

CHAPTER 6

Skeptical Repercussions

This chapter discusses what I consider to be a new argument in metaethics, the *skeptical argument from moral underdetermination*. The plan for the chapter is as follows. To prepare the ground, I turn once more to Duhem and Quine. Both were keenly aware of the repercussions that underdetermination has for our understanding of physical theories and their postulates, yet they once more differed strongly in their assessments. What this shows is that even if we grant underdetermination, it is less than obvious what follows from this. The same holds for ethics, and we thus need more argument. I will commence my case with a discussion of why we should think of the challenge that moral underdetermination poses as an epistemological, not a semantic, one. Building on this, I will then transfer the skeptical argument from science to ethics. I identify a basic structure of the skeptical argument and show how it can be adapted to the realm of ethics. Finally, I defend the argument from two objections that are closely inspired by analogous objections in the philosophy of science. Both, I argue, can be countered at least as plausibly as, if not more plausibly than, their counterparts in the philosophy of science.¹

6.1 Duhem and Quine Again

In order to motivate the turn from normative ethics to metaethics, let us first hear once more from the two progenitors of underdetermination in science. Both grappled with the question of whether and how the phenomenon of underdetermination impacts our understanding of the nature and the goals of scientific theorizing, yet they came to very different conclusions. In Chapter 1, we saw that underdetermination, for Duhem, is a pervasive phenomenon at least for physical theories, if not for all

¹ Parts of this chapter are based on Baumann (2021b) and Baumann (2022).

theories that include a high level of mathematization and symbolization. Does he think that this poses a problem for physics (or the mathematical sciences) itself? Not necessarily. In programmatic fashion, Duhem (1906, pp. 7 ff.) distinguishes between two views of what the objective of physics is. One holds that physical theory aims at explanation, laying open the invisible causes of what we observe. The other contends that it is merely the summarizing and logical classifying of groups of experimental laws. Duhem subscribes to the second view. Explanations typically point to some unobservable object that accounts for the observable phenomena. Yet, in cases of underdetermination, different theories postulate different unobservables that cannot both be the causes. Science, Duhem urges, should not be grounded on such an insecure basis. Instead, scientists should refrain from thinking that they are uncovering the true nature of the world and only aim at a (formal) system that optimally organizes observable phenomena and the outcomes of experiments.

At least on one plausible reading, this understanding of physical theories has been interpreted as an *anti-realist* move by Duhem. *Scientific realists* assume that our scientific theories aim at unveiling the hidden makeup of the world and often do, or at least come close to doing so.² When confronted with the problem that underdetermination poses, Duhem chooses to give up on this realistic picture of science. He instead opts for a form of instrumentalism.³ Science, on that view, instead of yielding a true description of an independent reality, should aim only at the most economical way of systematizing and classifying observable phenomena. In combative fashion, Duhem (1906, p. 32) likens the explanatory aspirations of physical theories to a parasite that attaches itself to the fully autonomous body that is the proper physical theory.

Duhem (1906, pp. 24–27) is well aware that most people, including scientists, tend to search for more and will naturally be inclined to interpret physical theories as a representation of a natural ontological order. He even sees this as an advantage, since taking theories at face value will lead us to trust them more when it comes to future cases. Nevertheless, at least on a philosophical level, he urges us to settle for a non-explanatory view of science.

² I will come to a more precise definition of scientific realism shortly.

³ Compare Horwich (1991, pp. 1–2), Psillos (1999, pp. 28–33), and Okasha (2002, p. 103). As Horwich (1991, p. 2) notes, there are various other names that have been given to Duhem's and similar positions. Since I am not interested in these terminological subtleties, I will mostly use the catch-all term *anti-realism*, unless it is important for an argument to distinguish more precisely between the different positions that anti-realism encompasses.

What about Quine? We saw that Quine is very radical in his insistence that underdetermination pertains to all fields of knowledge and that we can in principle hold on to any hypothesis. Yet, perhaps somewhat surprisingly, he is much more conservative when it comes to appraising the upshots of underdetermination for our understanding of science and, more specifically, for the issue of scientific realism.

This is often misunderstood because some of his early remarks could be read as pointing to an anti-realist stance. For example, Quine (1951, pp. 44 ff.) frequently talks of *posits* in TDE.⁴ From ordinary physical objects, to the microcosmic ones of science, to forces, he discerns posits everywhere. Quine (1951, p. 44) even provocatively compares these posits to the *Gods of Homer*. On a similar note, Quine's use of the term *pragmatic* can cause misunderstandings. Quine (1951, pp. 45–46) dubs the criteria by which we choose some theories over others, such as simplicity, *pragmatic* (or *merely pragmatic*). Such a classification of theoretical virtues, as we will see, often figures in anti-realist arguments. If theories are underdetermined, so the idea goes, our choosing between them on the basis of pragmatic criteria might tell us something about our actual practice, but it does not render the theories more believable.

However, we need to be very careful when interpreting these claims. As to the first point, Quine only likens the posits of science and other posits to the Gods of Homer *epistemologically*. He thinks that epistemologically, those posits only differ in degree and not in kind. They are all brought in to allow for a better handling of our past and future experiences, and in this context they are judged according to the same criteria.⁵ Yet Quine (1951, p. 44) is very clear that he considers it a scientific error to hold on to the Gods of Homer and that he, as a lay physicist, believes otherwise. On a more principled note, Quine later points out that the mere fact that something is posited in order to help us deal with our experiences does not speak against its reality at all. As Quine (1960, p. 22) puts it: “[t]o call a posit a posit is not to patronize it.” As to the second point, Quine subsequently clarifies that he considers it not to the detriment of a theory's claim to truth if the theory is chosen on a pragmatic basis. Indeed, he considers such considerations to be exactly what guide us in our search for truth.

In later texts, Quine subsequently makes it very clear that he does not draw anti-realist conclusions from his considerations on

⁴ Compare also Quine (1966a, p. 210).

