

A New Position in Metaethics

If the skeptical argument is indeed successful, we are facing an interesting situation. On the one hand, we are no longer justified to believe the explanations that moral theories offer for why certain acts are right or wrong. On the other hand, nothing that has been said so far (directly) impacts our justification to believe that we can know which acts are right or wrong.¹ Indeed, one might think that the fact that the best theories agree about what we should do strengthens our justification to think that these are indeed the correct actions. The position that results from the skeptical argument might thus be that whereas we are entitled to believe that we have knowledge about which acts are correct, we are not entitled to believe the explanations that moral theories put forward for why these are the correct acts. Such a position, as far as I know, has not received much attention in the literature.² This is somewhat surprising since, once more, there is a structural analogy to science. That analogy is to Bas van Fraassen's *constructive empiricism*. What is more, constructive empiricism is generally considered to be one of the most promising anti-realist views in the philosophy of science.³ It should therefore prove interesting to consider that final analogy more closely. What it unveils, I argue, is a mirror view to van Fraassen's view in ethics, which I call *constructive deonticism*.

This chapter mainly aims to introduce and classify that view. I start with a brief characterization of constructive empiricism. I then flesh out the new position in ethics, using van Fraassen's view as a foil. Next, I discuss where the new position should be put on the metaethical map. Constructive empiricism is clearly an anti-realist view in science, but, as we shall see, the different dialectical situation in metaethics makes it less clear how to think

¹ I'll consider a way in which the argument could be thought to impact the justification of the deontic content of theories as well in Section 7.3

² I have noted the possibility of such a position in Baumann (2021b, 2022) but have not elaborated on it in any detail yet.

³ Compare Rosen (1994, p. 143).

of constructive deonticism. Indeed, as I shall try to argue, the possibility of such a view might prompt us to reconsider the traditional way that the realism debate has been framed in metaethics. Finally, I consider two worries that a future, more detailed, defense of constructive deonticism would have to grapple with.

Two short remarks before I begin. First, this is the most speculative part of the book and also the one that outlines what are arguably the most controversial views. Some of them, especially those pertaining to moral explanation, are in tension with other views that have so far taken a prominent place in the book. However, I don't think that this is in any way inconsistent. So as not to endanger the lessons for normative ethics, where necessary, I have attempted to base my arguments on metaethical views that I consider to be as uncontroversial as possible. Here, I entertain some more revisionary ideas, and defending these requires me to take a stronger stance on some matters. However, this does not nullify what has been said so far; indeed, I believe that much of what I have set out so far is correct even if it should turn out that what follows is not. Second, what follows is at most a tentative endorsement of a new position in metaethics. I remain torn as to whether I should accept the position myself, and the last section of the chapter will consider two reasons for this hesitation. What I am positive about, however, is that this view merits consideration even if it should ultimately turn out to be untenable. This is, not least, due to the aforementioned fact that its mirror view in science holds such a prominent place that it would betray a lack of curiosity not to inquire more deeply into the merits of the analogous view in ethics.

7.1 Constructive Empiricism and Constructive Deonticism

Underdetermination in science has been linked to a variety of views, such as *conventionalism*, *instrumentalism*, and *social constructivism*.⁴ However, the position that philosophers of science probably associate most closely with underdetermination today is Bas van Fraassen's constructive empiricism. His outline of this view in *The Scientific Image* has been widely credited with rehabilitating anti-realism in the philosophy of science.⁵ Let us thus begin with a closer look at that view.

⁴ Compare Park (2009, pp. 116–119) for a discussion of these positions as they relate to underdetermination.

⁵ Compare Rosen (1994, p. 143) and Rosen (2005, p. 17).

The Foil: Constructive Empiricism

As an introduction to constructive empiricism, it makes sense to follow the structure of van Fraassen's own exposition and start with the two positions from which he distinguishes his view, namely *scientific realism* and *logical positivism*.

