higher-order defeat of morality, which is also discussed in Part II of this volume. The main innovation of their chapter, however, is that they can show that level splitting can be a rationally permissible, and undogmatic, reaction to higher-order evidence. The key element of their argument is their distinction between

higher-order evidence about one's assessment of a given body of evidence (which is what most epistemologists have focused on) and higher-order evidence about whether one reasoned from an illusionary or incomplete body of evidence. Noting this distinction, they argue that it can be rational to maintain what one's first-order evidence supports and also maintain that taking different evidence into account supports a different belief. They defend their argument by showing that two objections to level splitting by Horowitz (2014) do not apply to their version of level splitting. The authors thereby carve out space for rational level splitting and, furthermore, argue that adopting such beliefs allows level splitters "to occupy the narrow territory of humility, between servility and arrogance" (21).

Almost all of the contributions to this volume implicitly assume that how we should settle how to change our moral beliefs – that is, what epistemically permissible attitudes to adopt – involves analysing concepts such as justified belief. **Justin Clarke-Doane** in his chapter "Epistemic Non-factualism and Methodology" goes against this assumption. He asks us to consider normative questions, about epistemology and metaethics, as questions that concern what to believe and what to do. The different chapters in this volume propose different answers to this question, which can be understood as suggesting different theories about what it takes for beliefs to be justified (e.g. Silva's adoption, for the sake of argument, of knowledge-centric anti-luck epistemology or Carter and Mortini's adoption of virtue epistemology). Clarke-Doane argues, however, that an open question remains about whether we must to adopt any particular understanding of justification or whether there might be equally good grounds to adopt another one. In doing so, Clarke-Doane transposes the well-known open-question

argument from descriptive to evaluative properties. If Clarke-Doane is right, then we face an under-studied epistemological question that concerns how we should pick our epistemic concepts, which clearly has a bearing on how we could settle on solutions to the classification problem and the accounting problem.

The sets of issues included in this volume have been chosen both for their intrinsic interest and for their importance in central topics that are included in the current metaethical debate. I hope that they give the reader a sense of how progress is possible in the debunking debate, by using the finer epistemological tools at hand and a sense of excitement offered by the higher-order evidence perspective on other topics in moral epistemology. The contributions in this volume demonstrate valuable starting points for future examinations.

Notes

1 Thanks to Jaakko Hirvelä, Herman Philipse, Ibo van de Poel, Martin Sand, Steffen Steinert, and Marco Tiozzo for helpful comments on a previous version of this chapter and to Liam Deane for an instructive discussion.

2 Ordinarily speaking, objects (e.g. a gun with fingerprints on it at a crime scene) are taken to be evidence, though prevalent epistemological approaches maintain that one's knowledge or (subsets of) one's mental states constitute one's evidence; see Williamson (2000) and Conee and Feldman (2004), respectively. Intuitions about moral principles (e.g. harming sentient beings is wrong), knowledge of facts related to those principles (e.g. "X is a sentient being"), and sound arguments that suggest that a certain state of affairs obtains are thus evidence in this general sense.

3 E.g. Cameron, J. and Cameron, S. Amis (December 4, 2017).

4 Early proponents of non-cognitivism (e.g. Ayer 1971 [1936]; Hare 1963) may deny this, but most theories across the metaethical spectrum make room for speaking about moral truth; this includes modern versions of non-cognitivism (e.g. Gibbard 1990; Blackburn 1998) and non-naturalist moral realism (e.g. Shafer-Landau 2003; Enoch 2011; Huemer 2005), error theory (e.g. Mackie 1977), and naturalist realism (e.g. Brink 1989; Railton 1986) and of more subjectivist leanings (e.g. Prinz 2007). To consider evidence to be any factor that makes a given state of affairs more probable, they might need to adopt a correspondence theory of truth, which would also exclude modern versions of non-cognitivism.

5 Such factors might include intuitions, beliefs, and other mental states, as well as facts, propositions, events, and worldly objects; see Kelly (2016). It is commonly understood that the truth of most moral views depends on both descriptive and normative factors. For example, the truth of the view that eating meat is impermissible plausibly depends, in part, on the descriptive fact that animals feel pain. Hence, there can be descriptive evidence in support of this view (e.g. scientific studies about the extent to which animals feel pain). But that view also depends on normative principles (e.g. a prohibition against causing harm). What counts as evidence for such principles will be determined by one's account of the nature of moral truth. Depending on how we fill in the details, intuitions, and emotions, which can be socially informed, can be interpreted as evidence according to metaethical views from expressivism (Gibbard 1990) to non-naturalist realism (Enoch 2011); see also (Roeser 2011).