⁵ See also Quine (1966a, pp. 210–211) for a reiteration of these sentiments.

underdetermination. Quine (1966b, p. 241) is aware that an inkling of the phenomenon of underdetermination has contributed to some people seeing, for example, the particles of physics as a fiction. However, he himself considers this tendency to be a perverse one. The problem, as Quine (1975, pp. 327–328) sees it, is a looming threat of relativism inherent in versions of the underdetermination thesis that assume logical incompatibility. Quine holds that if two theories predict exactly the same data and there is no additional evidence that could decide between them, then both have the same claim to being true. However, in cases of radical underdetermination, those theories also contain propositions that cannot be true at the same time. Thus, Quine thinks, the propositions these theories contain can only be true relative to the theory itself. This consequence, however, is unacceptable to him.

Instead, Quine opts to give up his earlier thesis of radical underdetermination, instead making the claim that there can at most be underdetermination between theories that are notational variants of each other. The argument that he offers in support of this view Quine (1981, p. 30) calls the *trivial expedient*. The gist of the argument is as follows. Since the theories are empirically equivalent, they must, by hypothesis, agree on every sentence that has any observational upshots. If they nevertheless disagree, they must disagree about sentences that contain terms that are short on observational criteria. However, if observation does not determine the exact meaning of those terms, we do not have to consider them to be the same. Instead, we can treat those terms as different concepts that happen to be expressed by the same word. The radical conflict between the theories thus vanishes, and we can conceive of them as merely different in notation but not logically incompatible.

Quine (1975, p. 319) provides an example. Imagine someone taking a theory and randomly exchanging two words, say, ‘electron’ and ‘molecule.’ The resulting theory and the original one would disagree about many sentences. However, no one would be under the impression that they substantively disagree; they are not *radically* different. Instead, they just happen to express different concepts with the same word. To mark this, we can rename the words in the second theory as ‘electron*’ and ‘molecule*.’ This way, it becomes obvious that they do not disagree about anything beyond words. One theory would claim, for example, that electrons are smaller than molecules, whereas the other would claim that molecules* are smaller than electrons*. Quine holds that we can do this with any two

theories that are empirically equivalent, since the sentences they disagree on must be without observational content. We can thus just treat them as speaking in the same words about different concepts.⁶

Underdetermination, although it has close connections to skepticism, should thus not automatically be understood to have skeptical repercussions in science. Showing that it has skeptical upshots needs further argument. The same, I think, goes for ethics: Further analysis and argument is needed before we can appreciate the upshots of underdetermination for metaethics. To this I will now turn.

6.2 From a Semantic to an Epistemic Challenge

Appreciation for possible metaethical upshots of the consequentializing and deontologizing projects has been rather restricted so far. Dietrich and List (2017, p. 425), when making the analogy to science explicit, do mention how a parallel view to scientific instrumentalism in ethics might clash with their reason-based representation of moral theories. However, due to their interest being of a more formal nature, they do not pursue this line of thinking any further.

The most notable exception is Dreier.⁷ The original reason Dreier (1993) gives for consequentializing has to do with a structural point about normative ethics, that is, that the distinction between consequentialist and non-consequentialist views doesn't cut deeply. Yet, Schroeder suggests that at least when it comes to Dreier's later view, metaethical considerations are likely to have played a role as well:

I wonder if Dreier might offer such an argument based on a kind of broad metaethical skepticism (or at least agnosticism) which surfaces toward the end of his chapter. If Dreier could argue that there is no meaningful sense in which we can distinguish different types of ethical properties – that ethics in a fundamental sense is just about deontic verdicts, and that there isn't any further “real” structure – that might support the assimilation argument. (Schroeder, 2017, p. 1484)

I think that Schroeder is clearly right in his estimation that metaethical considerations play a role in Dreier's understanding of consequentializing, although I do not think that it is *skeptical* considerations. Dreier's point

⁶ It is unclear what role should be accorded to the trivial expedient in Quine's overall philosophy. Adeel (2010, p. 7) suggests that Quine only makes use of it in what might be considered his mid-period thinking about underdetermination but discards it toward the end of his career.

⁷ Suikkanen (2021) is another exception, which I will come back to in Chapter 7.

is not that we do not know which explanations, or which theories, are the correct ones. Instead, when putting forward his notational variants view, he offers a semantic argument about the conceptual connection between the notion of a good outcome and action-guidance. As I have remarked in Chapter 3, I consider Dreier's view to be the most important rival to the underdetermination interpretation and one that isn't easy to refute precisely because what are at stake between the two views are deep metaethical questions. For this reason, I have so far staked my case against Dreier's view on the fact that Dreier (2011, p. 114) himself acknowledges that his Extensionality Thesis (i.e., the claim that nothing but extension matters) is even more controversial than the Deontic Equivalence Thesis (i.e., the claim that we can build consequentialist counterparts to any non-consequentialist theory). The former thus provides little in support of the latter. However, this has merely been a dialectical move. We are now in a position to consider a stronger argument.

That argument starts from the observation that Dreier's reaction to extensionally equivalent theories in ethics is noticeably similar to Quine's reaction to extensionally equivalent theories in science. When faced with theories that can account for the same evidence – empirical observations on the one hand, intuitions or considered judgments about particular moral cases on the other – Quine and Dreier opt for the same solution. The alleged incompatibility between such theories must be an illusion; the disagreements merely notational. Moreover, Dreier's argument to this conclusion echoes Quine's trivial expedient. Quine thinks that we can dissolve the conflict in science because we can explain away any difference between these theories that is not sufficiently grounded in empirical observations. Similarly, Dreier (2011, pp. 114–115) thinks we can dissolve the conflict in ethics because their differences have no consequences for what they imply about action-guidance. Both authors are inclined to think that theoretical differences that do not amount to direct observational/action-guiding differences are to be explained in terms of a certain slack in theory building. If observation/action-guidance does not decide between the theories, we do not have to consider them to be incompatible even if they make seemingly incompatible claims.