Van Fraassen's main aim is to come up with an alternative to what he considers the dominant view of his time. Scientific realism, van Fraassen holds, is the view that:

Science aims to give us, in its theories, a literally true story of what the world is like; and acceptance of a scientific theory involves the belief that it is true. (van Fraassen, 1980, p. 8)

This is just a minimal definition of scientific realism, of course. Still, van Fraassen thinks, it already includes a fundamental mistake. Scientific theories, in their explanations, typically posit unobservable entities (such as electrons). If we assume that scientific theories aim to give a true account of the world, we must, on the scientific realist's view, take these claims about unobservable entities as unveiling the true nature of the world. Yet, what kinds of unobservable entities we should assume in our theories, van Fraassen argues, is not mandated by the evidence. The empirical evidence only tells us whether a theory is *empirically adequate*, whether it *saves the phenomena*. Scientific realists, by buying into the truth of claims about unobservables, additionally commit themselves to controversial metaphysical assumptions. Yet with this, they overburden science. Their picture of science becomes too heavily metaphysical, and one of the core motivations behind his development of constructive deonticism, van Fraassen (2002, p. xiii) later makes explicit, is a revolt against metaphysics.

This last feature is something that constructive empiricism shares with another view of science, that of logical positivism. Positivists, van Fraassen (1980, pp. 3–4) tells us, are right in resisting the metaphysical presumptions of scientific realism. However, in order to do so, they rely on a theory of meaning and language that van Fraassen considers to be thoroughly refuted. In order to explain away some of the metaphysically charged aspects of scientific language, positivists attempted to prove that all that is involved in the theoretical claims of scientific theories must be firmly moored in observations. They thought that the meaning of theoretical concepts, such as electron, could be reduced to what the theory entails about observables. Yet this undertaking, van Fraassen thinks, failed completely. In contrast to positivists, van Fraassen (1980, pp. 80–83) accepts that our observational

reports are given in a language that is thoroughly theory-laden and cannot be reframed in a theory-independent way. Hence the positivists' semantic proposal must fail.

Van Fraassen concludes that just as scientific realists are committed to an implausibly metaphysical picture of science, so positivists are committed to a highly implausible view of the language of scientific theories. In his own view, he therefore aims for a less metaphysical view of science while at the same time eschewing the implausible semantic views of logical positivism. What could such a view look like? Here is the essence of what van Fraassen offers:

Science aims to give us theories which are empirically adequate; and acceptance of a theory involves as belief only that it is empirically adequate. (van Fraassen, 1980, p. 12)

This is the (short) definition of constructive empiricism.⁶ When comparing it to the definition of scientific realism above, it becomes obvious that constructive empiricism is defined in opposition to the former. While scientific realists think of science's aim as discovering an underlying reality, van Fraassen urges scientists to restrict their ambitions to the empirical adequacy of their theories. Accordingly, to accept a theory implies nothing more than believing its empirical predictions.

Van Fraassen inscribes his position into the empiricist tradition, the main rival of scientific realism. The added *constructive*, van Fraassen (1980, p. 5) informs us, is supposed to highlight that scientific activity is not really one of discovery but rather something more active. At the same time, and *pace* logical positivism, scientific explanations include claims about unobservables, and these claims, van Fraassen thinks, should be interpreted in a literal way. If a theory says that electrons exist, that does mean that it postulates the existence of some unobservable entity. Consequently, the theoretical claims of scientific theories can be true or false; they are truth-apt.

The last point might seem puzzling. If we are to construe scientific theories literally and they do entail claims about unobservables, why should we restrict our beliefs to what they tell us about observables and ignore what they entail about unobservables? This is precisely where considerations of theory underdetermination come into play. Van Fraassen thinks that if theoretically incompatible scientific theories can account equally well for

⁶ In what follows, I consider constructive empiricism to be a view or a position in the philosophy of science. In *The Scientific Image*, van Fraassen (1980, p. 5) does speak of it as a position, although one with only a "momentary name." Later, van Fraassen (2002, pp. 46 ff) frames empiricism as a *stance*. He repudiates the idea that empiricism could consist in holding any particular thesis or dogma. For the sake of simplicity, I gloss over this issue in the remainder.

all the observable evidence, this should lead us to withhold belief in what these theories tell us about unobservables. Yet the reason for this is not a semantic one as the logical positivists thought; rather, it is epistemic:⁷ Given underdetermination, we simply cannot know whether claims about unobservables are true or false and should therefore remain agnostic.