6 See Climenhaga (2018) and Bedke (2008).

7 Compare Durkheim (1995), based on whose work one might provide the defence for such a view.

8 See, for example, Greene (2013).

9 See Leiter (2004) for a discussion of a related genealogical suspicion that is especially highlighted by Friedrich Nietzsche, Sigmund Freud, and Karl Marx.

10 Friedrich Nietzsche (1887 [2013]) coined the term genealogy as applied to concepts, beliefs, and values in his endeavor to discredit what he took to be Christian values based on an analysis of their origins (see Klenk 2018b for more discussion). I use the term genealogy in a wider sense. They can take many forms; although they typically concern the history and development of a concept, belief, or value, I interpret the claim widely, so that they encompass information about how individuals attained a concept, formed a belief, or endorsed a value, cf. Queloz (forthcoming).

11 We get higher-order evidence whenever we learn about the circumstances in which we formed a belief, about the causes of our beliefs, or about the way we interpreted the evidence. The term higher-order evidence is used in several non-equivalent ways in the literature. Some take it to mean evidence that has a bearing on evidential relations. Others focus on evidence about the rationality of the person's thinking or on evidence about the reliability or accuracy of the person's thinking (Christensen forthcoming).

12 In contrast to first-order evidence, which affects it directly.

13 Compare the following: whether or not you interpreted Peter Singer's (1975) arguments correctly does not determine whether eating animals is permissible, but it does affect how confident you should be about that question.

14 I am indebted to Tiozzo (2019, ch. 2) for the setup of this section.

15 However, see Whiting (2017) for an argument rejecting the significance of higher-order evidence.

16 See Elga (2007), Christensen (2010, 2011), and Sliwa and Horowitz (2015) for similar bootstrapping cases.

17 Insofar as some moral truths are necessary, there would not be a rational change in a moral view about these either.

18 Horowitz (2014); see also Tiozzo (2019, ch. 2.2).

19 See Wittwer's contribution to this volume for further discussion of the total evidence view.

20 See the contributions by Risberg and Tersman and by Turnbull and Sampson for further discussion of level-splitting views.

21 Though see Silva (2017).

22 It is not my aim to provide a thorough overview of moral epistemology. See Campbell (2015) and Zimmerman, Jones, and Timmons (2018) for useful introductions and overviews.

23 As Richard Joyce (2016a, 2) notes, genealogical debunking arguments are commonplace and were made much earlier in recorded scholarship. That they are part and parcel of our everyday lives should be clear by the examples given in the introduction.

24 See Norbert Paulo's contribution to this volume, as well as Klenk and Sauer (2019), for recent evaluations of the situationist evidence.

25 I am treating genealogical findings as a class here to find common themes in how philosophers have reacted to them. Of course, individual reactions will look

quite different, depending on whether they concern, for example, Darwin's early formulation of the evolution of morality or later improvements of the theory.

26 This is because of well-known problems, and worries, about attempts to derive first-order moral claims from purely descriptive premises. See Farber (1994) for a thorough discussion of the historically best-known case of evolutionary ethics.

27 Though see Coates (2012). See Paul Silva's contribution to this volume for a relevant argument that suggests a way, though Silva does not explicitly endorse it, for us to go from higher-order evidence to first-order conclusions.

28 Most arguments along these lines appeal to a variant of the parsimony principle; see, for example Harman (1977) and the argument about moral disagreement in Mackie (1977). See Lillehammer (2016) and Clarke-Doane (2016) for discussions on the view that several recent attempts to offer "debunking explanations" of morality, such as those offered by Joyce (2006), are instances that belong to this category.

29 Initially, scholars were most impressed with the alleged normative implications of genealogical findings (cf. Farber 1994). Later developments brought with them not only novel descriptive perspectives on morality but also new perspectives on how to interpret these findings metaethically; see Klenk (forthcoming) for an introduction.

30 It is another question, of course, how important the descriptive premise is to the soundness of arguments based on the epistemological debunking scheme; see Klenk (2017)

31 See Street (2006) and Joyce (2006) for the starting points of the modern discussion on evolutionary debunking arguments. See Kahane (2011) Wielenberg (2016) for recent overviews.