Since Dreier's argument is structurally similar to Quine's trivial expedient in this way, we might learn something from how philosophers of science have evaluated the plausibility of the latter argument. This, I think, puts pressure on Dreier's view. Let us thus look at how philosophers of science have reacted to the trivial expedient. Bonk (2008, pp. 178–180) provides an illuminating discussion. He points out that Quine's strategy is based

on a form of *descriptivism* about the meaning of scientific terms. (This form of) descriptivism, Bonk (2008, p. 178) explains, entails that: “[...] what an object is, is fully determined by its place in the closely knit structure of explanatory hypotheses [...].” To illustrate what this means, consider Quine’s electron example again. The notion of an electron, on the descriptivist view, is one that we cannot characterize in a pre-theoretic way. Instead, we only have an understanding of it via the place it occupies within our scientific theories. Hence what the term means is completely determined by the place it takes in those theories. This is crucial for the success of the trivial expedient strategy because the object about which the rival scientific theories disagree does not occupy the same place in the respective theories (remember: the terms electron and molecule have been switched in one theory). Since the place in the theory determines the meaning of the object and that place differs in the example, we are thus entitled to give different names to the objects, electron and electron*. This way we see that the theories do not really disagree. Still, there is an alternative to descriptivism. On this view, Bonk (2008, p. 178) informs us: “[w]hat an object is, what properties it has, how it can be identified, is partly determined by salient phenomena and experiments [...].” These, importantly, are not determined by the theory. That is, the meaning of the terms on which the theories disagree is not wholly given by the theories themselves but tracks something that is (at least partly) defined independently of the theory. If this is so, we cannot simply rename the concepts on which the theories disagree. We were only allowed to do this under the assumption that the meaning of the terms was entirely given by their place in the overall theory. If this is not true, and the meaning is (at least partly) given independently of those theories, then the disagreement remains. Hence, whether the trivial expedience strategy works depends on our semantic theory of the terms used in scientific theories.

How does this all relate to our case in ethics? Note first that Dreier (2011, pp. 114–115) is more guarded than Quine. He explicitly distinguishes his view from that of Foot (1985), who argues that we have *no* independent notion of goodness as it relates to states of affairs. In contrast, Dreier only claims that the notion of a good outcome is *firmly moored* in action-guidance. Still, for his argument to succeed, Dreier has to assume that the notion of a good outcome cannot substantially come apart from what it implies about what we should do. And this, I think, is where he runs into problems. What makes Quine’s trivial expedient seem at least initially promising is that it is quite clear that we have no theory-independent grasp of a concept like electron. Since this is the case, Quine can dissolve any

seeming disagreement by showing that this concept plays a different role in the two theories. However, that does not seem to be the case with regard to the disagreement in ethics. Consequentialist moral theories employ the notion of a good outcome in their explanations, but the notion is not provided to us *by* those moral theories. If it were, Dreier could take a chapter from Quine's book. He could claim that non-consequentialists assume that the status of some acts is independent of the 'good outcome,' whereas consequentializers claim that the status of all acts is dependent on the 'good outcome*.' Since these two expressions do not pick out the same notion of a good outcome, they do not disagree. Indeed, non-consequentialists could accept that, on this wide understanding, they are indeed evaluating all acts according to their contribution to 'good outcomes*.' But this does not seem to portray the debate in an accurate way. Most non-consequentialists would repudiate the description that they are in fact evaluating acts exclusively according their contribution to the best outcome, *no matter how wide* the concept of a good outcome may be. Instead, they would insist that there is a real disagreement about an (at least partly) antecedently fixed concept of a good outcome. Put in another way, they wouldn't accept the consequentializers' framing as one that just employs a different notion of good outcomes. They would insist that on top of wrongly assuming that outcomes are all that matters, consequentializers also have a wrong view about what constitutes good outcomes. And they would defend these charges by appealing to a notion of a good outcome that they take to be independent of the role that notion plays in the explanation of which actions are right or wrong. In this undertaking, I think we can safely say, they are better off than someone trying to argue the same case regarding the notion of an electron. Our understanding of the notions that figure in normative moral explanations is much less theory-dependent than our understanding of some of the terms that are used in scientific theories. This makes arguments like the trivial expedient strategy much less promising in ethics.

The parallel to the philosophy of science thus helps us to focus on the real problem with Dreier's interpretation, which is an implausible semantics of the terms that figure in moral explanations. If the underdetermination interpretation can avoid such a view, this is an important advantage. And, indeed, I think that it can. It does so by shifting our focus away from semantics and to epistemology. As I am about to argue in the next section, the problem with deontically equivalent theories is not that we don't understand how they could nevertheless differ in their explanations. The problem is that we might not know which explanation is correct if the

theories turn out to be deontically equivalent. Very basically, the fact that we do not know which of two explanations is true does not mean that we have no clear concept of how the explanations differ. The problem, it seems to me, does not have to do with meaning. Instead, and this is what I will argue next, it is an epistemological one: We do not know which explanation is right, and this, in turn, might make us skeptical about moral explanation altogether.

This epistemological perspective is also the one that Parfit takes with regard to the metaethical upshots of his project. Parfit fears that widespread disagreement among ethicists gives rise to skeptical doubts. The success of his Convergence Argument, he hopes, helps alleviate these doubts by proving that agreement is after all attainable in ethics. I will argue that this is not so. The fact that deontically equivalent moral theories nevertheless disagree radically about what makes acts right gives rise to a new skeptical challenge.

6.3 Adapting the Skeptical Argument to the Realm of Ethics

In the philosophy of science, the skeptical streak in the discussion of underdetermination has been present since Duhem and arguably explains much of the interest that the phenomenon has generated. This interest was reinvigorated by van Fraassen's 1980 book *The Scientific Image*, which (re-)established the idea that underdetermination poses a direct threat to the position of *scientific realism* in the contemporary debate.⁸ Henceforth, the argument from underdetermination has often been seen as one of the two main arguments against scientific realism.⁹ If we want to understand what metaethical repercussions moral underdetermination might have, we should thus start here.

Let us thus start with the position under attack, that is, scientific realism. As Chakravartty (2017, p. 2) informs us, there are almost as many definitions of scientific realism as there are authors writing about it. However, he suggests a useful way of understanding the position in terms of three components:

The Semantic Component: Semantically, realism is committed to a literal interpretation of scientific claims about the world. [Scientific claims] should be construed literally as having truth values, whether true or false.