Van Fraassen does not consider this to be a problem since, in his view, whether a theory is true is just not that important for whether it is a good theory. This claim might seem even more puzzling. What does it even mean to accept a theory according to such a view? At this point, van Fraassen introduces a crucial distinction. When it comes to what we believe, accepting a theory only entails that it is empirically adequate.⁸ No belief in unobservables and no metaphysical assumptions are entailed. Yet acceptance of a theory is not exhausted by belief:

To accept a theory is to make a commitment, a commitment to the further confrontation of new phenomena within the framework of that theory, a commitment to a research programme, and a wager that all relevant phenomena can be accounted for without giving up that theory. (van Fraassen, 1980, p. 88)

Such commitments, van Fraassen thinks, matter and, consequently, it does matter which theory we accept:

Even if two theories are empirically equivalent, and acceptance of a theory involves a belief only that it is empirically adequate, it may still make a great difference which one is preferred. (van Fraassen, 1980, p. 4)

What could that difference be if it concerns neither empirical predictions nor metaphysics? Van Fraassen thinks that it is one of *pragmatics*: Working in one or another scientific framework might be much more congenial to some pragmatic goal the researcher has. For example, even though scientists know that the Newtonian framework is wrong, strictly speaking, due to its relative simplicity it can still make sense to employ it on many occasions. Yet even if we might say that the Newtonian theory is better in this sense, that emphatically does not mean that it provides a better explanation in terms of the underlying structure of the world. Instead, what we mean when we say that the theory is better is just that it better serves one of our pragmatic goals.

⁷ Compare Ladyman (2002) for the assessment that what van Fraassen attacks is the epistemic component of realism.

⁸ Note that this is not just the familiar point that we are sometimes satisfied with *tentatively* accepting an empirically adequate theory without knowing whether it is true. Instead, as Rosen (2005, pp. 15–17) helpfully points out, the point is about *full* or *ideal* acceptance.

This pragmatic view of explanation has two important implications. The first is one that we have already encountered in Chapter 6. For van Fraassen (1980, pp. 87–96), the theoretical virtues of a theory, for example, simplicity, elegance, etc., have nothing to do with the theory being true. Although such virtues might give us a pragmatic criterion for which theory to use, they do not give us any additional reason to believe that it is true. Take a virtue like simplicity. Granted that simplicity is a factor that plays a large role in our actual choice of theories for practical matters, why should that be an indication of truth? Why should we make our own preferences for simple theories the arbiter of truth? In van Fraassen's view, what ultimately speaks to a theory's veracity is empirical adequacy, that is, a theory's capacity to account for the evidence. Theoretical virtues might explain why scientists accept some theory in practice, but the criteria for acceptance of a theory are not identical to the criteria for believing in the truth of a theory.

The second, and related, implication is that explanation, in this view, is radically context-dependent. Whether a theory serves a pragmatic goal of a researcher better or worse depends on what that goal is. Whereas all researchers share the goal of coming up with empirically adequate and internally consistent theories, their goals can vary substantially in other respects. Thus, as the example above illustrates, the Newtonian framework might be preferable if our context is that of an engineer occupied with the mechanics of regular-sized terrestrial bodies. It is not preferable for the astrophysicist trying to understand issues related to space-time. What counts as a good explanation, according to van Fraassen, is a question that can only be decided in a specific context.

The New Position: Constructive Deonticism

If underdetermination in science should lead us to accept constructive empiricism, is there an analogous position that ethicists facing moral underdetermination can take? I think that the answer is positive, and I also think that this position is both interesting and novel.

Adapting van Fraassen's definition of constructive empiricism to the moral realm, *constructive deonticism* holds that:

Ethics aims to give us theories which are deontically adequate; and acceptance of a theory involves as belief only that it is deontically adequate.⁹

⁹ For the sake of simplicity in exposition, I will not quibble with van Fraassen's formulation when adapting the position to ethics. However, it should be noted already here that, for reasons to become clear in the next section, we may prefer the formulation "Ethics (or Ethicists) *should* aim to give us theories..."