32 See Rowland (2017) for a helpful overview.

33 The argument from moral disagreement rose to popularity with the work of Mackie (1977). As Tiozzo shows in his contribution to this volume, however, the argument aims for a metaphysical conclusion, whereas the moral disagreement argument as understood here proceeds on epistemological terms.

References

- Adler, Jonathan Eric. 2002. "Akratic Believing?" *Philosophical Studies* 110 (1): 1–27.
- Anderson, Elizabeth. 2016. "The Social Epistemology of Morality." In *The Epistemic Life of Groups: Essays in the Epistemology of Collectives*, edited by Michael Brady and Miranda Fricker, 75–94. Oxford: Oxford University Press.
- Ayer, A.J. 1971 [1936]. *Language, Truth, and Logic*. Pelican Books. Harmondsworth: Penguin Books.
- Barkow, Jerome H., Leda Cosmides, and John Tooby. 1995. *The Adapted Mind: Evolutionary Psychology and the Generation of Culture*. Oxford: Oxford University Press.
- Baumard, Nicholas. 2016. *The Origins of Fairness: How Evolution Explains our Moral Nature*. Oxford: Oxford University Press.
- Bedke, Matthew S. 2008. "Ethical Intuitions: What They Are, What They Are Not, and How They Justify." *American Philosophical Quarterly* 45 (3): 253–69.
- Blackburn, Simon. 1998. *Ruling Passions: A Theory of Practical Reasoning*. Oxford: Oxford University Press.
- Bogardus, Tomas. 2016. "Only All Naturalists Should Worry about Only One Evolutionary Debunking Argument." *Ethics* 126 (3): 636–61. <https://doi.org/10.1086/684711>.
- Brink, David O. 1989. *Moral Realism and the Foundations of Ethics*. Cambridge: Cambridge University Press.

Buchanan, Allen, and Russell Powell. 2015. "The Limits of Evolutionary Explanations of Morality and Their Implications for Moral Progress." *Ethics* 126 (1): 37–67. <https://doi.org/10.1086/682188>.

Cameron, James, and Suzy Amis Cameron. 2017. "Animal Agriculture Is Choking the Earth and Making Us Sick: We Must Act Now." *The Guardian*, December 4. Accessed September 24, 2019. www.theguardian.com/commentisfree/2017/dec/04/animal-agriculture-choking-earth-making-sick-climate-food-environmental-impact-james-cameron-suzy-amis-cameron.

Campbell, Richmond. 2015. "Moral Epistemology." In *Stanford Encyclopedia of Philosophy*, edited by Edward N. Zalta. Winter. Accessed July 30, 2018. <https://plato.stanford.edu/archives/win2015/entries/moral-epistemology>.

Christensen, David. forthcoming. "Formulating Independence." In *Higher Order Evidence: New Essays*, edited by Mattias S. Rasmussen and Asbjørn Steglich-Petersen. Oxford: Oxford University Press.

Christensen, David. 2007. "Epistemology of Disagreement: The Good News." *The Philosophical Review* 116 (2): 187–217.

Christensen, David. 2009. "Disagreement as Evidence: The Epistemology of Controversy." *Philosophy Compass* 4 (5): 756–67. <https://doi.org/10.1111/j.1747-9991.2009.00237.x>.

Christensen, David. 2010. "Higher-Order Evidence." *Philosophy and Phenomenological Research* 81 (1): 185–215. <https://doi.org/10.1111/j.1933-1592.2010.00366.x>.

- Christensen, David. 2011. "Disagreement, Question-Begging, and Epistemic Self-Criticism." *Philosopher's Imprint* 11 (6): 1–22.
- Clarke-Doane, Justin. 2016. "Debunking and Dispensability." In Leibowitz and Sinclair 2016, 23–36.
- Clarke-Doane, Justin, and Dan Baras. 2019. "Modal Security." *Philosophy and Phenomenological Research* 65 (1): 87. <https://doi.org/10.1111/phpr.12643>.
- Climenhaga, Nevin. 2018. "Intuitions Are Used as Evidence in Philosophy." *Mind* 127 (505): 69–104. <https://doi.org/10.1093/mind/fzw032>.
- Coates, Allen. 2012. "Rational Epistemic Akrasia." *American Philosophical Quarterly* 49 (2): 113–24.
- Cohen, G.A. 2001. *If You're an Egalitarian, How Come You're so Rich?* Cambridge, MA: Harvard University Press.
- Conee, Earl, and Richard Feldman. 2004. *Evidentialism: Essays in Epistemology*. Oxford: Oxford University Press. www.loc.gov/catdir/enhancements/fy0620/2004300585-d.html.
- Doris, John M., ed. 2010. *The Moral Psychology Handbook*. Oxford: Oxford University Press.
- Durkheim, Émile. 1995. *Erziehung, Moral und Gesellschaft: Vorlesung an der Sorbonne 1902/1903*. Frankfurt am Main: Suhrkamp.
- Elga, Adam. 2007. "Reflection and Disagreement." *Noûs* 41 (3): 478–502. <https://doi.org/10.1111/j.1468-0068.2007.00656.x>.
- Enoch, David. 2011. *Taking Morality Seriously: A Defense of Robust Realism*. Oxford: Oxford University Press.