⁸ For an elucidation of van Fraassen's attack on realism, see Ladyman (2002, p. 185 ff).

⁹ The other being the so-called *pessimistic metainduction*. Compare Stanford (2006, pp. 6–8) and Bortolotti (2008, pp. 98–104). Compare also Kukla (1998, p. 58) for an assessment that underdetermination is *undoubtedly* the threat to scientific realism.

The Metaphysical Component: Metaphysically, realism is committed to the mind-independent existence of the world investigated by the sciences.

The Epistemological Component: Epistemologically, realism is committed to the idea that theoretical claims (interpreted literally as describing a mind-independent reality) constitute knowledge of the world. (Chakravartty, 2017, pp. 4–7)¹⁰

Rival definitions of scientific realism, on Chakravartty's understanding, can illuminatingly be understood as differing mainly in how they emphasize these three components. For our purposes, the epistemological component is of primary interest. In the form of a skeptical challenge, underdetermination is levied to attack the notion that our theories can provide us with knowledge about the world. How does it do that?

Just as there is no single accepted definition of its target, so there are different reconstructions of the skeptical argument from underdetermination itself. These vary considerably, depending on how the target is defined and what kind of underdetermination the challenger considers to have established. There is, however, something like a core structure of the argument.¹¹ That structure is as follows:

(P1) *Extensional Equivalence:* There are empirically equivalent rivals to even our best scientific theories.

(P2) *Evidential Equivalence:* Some of those empirically equivalent rival theories are equally believable.

(P3) *Withholding of Belief:* When facing two equally believable rival theories, belief should be withheld.

therefore

(C) Belief in even our best scientific theories should be withheld.

Before we transfer the argument, let me say a bit more about the premises and the way the argument is framed.

The first premise is a specific version of the underdetermination thesis, focusing on our best theories. To recapitulate, two theories are *empirically*

¹⁰ Compare also Ladyman (2002, p. 158) for a definition of realism in terms of these three components.

¹¹ Compare Kukla (1998, p. 58), Psillos (1999, p. 164), and Douven (2008, pp. 294–295). So as not to complicate matters too much, I will work with this structure when transferring the argument. Readers who have their own preferred version of the argument might try to substitute it here and see whether the adaption to the moral realm works for their version, too.

equivalent if and only if they entail the same set of propositions about what we can observe. They are *rivals* if and only if they are radically different and thus cannot be true at the same time. Many commentators think that this premise is uninterestingly true because it does not require more than *deductive* underdetermination. Deductive underdetermination, as we have noted, is not difficult to establish, because logic alone does not restrict our choice of theories very much.

The second premise stipulates that the underdetermined theories are equally believable. It thus goes beyond deductive logic and takes into consideration ampliative forms of reasoning that become relevant when we consider the plausibility of different theories. Obviously, this premise is more difficult to establish. As we shall shortly see, defenders of the skeptical argument basically have two options, but since these are not compatible, they have to make up their mind.

The third premise introduces the skeptical element, relating to our justification in believing what a theory tells us. I have expressed the skeptical element in the form that belief in theories should be withheld. Of course, some people don't think that the argument from underdetermination warrants such a skeptical conclusion, instead advising us to remain *steadfast* in our belief in one theory. Moreover, what exactly the appropriate epistemic stance regarding the competing theories is supposed to be is a matter of debate even among skeptics. Instead of withholding belief, one might also think that belief in the theories is rendered arbitrary or that we should merely downgrade our credences. What all skeptical arguments have in common, though, is that underdetermination negatively impacts our justification to believe in the theories' truth. Since skepticism is an epistemological position, this entails that the underdetermination argument is only a threat to scientific realism if that position does indeed include an epistemic component. Most philosophers of science do accept that there is at least a weak epistemological component to realism.¹² Chakravartty even takes the epistemological component to be the most important one, when he characterizes scientific realism as:

[...] [A] positive epistemic attitude toward the content of our best theories and models, recommending belief in both observable and unobservable aspects of the world described by the sciences. (Chakravartty, 2017, p. 2)

¹² For a textbook definition of scientific realism that includes both the metaphysical and the epistemological components, see Bortolotti (2008, p. 96).

Having introduced the skeptical argument from science, we next need to transfer it to the moral realm. To do this, we follow the analogy outlined in Chapter 2 and substitute *deontic* for *empirical* and *moral theory* for *scientific theory*. This yields:

(P1*) *Extensional Equivalence*: There are deontically equivalent rivals to even our best moral theories.

(P2*) *Evidential Equivalence*: Some of those deontically equivalent rival theories are equally believable.

(P3*) *Withholding of Belief*: When facing two equally believable rival theories, belief should be withheld.

therefore

(C*) Belief in even our best moral theories should be withheld.

We thus arrive at a skeptical argument for the moral realm that has the same structure as the one from science. Of course, the argument is by no means universally accepted in the scientific realm, and even if it were, simply transferring it from science to ethics does not guarantee that it will be successful in the latter domain as well. Just as there are objections to the argument in science, there are potential objections in ethics, two of which I will consider in the next section. But before doing so, there is a disanalogy between science and ethics that needs to be addressed.

In Chapter 5, I observed a major difference between scientific and moral theories as it relates to the local–global distinction. Scientific theories are typically *local*. They account for different subsets of the evidence in different domains, for example, biological theories for biology, chemical theories for chemistry, and so on. As a result, what philosophers of science mean when they claim that the best scientific theories are underdetermined is that there is a multitude of theories from very different domains of science for which there are empirically equivalent alternatives. In contrast, moral theories, as they are widely understood, are global (or near enough). According to the completeness condition, they specify deontic verdicts for (close to) all cases where there is a clear answer. As noted, that does not mean that they yield verdicts for every choice-situation, since there might be situations where there is no single correct answer. However, something close enough is most likely true. The great traditions of moral theorizing try to give answers not just to some specific subfield of morality but to what we are morally supposed to do *tout court*.