The name of van Fraassen's position immediately identifies his view as a member of the empiricist tradition. The question of under which metaethical tradition constructive deonticism falls is much more complicated (there is no tradition of *deonticism* in ethics), and I will discuss it at some length in Section 7.2. In contrast, we can understand the label *constructive* in much the same way as van Fraassen (1980, p. 5) intends it to be understood in science. Just as he thinks that the aim of scientific theories should not be thought of (primarily) in terms of discovering the hidden nature of the world, the constructive deonticist argues that ethics should not aim at uncovering the ultimate grounds of rightness, the right-makers. Instead, ethical theorizing should be concerned with constructing frameworks or models that are deontically adequate, that is, tell us to do the right things.

The reason why we should restrict the aims of ethics in this way stems from moral underdetermination. The fact (if it is one) that the best moral theories are underdetermined entails that we are entitled to believe in truths about which acts are right or wrong but should suspend belief in the explanations of *why* these are the correct acts. Constructive deonticism thus eschews an overtly metaphysical understanding of moral explanation. Yet the way we arrive at this conclusion, the constructive deonticist agrees with van Fraassen, is via a skeptical, hence epistemological, argument.

Constructive deonticism is, furthermore, not semantically revisionary. Just as van Fraassen avoids the logical empiricists' implausible semantics, so constructive deonticists avoid the semantics that undergird the notational variants view of the moral traditions. Instead, constructive deonticism accepts that much of the language used in moral theorizing (and arguably our everyday moral discourse as well) is soaked in metaphysical notions. If a theory claims that we have a duty or a right, it assumes that there are duties or rights. Similarly, if a theory claims that an act is good because it maximizes value, this claim should be construed literally and as truth-apt. Constructive deonticism doesn't take issue with this. What it does take issue with is the idea that we can know which theories, and hence which explanations, are correct. Given moral underdetermination, we cannot know whether the explanatory parts of our moral theories are right and, the constructive deonticist adds, we should therefore suspend belief in moral explanations.

This view of moral explanation, readers will have noticed, sits uneasily with the grounding picture that has so far played a large role in the book. I will attend to this worry at the end of this section. But first, let me characterize the view in some more detail. The present picture of moral theorizing immediately invites a fundamental question: What does it mean,

in this picture, to accept a moral theory? Once more, van Fraassen leads the way. The answer is that, when it comes to belief, acceptance of one of the moral traditions does not commit us to anything more than believing that the theory yields the correct deontic verdicts. We should believe in what a theory entails about the rightness or wrongness of particular acts. Yet we need not believe in theories' claims about what *makes* these acts right or wrong.

An outline of such a pragmatic view of moral theories has recently been given by Suikkanen (2021). Suikkanen (2021, pp. 20–22) illustrates the view with an example involving climate change. As Suikkanen sees it, a consequentialist framework is best suited to thinking about the moral issues relating to climate change. For one, we need to evaluate policies that have large-scale implications, and this is much more difficult within a contractualist or a Kantian framework. For another, issues of climate change often involve a high degree of uncertainty and risk, which can best be dealt with using an expected value version of consequentialism. So consequentialism, according to Suikkanen, seems the most useful framework when dealing with moral questions surrounding climate change.

However, at the same time, Suikkanen denies that this means that other frameworks are wrong or that consequentialism is preferable overall. This is so because Suikkanen, accepting the results of consequentializing and deontologizing (even adding *contractualizing* to the mix), thinks that other frameworks can in principle account for all the verdicts that the consequentialist theory comes up with. Moreover, in other cases, these frameworks might be preferable to consequentialism. For example, Suikkanen argues that a contractualist framework can better deal with small-scale issues of promise-breaking. Yet, again, this advantage is pragmatic, and it does not tell us that contractualism is more likely to be true. Instead, Suikkanen proposes that we should selectively use moral theories depending on how they can better help us approach specific moral issues.

This last point is interesting because it suggests that in ethics, we might be much more flexible in the way we use our theories than most scientists are. However, one might also think that in ethics, too, we often commit to a framework. For illustration, consider the human rights regime with its international political and legal framework. Arguably, this regime has been highly successful not only politically and legally, but also in securing some of our morally most valuable needs. Now think of some person, perhaps a human rights activist or lawyer, trying to secure that some other human need, which at present is not the content of a human right, will be fulfilled more effectively in the future. When asked why this human need should

be protected, we would expect this person to put their answer in terms of human rights. We would thus assume that this person will tend to think about and evaluate future cases using the conceptual framework that their already preferred theory employs. Furthermore, we will typically also think that this person thinks that doing so – connecting the fulfillment of the new need to the already existing political and legal regime – will be effective. We will expect that the person thinks that the best way forward is to enshrine that need by making it the content of another human right.