Farber, Paul Lawrence. 1994. *The Temptations of Evolutionary Ethics*. Berkeley, CA: University of California Press.

Feldman, Richard. 2005. "Respecting the Evidence." *Philosophical Perspectives* 19 (1): 95–119. <https://doi.org/10.1111/j.1520-8583.2005.00055.x>.

Forgas, Joseph P., Lee J. Jussim, and Paul A.M. van Lange, eds. 2016. *The Social Psychology of Morality*. New York, NY: Psychology Press.

Gibbard, Allan. 1990. *Wise Choices, Apt Feelings: A Theory of Normative Judgment*. Cambridge, MA: Harvard University Press.

Greene, Joshua D. 2008. "The Secret Joke of Kant's Soul." In *Moral Psychology: The Neuroscience of Morality: Emotion, Brain Disorders, and Development*, edited by Walter Sinnott-Armstrong, 35–79. A Bradford Book Vol. 3. Cambridge, MA: MIT Press.

Greene, Joshua D. 2013. *Moral Tribes: Emotion, Reason, and the Gap between Us and Them*. New York, NY: Penguin Books. Accessed September 23, 2019.

Greene, Joshua D., R.B. Sommerville, L.E. Nystrom, J.M. Darley, and J.D. Cohen. 2001. "An FMRI Investigation of Emotional Engagement in Moral Judgment." *Science* 293 (5537): 2105–8. <https://doi.org/10.1126/science.1062872>.

Hare, R.M. 1963. *The Language of Morals*. Oxford: Oxford University Press.

Harman, Gilbert. 1977. *The Nature of Morality: An Introduction to Ethics*. Oxford: Oxford University Press.

Hazlett, Allan. 2012. "Higher-Order Epistemic Attitudes and Intellectual Humility." *Episteme* 9 (3): 205–23. <https://doi.org/10.1017/epi.2012.11>.

- Hills, Alison. 2009. "Moral Testimony and Moral Epistemology." *Ethics* 120 (1): 94–127. <https://doi.org/10.1086/648610>.
- Hills, Alison. 2013. "Moral Testimony." *Philosophy Compass* 8 (6): 552–9. <https://doi.org/10.1111/phc3.12040>.
- Horowitz, Sophie. 2014. "Epistemic Akrasia." *Noûs* 48 (4): 718–44. <https://doi.org/10.1111/nous.12026>.
- Huemer, Michael. 2005. *Ethical Intuitionism*. Basingstoke: Palgrave Macmillan.
- Joyce, Richard. 2006. *The Evolution of Morality. Life and Mind*. Cambridge, MA: MIT Press.
- Joyce, Richard. 2016a. *Essays in Moral Skepticism*. Oxford: Oxford University Press.
- Joyce, Richard. 2016b. "Evolution, Truth-Tracking, and Moral Scepticism." In *Essays in Moral Skepticism*, 142–58. Oxford: Oxford University Press.
- Kahane, Guy. 2011. "Evolutionary Debunking Arguments." *Noûs* 45 (1): 103–25. Accessed October 3, 2014.
- Kahane, Guy, Jim A.C. Everett, Brian D. Earp, Miguel Farias, and Julian Savulescu. 2015. "Utilitarian' Judgments in Sacrificial Moral Dilemmas Do Not Reflect Impartial Concern for the Greater Good." *Cognition* 134: 193–209. <https://doi.org/10.1016/j.cognition.2014.10.005>.
- Kelly, Thomas. 2005. "The Epistemic Significance of Disagreement." In *Oxford Studies in Epistemology*. Vol. 1, edited by Tamar S. Gendler and John P. Hawthorne, 167–96. Oxford: Oxford University Press.