This has important consequences for what skeptics need to prove to establish the premises of the skeptical argument in ethics. Scientific realists need not feel threatened by just one example of underdetermination, even if it concerns a widely accepted scientific theory. The case is different in ethics, however. If moral theories are indeed complete, in the sense of covering (close to) all verdicts, the skeptics do not need to prove underdetermination for many different theories from many different domains. Instead, they only need to show that *the* best moral theory has at least one deontically equivalent and equally believable rival. Accordingly, the burden of proof is less heavy for moral skeptics, at least as regards the sheer number of theories for which they have to prove underdetermination.

This, in turn, influences the strategy that skeptics should preferentially pursue to establish their case. Many skeptics in science have felt the need to make use of the algorithmic strategy in order to prove underdetermination to be a sufficiently pervasive phenomenon to challenge realism. This makes perfect sense considering the number of theories for which underdetermination has to be proven in science. If the situation were analogous in ethics, skeptics might think that they have to rely on the algorithmic strategy as well. This, however, would make establishing the overall argument much more challenging. Remember that although we have been able to establish a pervasive thesis of deductive underdetermination, I have also noted that no such wholesale argument has been established for the ampliative thesis. However, if moral theories are global, skeptics do not have to rely on this strategy. It suffices to show that choice between the best theories is underdetermined. This means that the case for establishing the skeptical argument in ethics might already be quite far advanced. *If* Parfit has identified the best theories, and if those theories are indeed equally believable, convergence between them would suffice to establish the first and second premises. There would be no need to prove underdetermination for all other moral theories.

Granted, this is still a big *if*. Since I do not claim to have identified the best moral theory (nor that Parfit has done so), I cannot claim with any certainty that that theory has deontically equivalent rivals with which it still radically disagrees. Instead, what follows is much more speculative and (partly) based on preliminary results from normative ethical theorizing. This makes the following argument dependent on whether the developments in normative ethics are indeed on the right track. Still, the preceding considerations about the completeness of moral theories suggest that we might not be as far from establishing the first two premises of the skeptical argument as might be first thought.

There are other objections, however. It is these to which we turn next.

6.4 In Defense of the Argument

Transferring the skeptical argument from science to ethics is only half the story. We also need to consider the plausibility of its premises in the moral realm as well as defending the argument against potential objections. The first premise, I will assume, has been established. If the technical side of consequentializing and deontologizing is successful, extensionally equivalent theories can be produced for all (plausible) moral theories. This should not come as a big surprise if only for the fact that the first premise only requires deductive underdetermination, which, as the discussion in the philosophy of science has shown, is easy to establish. The controversial premises are the second and third ones. I therefore restrict my discussion to two objections to these premises that are inspired by similar concerns in the philosophy of science. In both cases, I will argue that the defense against these objections is at least as plausible as, if not more plausible than, its counterpart in the philosophy of science. I do not claim that these are the only objections to the argument. However, I do think that they are among the most important ones. If they can be responded to successfully, the skeptical argument in ethics is on relatively firm ground. Since the skeptical argument has commanded considerable attention in the philosophy of science, and it is at least as plausible in ethics, I thereby hope to show that this argument deserves to be taken very seriously in ethics as well.

Against Evidential Equivalence

Let us start with the objection to the second premise. Are there deontically equivalent rivals to our best moral theory (or theories) that are also equally believable? Not necessarily. The mere fact that theories are extensionally equivalent does not, one might think, entail that they are also equally believable. This is a message that philosophers of science have taken some pains to drive home – most forcefully perhaps Laudan (1990). Mere deductive underdetermination, Laudan argues, is a situation we can establish for almost any theory for the simple reason that deductive logic alone does not set high standards for what the alternative theories are. Even creationists, remember, can claim deductive underdetermination between their theory and the best scientific alternative. Yet no scientist is worried by this. The interesting conclusions follow from *ampliative* underdetermination.¹³

¹³ Compare also Bonk (2008, pp. 38–39) and Carrier (2011, pp. 190–191).

The underlying problem, according to Laudan, is one of methodology. When defending underdetermination, Laudan thinks, many theorists employ a *naïve hypothetico-deductivist* understanding of theory confirmation, according to which there is no restriction on theory choice except having to entail the correct predictions.¹⁴ On this understanding, whenever some piece of evidence is entailed by a theory, it also supports belief in that theory, and whatever amendments to a theory are logically possible are also rationally acceptable. However, Laudan objects, the scientists' toolkit includes more than just deductive logic. Theoretical virtues, such as simplicity or non-ad-hocness, allow us to distinguish between theories with the same empirical content. When we consider these, ampliative forms of underdetermination become increasingly unlikely (especially the broader versions like the *non-uniqueness* and the *egalitarian* versions).

I think that we should take this challenge very seriously in ethics as well.¹⁵ If scientists' toolbox is not restricted to deductive logic, the same might hold true for ethicists. Moral theories, as we have seen in Chapter 5, are often thought to have theoretical virtues as well, and these might be brought in to decide between theories that are deontically equivalent. If they do indeed tip the balance for one theory, the skeptical argument might fail since in such a case we are justified in believing in this theory.

What options are there for the defender of the skeptical argument? The philosophy of science presents us with two. The first reply accepts the idea that theoretical virtues could in principle tip the balance for one theory but attempts to cast doubt on whether such virtues will indeed be sufficient to decide the specific case at hand. Granting that different theories might showcase theoretical virtues and that these might in principle be decisive, this need not always be the case. It might be that different virtues are exhibited by different theories. If so, we'd need to find out which ones are more important. But that might prove very difficult. We would need to find a way to weigh these virtues against each other, and even if we can do so, the virtues might turn out to be of (approximately) equal weight and it might therefore not be clear which theory is to be preferred. More radically still, different theoretical virtues might turn out to be *incommensurable* in principle, as Tulodziecki (2012, p. 326) has recently argued with regard to scientific theories. Incommensurability has often been interpreted in terms of a lack of common understanding. On this

¹⁴ See also Norton (2008, pp. 26 ff.) who criticizes a general neglect of the literature on induction and confirmation in early discussions of underdetermination.