At the same time, however, it does seem quite clear to me that this fact, the higher likelihood of successful protection of a need when put in terms of human rights, does not at all speak for the truth of a moral theory employing human rights. What is more, it would not seem wrong to me if our human rights activist or lawyer were to say that they are agnostic about whether human rights actually exist. What this person should believe, if they pursue an honest goal, is that the need that will be protected by that human right is really one worth protecting. Thus they should believe that the human rights framework will have the correct implications about what we (individuals or, in this example, states and other organizations) should do. Yet, for this, they don't need to commit to any metaphysical views that attach to the notion of a human right.

If this reasoning is correct, it might provide evidence for the constructive deonticists' claim that acceptance of a moral theory entails no other beliefs than that the theory leads to the correct verdicts. One theory might prove much more promising when it comes to achieving a specific goal. Yet this advantage is solely pragmatic. At the same time, the constrictive deonticist can claim that accepting a moral theory, just like accepting a scientific theory, involves more. It means that we work in a certain tradition, employ its conceptual framework, and so on.

Such a pragmatic view of moral explanation has two further implications that it shares with the analogous view in science. First, the constructive deonticist accepts that the theoretical virtues that moral explanations might have do not give us any additional reason to think that these theories are also correct. The fact that utilitarianism is relatively simple might be a reason to prefer it in some contexts, but it does not mean that it is more likely to be true. The fact that rule-consequentialism is more readily applicable than act-consequentialism (since it does not ask us to do the math for particular actions) might make it preferable in everyday life, but does not mean it is more likely to be true. The usefulness of a theory for our pragmatic purposes is no indication of that theory being correct. Second, and relatedly, moral explanation comes out as radically context-dependent. We adopt a specific

moral framework in pursuit of a specific goal, and in the context of that goal, the framework might indeed be advantageous. But it need not be in another context. To use yet another example, one view put forward in the consequentializing debate is that putting the deontic content of non-consequentialist views in a consequentialist framework is advantageous since it makes these views amenable to decision theory.¹⁰ Granted, this can be a clear advantage. However, it is only an advantage in certain situations, for example, if we have the goal of applying our non-consequentialist theory in situations with imperfect information. Tellingly, Portmore (2022) calls this *pragmatic consequentializing*.

Summarizing what has been said so far, the constructive deonticist thinks that the aim of ethical theorizing consists in constructing frameworks that lead us to do the right acts, and accepting a specific moral theory need not entail more than thinking that it yields the correct verdicts on these acts. We should therefore withhold belief in what moral theories tell us about the (ultimate) right-makers and chose our moral theories based on how well they serve our pragmatic goals. This, in itself, is not unheard of in ethics, as is evidenced by the fact that Suikkanen employs such an understanding. However, if we combine this view of moral theory with the idea that we can nevertheless have knowledge of what the correct verdicts about particular acts are, we arrive at what, I think, constitutes a new perspective in metaethics. We will shortly see how constructive deonticism compares to other metaethical views.

But before we do this, let me address some worries. Some readers might feel a certain discrepancy between my insistence on the explanatory dimension in the earlier part of the book and the skeptical attack on it in the latter part, culminating in this revisionary view of moral explanation. Specifically, how does what I have said just now fit with the grounding picture of moral explanation that has played a large part throughout the book? The feeling that there is some friction here is certainly not unjustified. However, I do not think that the book's overall argument is inconsistent.

This is so because all that is needed for the arguments in the earlier parts of the book to work is a rather minimal understanding of the notion of grounding. Such an understanding has been defended by Dasgupta (2017), among others.¹¹ Dasgupta argues that grounding need not have any heavily metaphysical implications. We need neither take ground to

¹⁰ Compare Lazar (2017).