- Kelly, Thomas. 2010. "Peer Disagreement and Higher-Order Evidence." In *Disagreement*, edited by Richard Feldman and Ted A. Warfield, 111–74. Oxford: Oxford University Press.
- Kelly, Thomas. 2016. "Evidence." In *Stanford Encyclopedia of Philosophy*, edited by Edward N. Zalta. <https://plato.stanford.edu/archives/win2016/entries/evidence/>.
- Klenk, Michael. forthcoming. "Moral Naturalism and Evolutionary Ethics." In *Introduction to Philosophy: Ethics*, edited by Christina Hendricks, 1–16. The Rebus Community.
- Klenk, Michael. 2017. "Old Wine in New Bottles: Evolutionary Debunking Arguments and the Benacerraf–Field Challenge." *Ethical Theory and Moral Practice* 20 (4): 781–95. <https://doi.org/10.1007/s10677-017-9797-y>.
- Klenk, Michael. 2018a. "Evolution and Moral Disagreement." *Journal of Ethics and Social Philosophy* 14 (2): 112–42. <https://doi.org/10.26556/jesp.v14i2.476>.
- Klenk, Michael. 2018b. "Survival of Defeat: Evolution, Moral Objectivity, and Undercutting." PhD thesis, Utrecht University.
- Klenk, Michael. 2018c. "Third Factor Explanations and Disagreement in Metaethics." *Synthese*. <https://doi.org/10.1007/s11229-018-1875-8>.
- Klenk, Michael. 2019. "Objectivist Conditions for Defeat and Evolutionary Debunking Arguments." *Ratio*: 1–14. <https://doi.org/10.1111/rati.12230>.
- Klenk, Michael, and Hanno Sauer. 2019. *Situationism and the Control Requirement for Moral Progress*. Manuscript.
- Lasonen-Aarnio, Maria. 2014. "Higher-Order Evidence and the Limits of Defeat." *Philosophy and Phenomenological Research* 88 (2): 314–45.

- Leibowitz, Uri D., and Neil Sinclair, eds. 2016. *Explanation in Ethics and Mathematics: Debunking and Dispensability*. Oxford: Oxford University Press.
- Leiter, Brian. 2004. "The Hermeneutics of Suspicion: Recovering Marx, Nietzsche, and Freud." In *The Future of Philosophy*, edited by Brian Leiter, 74–105. Oxford: Oxford University Press.
- Liao, S. Matthew, ed. 2016. *Moral Brains: The Neuroscience of Morality*. Oxford: Oxford University Press.
- Lillehammer, Hallvard. 2016. "‘An Assumption of Extreme Significance’: Moore, Ross, and Spencer on Ethics and Evolution." In Leibowitz and Sinclair 2016, 103–23.
- Mackie, John Leslie. 1977. *Ethics: Inventing Right and Wrong*. London: Penguin Books.
- Matheson, Jonathan. 2009. "Conciliatory Views of Disagreement and Higher-Order Evidence." *Episteme* 6 (3): 269–79. <https://doi.org/10.3366/E1742360009000707>.
- May, Joshua. 2018. *Regard for Reason in the Moral Mind*. Oxford: Oxford University Press.
- McGrath, Sarah. 2008. "Moral Disagreement and Moral Expertise." In *Oxford Studies in Metaethics*. Vol. 3, edited by Russ Shafer-Landau, 87–108. Oxford: Oxford University Press.
- McGrath, Sarah. 2009. "The Puzzle of Pure Moral Deference." *Philosophical Perspectives* 23 (1): 321–44.

- Mogensen, Andreas L. 2016. "Contingency Anxiety and the Epistemology of Disagreement." *Pacific Philosophical Quarterly* 97 (4): 590–611. <https://doi.org/10.1111/papq.12099>.
- Montaigne, Michel de. 1877. "Of Cannibals." In *Essays of Michel De Montaigne*, edited by William C. Hazlitt, 1–10.
- Nietzsche, Friedrich Wilhelm. 1887 [2013]. *On the Genealogy of Morals: A Polemic*. London: Penguin Books. translated by Michael A. Scarpitti; with an introduction and notes by Robert C. Holub.
- Pollock, John L., and Joseph Cruz. 1999. *Contemporary Theories of Knowledge*. Lanham, MD: Rowman & Littlefield.
- Prinz, Jesse J. 2007. *The Emotional Construction of Morals*. Oxford: Oxford University Press.
- Pritchard, Duncan. 2012. "Anti-Luck Virtue Epistemology." *The Journal of Philosophy* 109 (3): 247–79. <https://doi.org/10.5840/jphil201210939>.
- Queloz, Matthieu. forthcoming. "Nietzsche's English Genealogy of Truthfulness." *Archiv für Geschichte der Philosophie*.
- Railton, Peter. 1986. "Moral Realism." *The Philosophical Review* 95 (2): 163–207. <https://doi.org/10.2307/2185589>.
- Rini, Regina A. 2016. "Debunking Debunking: A Regress Challenge for Psychological Threats to Moral Judgment." *Philosophical Studies* 173 (3): 675–97. <https://doi.org/10.1007/s11098-015-0513-2>.
- Roeser, Sabine. 2011. *Moral Emotions and Intuitions*. Basingstoke: Palgrave Macmillan.