¹⁵ Van Dyck (2007, pp. 12–15) goes so far as to claim that this is the only objection against underdetermination that realists should really want to make.

interpretation, the virtues would be incommensurable because scientists employing different theoretical frameworks do not sufficiently understand each other's frameworks and are therefore unable to arrive at a common assessment of the relative value of the virtues of their theories. This interpretation of incommensurability is rather dubious, however, as in most cases, a true lack of understanding doesn't seem to be the real problem. The more plausible way to understand incommensurability is in terms of standards. What counts as a good explanation might be decided by standards that are built into theoretical frameworks themselves. Thus, in one framework, simplicity might count for more, whereas another theory values predictive fruitfulness more highly, and so on. Since the assessment is based on the framework itself, no common standard can be appealed to in assessing the relative value of a theoretical virtue, and we are thus presented with a case of incommensurability.

The second reply tackles the objection head-on by explicitly denying that theoretical virtues are any indication of the truth of a theory. The most famous contemporary proponent of this strategy is van Fraassen (1980, pp. 87–96). In van Fraassen's view, what ultimately speaks to a theory's veracity is empirical adequacy alone. In contrast, theoretical virtues only provide us with a *pragmatic* criterion for theory choice. They might explain why a scientist accepts a theory in practice, but they do not give us any additional reason to believe the theory to be true. Hence, even if one theory were to be clearly preferable regarding its theoretical virtues, that would not entail that we have more reason to believe it. I will have much more to say about this strategy in Chapter 7. For now, it is only important to note that there are two possible strategies to defend the skeptical argument against the charge that extensionally equivalent theories fail to establish an interesting form of underdetermination. Note also that the two strategies are not combinable. One accepts that theoretical virtues are relevant to the truth of a theory, whereas the other denies this. Defenders of the skeptical argument thus have to make up their mind. The same is true for ethics. Skeptics in ethics have to make a decision about how to counter the objection if they want to defend the underdetermination argument for the moral realm. Fortunately, I think that no matter how they decide, the case looks more promising for them than for skeptics in science. The reasons for this are as follows.

First, note that the second strategy is more aggressive. It flat out denies one of the premises of the objection, that is, that the theoretical virtues of theories are indicative of their truth. Obviously, this makes the second strategy much more controversial. As critics of van Fraassen have pointed out, many scientists do in fact accord much importance to ampliative forms

of reasoning, such as *inference to the best explanation*.¹⁶ When different explanations for some phenomena are available, we routinely choose the most simple, or the one that best accords with other theories, and so on. This, critics have argued, amounts to a principle of rationality that should also guide theory choice in science. Denying the importance of theoretical virtues might thus come at the very high price of contradicting one of our basic rules of rationality.¹⁷ The good news for skeptics in ethics is that it might very well not be necessary to make use of this aggressive rejoinder at all. This brings us back, once more, to the issue of *global* vs. *local* theories. One of the main reasons why proponents of the skeptical argument in science deny the relevance of theoretical virtues is the sheer number of theories for which they would otherwise have to claim that theoretical virtues are indecisive. If skeptics in science did not deny in principle that theoretical virtues might break the tie, they would have to prove their ineffectiveness for all the best scientific theories in isolation (i.e., for theories in biology, chemistry, and so on). Instead of taking on that Sisyphean task, it is tempting to deny the importance of theoretical virtues from the outset. Yet this, as we have seen, comes at a price. The situation is different in ethics. Due to the *global* nature of moral theories, what needs to be shown is only that theoretical virtues fail to decide the case between our best theory and some deontically equivalent rival. That is certainly more easily done than showing the same thing for all the best theories of science. Thus, in ethics, there is less reason to adopt the second, more controversial, strategy in the first place.

Granted, at this stage, we cannot say for sure whether the first strategy succeeds since we do not know what the best moral theory is. However, as we saw in Chapter 5, different virtues do indeed come down on different sides of the traditional theoretical trenches. For example, consequentialists have argued that their theory can accommodate the Compelling Idea. Deontologists can counter that their theory has the advantage of Independence from Axiology. So far, none of these references to non-deontic criteria has decided the debate, and it is less than clear that they will do so once we have identified the best theories.¹⁸ Thus, the first reason to think that the skeptical argument is on at least as good a footing in ethics as in

¹⁶ See especially Harman (1965), who coined the term.

¹⁷ Churchland (1985, p. 42) goes so far as to claim that super-empirical virtues are more fundamental than empirical adequacy.

¹⁸ Schroeder (2017, p. 1479) notes that the problem is one that we face in respect of many philosophical issues. He argues that weighing the CI versus a more intuitive theory of the good proves highly difficult because it is generally difficult to make philosophical arguments about the relative plausibility of different considerations. Schroeder (2017, p. 1479) also thinks that this might

science is that the less controversial way to counter the objection seems more promising in ethics.

The second reason why the case looks brighter for skeptics in ethics is that even if the first strategy fails, the second, head-on, defense looks more promising in ethics than in science. In science, it seems that virtues such as simplicity have in practice often made the difference, and reliance on them has been corroborated retrospectively when further evidence has become available. A famous example of this comes from astronomy, where Kepler's elliptical model did away with epicycles, thereby allowing for a much simpler picture of planetary motion. No such dynamics can be seen in ethics. Consequentialists, for example, have often argued that their theory is simpler, yet this has not led to rival theorists giving up their theories. Instead, deontologists have doubted whether simplicity is even desirable considering that more complicated theories might be more phenomenologically adequate. The lack of effectiveness of these kinds of arguments throughout the history of ethical theorizing might thus be taken as an indication that such theoretical virtues are indeed not indicative of the truth of moral theories. This is very quick, of course. Someone like van Fraassen would deny that there can be a relevant difference here because theoretical virtues are not indicative of the truth of a theory in science either. We will consider his view in much more detail in Chapter 7. Here, I only want to make the more modest argument that when compared to science, it seems that theoretical virtues play an even smaller role in theory choice in ethics, and this provides a stronger indication that theoretical virtues are indeed irrelevant when it comes to adjudicating which moral theory is the correct one.