¹¹ See Kovacs (2023) for a discussion of other deflationary accounts of grounding.

correspond to any part of reality, nor assume that it is metaphysically primitive. Instead, Dasgupta argues for a significantly deflated conception. Still, that conception does important work. One way it does so is especially interesting for our purposes. Dasgupta writes:

So put, the grounding thesis may sound trivial, hardly worth stating. Who would deny that questions of explanation divide logical space at important joints? But here, as in life, what's worth saying depends on one's audience. Suppose someone insists that the question of human origins concerns the meaning of "humanity," of whether it can be analyzed in design-free terms. Then one should reply that that distorts the issue; that the issue concerns the world and not our concepts; that Darwin's contribution had nothing to do with the meaning of "humanity" and everything to do with a profoundly brilliant explanation of how complex biological life came to be. (Dasgupta, 2017, p. 77)

This passage is kindred in spirit to what I take the grounding picture to do in ethics. It indicates why the notational variant view seems to miss the point. The grounding picture of moral explanations delineates important distinctions. To do this, it need not be wedded to any metaphysically heavy picture of moral explanation. The notion of grounding indicates that what is at stake are explanatory questions, in contrast to questions of meaning. But how we understand that notion of explanation is left open. As Dasgupta (2017, p. 76) surmises, use of the notion of a metaphysical ground does not commit us to a specific theory of explanation, in just the same way as use of the notion of a causal ground does not commit us to any theory of causal explanation (such as the DN model).

Hence, at least on an understanding of moral grounding like Dasgupta's, there is no inconsistency in my overall argument. The case against Dreier's notational variant view can be based on a deflationary account of moral grounding that is compatible with the pragmatic view of moral explanation that I have outlined above.¹² At the same time, as I will discuss in the following section, we should acknowledge that most ethicists have a less deflationary view of moral explanation. A metaphysically more robust understanding of the grounding model is thus probably semantically more accurate. But granting this does not mean that one cannot, at the same time, be skeptical about the prospects of such a metaphysically heavy grounding view of moral explanation. Indeed, positions of this kind (semantic agreement with a metaphysical understanding of moral discourse paired with

¹² Indeed, the same point can arguably be made without any reference to a metaphysical notion like grounding or to hyperintensionality. See Depaul (1987).

epistemological skepticism about this understanding) are not at all unheard of in ethics, as will become clear in the following section.

Finally, readers might think that accepting the pragmatic view of moral explanation contradicts the way I outlined the skeptical argument in the previous chapter. When I mentioned van Fraassen's view of explanation in the defense of the skeptical argument in ethics in the last chapter, I said that the fact that we need not necessarily rely on such a controversial view counts as an advantage of the skeptical argument in ethics vis-à-vis the skeptical argument in science. Isn't this in contradiction with constructive deonticism now advancing that very view? I don't think it is. The skeptical argument and its defense are formulated in a maximally general way in the preceding chapter. Skeptics about moral explanation who do not share the van Fraassen-inspired view of moral explanation thus have the described advantage of not having to commit to this controversial view about the pragmatic character of moral explanation. By contrast, constructive deonticists do. But that is no contradiction. It just means that when we accept a more specific theory, that usually commits us to more specific views on specific matters – in this case, moral explanation – and these views are usually more controversial. As I have furthermore observed in Chapter 6, I also think that the idea that theoretical virtues play no role in ethics is more plausible than in science. Thus, *some* advantage of the moral version of the skeptical argument is preserved.

7.2 Classifying the New View

With the characterization from the last section in mind, we can next turn to classifying constructive deonticism among the families of metaethical traditions. For reasons of space, this will have to be a very schematic discussion where I can only discuss issues in very general terms and based on very rough distinctions. However, I hope that this will suffice to emphasize what I consider to be novel about constructive deonticism as well as bringing out in what way classifying the new position might pose difficulties for our classificatory scheme.

Constructive Deonticism as a Form of Anti-Realism

Since constructive empiricism is an anti-realist position in the philosophy of science, it stands to reason to ask first whether its mirror position in ethics should be classified as an anti-realist position in ethics as well. Since anti-realist positions in metaethics abound, I won't be able to do direct

comparisons. Instead, I will restrict myself to a discussion of whether and how constructive deonticism fits into what is perhaps the most general distinction: that between expressivism and error theory.¹³