- Rowland, Richard. 2017. "The Epistemology of Moral Disagreement." *Philosophy Compass* 12 (2): e12398. <https://doi.org/10.1111/phc3.12398>.
- Ruse, Michael, and Edward Osborne Wilson. 1986. "Moral Philosophy as Applied Science." *Philosophy* 61 (236): 173–92. Accessed December 13, 2016.
- Sauer, Hanno. 2018. *Debunking Arguments in Ethics*. Cambridge: Cambridge University Press.
- Shafer-Landau, Russ. 2003. *Moral Realism: A Defence*. Oxford: Oxford University Press. <http://site.ebrary.com/lib/alltitles/docDetail.action?docID=10266702>.
- Sher, George. 2001. "But I Could Be Wrong." *Social Philosophy and Policy* 18 (2): 64. <https://doi.org/10.1017/S0265052500002909>.
- Silva, Paul. 2017. "How Doxastic Justification Helps Us Solve the Puzzle of Misleading Higher-Order Evidence." *Pacific Philosophical Quarterly* 98: 308–28. <https://doi.org/10.1111/papq.12173>.
- Singer, Peter. 1975. *Animal Liberation: A New Ethics for Our Treatment of Animals*. New York, NY: Random House.
- Sliwa, Paulina, and Sophie Horowitz. 2015. "Respecting All the Evidence." *Philosophical Studies* 172 (11): 2835–58. <https://doi.org/10.1007/s11098-015-0446-9>.
- Street, Sharon. 2006. "A Darwinian Dilemma for Realist Theories of Value." *Philosophical Studies* 127 (1): 109–66. <https://doi.org/10.1007/s11098-005-1726-6>.
- Sunstein, Cass R. 2009. *Going to Extremes: How Like Minds Unite and Divide*. Oxford: Oxford University Press. <http://site.ebrary.com/lib/academiccompletetitles/home.action>.

Higher-Order Evidence and Moral Epistemology (Routledge, forthcoming)

Tersman, Folke. 2006. *Moral Disagreement*. Cambridge: Cambridge University Press.

Tiozzo, Marco. 2019. "Moral Disagreement and the Significance of Higher-Order Evidence." PhD thesis, Gothenburg University.

Tomasello, Michael. 2016. *A Natural History of Human Morality*. Cambridge, MA: Harvard University Press.

Vavova, Katia. 2015. "Evolutionary Debunking of Moral Realism." *Philosophy Compass* 10 (2): 104–16. <https://doi.org/10.1111/phc3.12194>.

Vavova, Katia. 2018. "Irrelevant Influences." *Philosophy and Phenomenological Research* 96 (1): 134–52. <https://doi.org/10.1111/phpr.12297>.

Weatherson, Brian. n.d. Do Judgements Screen Evidence?

White, Roger. 2010. "You Just Believe That Because . . ." *Philosophical Perspectives* 24: 573–615.

Whiting, Daniel. 2017. "Against Second-Order Reasons." *Noûs* 51 (2): 398–420. <https://doi.org/10.1111/nous.12138>.

Wielenberg, Erik J. 2016. "Ethics and Evolutionary Theory." *Analysis* 76 (4): 502–15. <https://doi.org/10.1093/analys/anw061>.

Williamson, Timothy. 2000. *Knowledge and Its Limits*. Oxford: Oxford University Press.

Wong, David B. 2006. *Natural Moralities: A Defense of Pluralistic Relativism*. Oxford: Oxford University Press.

Zimmerman, Aaron, Karen Jones, and Mark Timmons, eds. 2018. *Routledge Handbook of Moral Epistemology*. New York, NY: Routledge.