Summing up, we can state that when it comes to the first kind of objection to the skeptical argument, skeptics in ethics are in a position that is at least as good as, if not better than, the position of skeptics in science.

Against Withholding of Belief

The second objection is to the third premise. This premise introduces the skeptical element. According to it, the correct reaction in the described case of moral underdetermination is to withhold belief. But is this really the (only) rationally acceptable reaction?

be a reason why these kinds of considerations have yet to be explored in detail in the literature; commentators might simply not have known how to productively engage with them.

The idea that (persistent) disagreement in ethics should have a negative impact on the justification of our moral beliefs is, of course, not unheard of. Here are two examples:

[...] the denial by another of a proposition that I have affirmed has a tendency to impair my confidence in its validity [...] And it will be easily seen that the absence of such disagreement must remain an indispensable negative condition of the certainty of our beliefs. For if I find any of my judgements, intuitive or inferential, in direct conflict with a judgement of some other mind, there must be error somewhere: and if I have no more reason to suspect error in the other mind than in my own, reflective comparison between the two judgements necessarily reduces me [...] to a state of neutrality. (Sidgwick, 1981, p. 342) (quoted in Kelly (2005))

and

If we had strong reason to believe that, even in ideal conditions, we and others would have deeply conflicting normative beliefs, it would be hard to defend the view that we have the intuitive ability to recognize some normative truths. We would have to believe that, when we disagree with others, it is only we who can recognize such truths. But if many other people, even in ideal conditions, could not recognize such truths, we could not rationally believe that we have this ability. How could *we* be so special? And if none of us could recognize such normative truths, we could not rationally believe that there *are* any such truths. (Parfit, 2011b, p. 546)

These are weighty voices. Still, they don't prove the case, of course. Instead, the underlying issue can be framed in a way that ties into a wider recent debate, the debate about *peer disagreement*. Here, philosophers have been asking what the rational reaction should be to cases where we disagree with *epistemic peers*, that is, people who are as knowledgeable as ourselves, have access to the same evidence, and are generally equally reliable as ourselves in reacting correctly to the evidence.¹⁹

Two opposite positions can be distinguished. *Conciliationists* about peer disagreement think that if we learn that we have epistemic peers who disagree with us about some topic, we should withhold belief. *Steadfasters*, in contrast, hold that the mere fact that someone disagrees with us does not impact our justification in our own belief.²⁰

¹⁹ There are some issues regarding how best to characterize the notion of *peerhood*. See Rowland (2017, p. 2) and Matheson (2015). I gloss over this in the remainder.

²⁰ This is a great simplification of the debate. As Christensen (2009, p. 756) suggests, we can rather think of these two positions as the endpoints of a spectrum. For example, instead of thinking that peer disagreement should lead us to withhold belief, we might think that it only requires that we lessen our degree of certainty in a belief. The reason I will not consider such views is that I cannot make much sense of them with regard to the case at hand, that is, peer disagreements regarding

As peer disagreements can arise in almost any domain of knowledge inquiry, the issue is one of general epistemology. The debate around this issue has quickly become voluminous, and arguments have become so sophisticated that I cannot attempt to do justice to them here.²¹ Yet, for our purposes, we are interested in a more restricted question: Is the conciliationist reaction in the case of moral underdetermination at least as plausible as, if not more plausible than, in the case of scientific underdetermination? This question, I posit, is very likely to have a positive answer.

The main reason for this has to do with the idea of evidence transcendence. In a case of underdetermination, it seems likely that the only way in which one could claim to know that one of the theories is true is if one's knowledge transcends the evidence, in the sense that other people might forever be unable to attain that knowledge. But this notion of evidence transcendence is much more difficult to defend in ethics than in science, as has been convincingly argued by Wright (1992, 1995). Moreover, the specifics of the case of moral underdetermination, that is, that it is evidence transcendence about moral explanatory claims only, makes the overall argument even stronger.

So here is the basic argument. Many disagreements in ethics (and elsewhere) are not particularly threatening to the assumption that we can have knowledge because there are clear factors to which we can appeal in order to explain them. One set of such factors has to do with *cognitive shortcomings*, such as biases. If we know that someone is a (male) chauvinist, learning that they disagree with us about issues of gender equality will not be a ground for us to revise our own views. The second set concerns knowledge about non-moral facts. If we learn that someone thinks that animals have neither the capacity for pain nor self-consciousness, it will not unsettle us to hear that they think there's no problem in gratuitously killing animals. The third class concerns issues of vagueness. Perhaps reasonable people can disagree about when it is justifiable to terminate a pregnancy because there is vagueness in the ascription of personhood to the fetus. Yet other disagreements are not of such nature. Disagreements can persist about issues where there is no indication of vagueness being a problem and

large groups of philosophers (Kantians, consequentialists, contractualists). It seems quite obvious to me that the correct reaction here is *not* a fine-grained splitting of credences (e.g., that we become 30% Kantians, 25% consequentialists, 15% contractualists, and so on). However, I also accept, with Christensen (2007, p. 213), that there is no uncontroversial view about the relationship between graded belief and all-or-nothing belief. In general, my discussion here follows Rowland (2017, p. 3), who also defines conciliationism in terms of withholding belief and thinks that many ethicists do so too, because these are simply the more interesting cases when it comes to moral epistemology.

²¹ Feldman and Warfield (2010) provides a good overview.

between people who have spent their lives thinking about these issues, who are beyond the doubt of simple prejudice, who show awareness of all the relevant evidence, and so on.²² These disagreements are potentially more threatening. The fact that the people who should know best disagree might make us doubt that there is any reliable way of arriving at knowledge in the domain.

To be sure, disagreements among peers are by no means limited to the realm of ethics. Shafer-Landau (2012, p. 328) gives the examples of physicists disagreeing about string theory or archeologists disagreeing about the purpose of some ancient site. Still, the fact that such disagreements persist calls for an explanation. And here, it gets tricky for ethicists. For to remain steadfast in such a case, we would have to assume that we have knowledge that *transcends* the evidence in the sense that our peers might forever be banned from achieving the kind of knowledge. Yet this, Wright argues, is an extremely implausible position to take in the ethical domain. Comparison to evidence transcendence in science makes clear why. In science, it often makes sense to think that we can have knowledge about some state of affairs that a peer might never be able to attain. Why? Because we can offer a plausible explanation for this differential situation. In the scientific case, evidence transcendence can be due to *contingencies of epistemic opportunity*: our measuring instruments or spatio-temporal situations not allowing us to attain knowledge about a state of affairs that others might be able to attain. In contrast, ethicists arguably have no such explanation at hand. The disagreements in ethics seem to have nothing to do with such contingencies of epistemic opportunity. We can quite safely assume this because our moral claims are much more intimately related to our everyday lives. They do not pertain to remote or otherwise inaccessible events, which some of us might be barred from accessing. Instead, this kind of evidence transcendence has a more malign form than the one in science. This should make us withhold our own beliefs.

Wright (1992, p. 82) offers this argument in a very general way, writing that, as far as ethics is concerned: “evidence-transcendence is simply not in view.”²³ I tend to agree with this. The idea of evidence transcendence in

²² Indeed, for most of us, the situation is even worse: Some of the people with whom we disagree are clearly our moral epistemic *superiors*.

²³ Note that Wright thinks that ethicists can nevertheless hold out for some notion of truth if they give up a representational understanding and instead understand truth in terms of superassertability. I gloss over this since my interest here is only in Wright’s argument against evidence transcendence.

ethics is problematic *full stop*. However, for our purposes, the argument can be strengthened considerably by restricting it to the present case of moral underdetermination. By doing so, we can block two of the most promising objections.

The first objection aims directly at Wright's argument. Tersman (1998, pp. 360–362) has argued that Wright is too quick to assume that ethicists have to appeal to evidence transcendence to explain their disagreements. This is because, apart from factors like biases and lack of non-moral knowledge, ethicists might also point to the fact that others have false *moral beliefs* in order to explain their disagreements. In particular, if we assume that some ethicists accept false ultimate principles, this may explain why they are led to disagree with the correct verdicts in a particular case. Now, without assessing whether Tersman's move might work in the case he describes, it should be clear that this move is not available in the case of moral underdetermination. The simple reason for this is that in a case of moral underdetermination, we are simply stipulating that there are no disagreements about particular cases remaining. The only disagreements that remain are those about the fundamental explanatory claims. But to explain why these persist, the realist cannot fall back on disagreements about other (more fundamental) moral principles. Thus, the fact that the disagreements are about the fundamental explanatory claims blocks this counter to Wright's argument.

Granted, one could simply insist that one has direct intuitive knowledge about which moral explanation is true. But note what one would have to claim here. One would have to claim that one has knowledge about the correct explanations in ethics while at the same time conceding two things. First, that the explanation does not entail any difference when it comes to the deontic verdicts that follow from the theories. The explanation thus turns out to be unfalsifiable by the deontic evidence, that is, our considered judgments and intuitions about particular cases. That is what the first premise of the skeptical argument entails. Second, one would have to concede that the preferred explanation does not turn out to be better in any way that can be framed in terms of theoretical virtues, that is, it isn't simpler, and so on. That is what the second premise of the argument entails. The objector to the skeptical argument thus asks that we simply take people at their word when they tell us that they know what the correct explanations in ethics are, even though their peers disagree, they aren't able to point to an overall advantage of their preferred explanations, and the truth of their explanation is unfalsifiable by other moral intuitions we have. This, as I anticipated in Chapter 2, is a bridge too far.

The second objection is to the overall argument for conciliationism. Even if evidence transcendence seems like a problematic feature in ethics, potentially supporting a conciliationist view, might there not be other, even stronger, considerations that speak against conciliationism in ethics? Indeed, some authors think that there are, and one such consideration is particularly interesting for our purposes. Conciliationism about peer disagreement in ethics might be more problematic because it violates one's autonomy. In other cases, it might be fine to change one's views on the basis that other people disagree. However, or so one might think, ethics is different in this regard. Autonomy requires that our moral judgements are our own, in the sense that they are due only to our own reasoning. Conciliationism, by requiring us to withhold a moral belief on discovering that others disagree, violates our autonomy.²⁴

Again, I am not going to take a stance on whether this argument might work in some cases of moral peer disagreement. Instead, I only want to point out that the argument is very unlikely to be successful in the case of moral underdetermination. The reason, once more, has to do with the fact that in such a case, the only remaining disagreements are of the explanatory kind. Yet conciliationism about these, I'd argue, is clearly less problematic (if it is problematic at all). Perhaps autonomy does indeed require us to make up our own minds about what we ought to do; the fact that someone disagrees with me about what I ought to do should thus be disregarded. Yet moral underdetermination is not about such disagreements. The remaining peer disagreements in ethics in the case of underdetermination are about what explains why some acts are right or wrong, not which acts are right or wrong. The remaining disagreements are thus about explanatory beliefs, and in this, they are not in a relevant way distinct from explanatory disagreements in other domains. They are about whether Kantianism, consequentialism, or contractualism provide the correct explanation. This, I would argue, cannot sensibly be framed as infringing on one's autonomy.

In sum, I have argued that the idea that the correct reaction to cases of underdetermination is to withhold belief in ethics is at least as plausible as, if not more plausible than, the same idea in science. If, as I have earlier argued, the same goes for the defense of the second premise, the skeptical argument looks to be on sufficiently strong grounds to warrant serious consideration. Of course, other objections have been leveled at the underdetermination

²⁴ See Nguyen (2010) for a critical discussion of this argument. Compare also Rowland (2017, pp. 4–5), who calls this the *no requirement of pure moral deference* view.

argument in science, and it would be useful to analyze these in more detail in order to assess whether they translate to ethics. However, at least as regards the two objections I have considered, it seems that the prospects for the skeptical argument in ethics are at least as good as, if not better than, the prospects in science. The next chapter considers what this might entail.