Let us start with expressivism because it is the more straightforward case. Expressivism, we might say, is to constructive deonticism what logical positivism is to constructive empiricism.¹⁴ Expressivism shares with constructive deonticism a general hostility toward what both conceive as a too heavily metaphysical picture of morality that is manifested in moral realism. But the two differ when it comes to semantics. Different expressivists do of course disagree about semantics themselves. But one core element is that they are not *factualists* about moral discourse. They don't take the fundamental function of moral statements to be one of representing some state of affairs. Moral statements don't try to describe some kind of moral reality and they are not, in any substantive way, truth-apt.¹⁵ Instead, the core function of moral discourse is to express some sentiment, command, or plan. In contrast, according to constructive deonticism, we can take the seemingly factual claims in ethics at face value. Moral discourse is factualist in that it attempts to represent moral states of affairs, and its statements can, in principle, be truth-apt in a familiar, non-deflationary, sense of truth. That constructive deonticism and expressivism differ in this way seems only logical when we recall that what brought us to the novel position was an epistemic, not a semantic misgiving about moral theorizing. Since I have already said quite a bit about this point, I won't repeat it here.

Instead, let me focus on one area where there can be considerable overlap, that is, the account of moral explanation. We have already gotten a glimpse of this in Chapter 2, when I mentioned that Baker (2021) has criticized the grounding view of moral explanation as being biased against expressivism, proposing a metaphysically less heavy view as an alternative. Yet the parallel goes deeper still. Baker thinks of his alternative view of explanation as an explicitly *pragmatic* one, citing van Fraassen as one prominent inspiration.

¹³ Compare Shafer-Landau (2012, pp. 306–319) and Chrisman (2016, p. 12).

¹⁴ This is not, of course, to imply that expressivism is a (moral) form of logical positivism or that the two are equally (im)plausible. It is just an observation about the relative positions of these views in the logical space of theories.

¹⁵ The clarification is needed because quasi-realists do think that moral statements are truth-apt, albeit on a deflationary understanding of truth. Definitions are further complicated by what Dreier (2004) calls *creeping minimalism*, that is, the possibility that we can also give minimalist definitions of other concepts, such as “facts” or “representation,” further muddying the waters. I will gloss over these complications for the remainder.

On Baker's view, moral explanations are unifying generalizations. We explain why a specific injunction – lying is wrong – holds by subsuming it under a more general injunction – treating people as mere means is wrong. When we do this, we don't pick out a metaphysical dependency relation. Instead, unifying generalizations, on Baker's view, are explanatory because they respond to a particular interest we have in normative discourse. Normative discourse is supposed to facilitate social cooperation. To do this, we make commitments to act (and evaluate actions) in certain ways. Yet to do so in an effective way, we cannot enumerate all the particular ways we intend to act. Instead, we need to agree on unifying principles. To know that an act falls under such a principle is thus explanatory because it answers to an interest we have. Furthermore, knowing that some injunction (not to lie) falls under a more general one (not to treat people as mere means) can help alleviate the doubt that a person insisting on the less general injunction is doing so just for self-interested reasons. Knowing this once more facilitates social cooperation and thus speaks to an interest that we have. Baker goes on to note that:

[...] normative principles are always or nearly always explanatory when they are unifying, because our interests in normative discourse ensure that unity is always salient. (Baker, 2021, p. 9)

We thus arrive at a quite general pragmatic account of moral explanation.

This account, in highlighting the role that our interests play in explanation, is obviously akin to what I have outlined in the previous section. On further reflection, this should not be too surprising. Even though Baker offers his pragmatic view as a way for expressivists to make sense of moral explanation, there is no reason why only expressivists should be able to offer such an account. Indeed, pragmatic accounts will be welcome to anyone in search of an alternative to the overtly metaphysical emphasis that the grounding model, at least on a non-deflationary understanding, gives to moral explanation. Still, the way we get to this account matters. Baker, for reasons too complicated to recount here, thinks that expressivist ways of capturing the notion of moral grounding must fail. Instead, expressivists should deny that our moral statements should be interpreted in this way in the first place, instead modeling their understanding of explanation on pragmatic accounts in science and mathematics. Constructive deonticists, in contrast, do not deny that the grounding model is an accurate way to model how most ethicists think of moral explanation. However, due to considerations stemming from moral underdetermination, they become skeptical about moral explanation *construed in this way*. This inevitably

leads to the question of what the point of moral explanation is, if we don't want to completely abandon the notion. And here, the pragmatic notion makes its entrance for the constructive deonticist. Constructive deonticists, even though they can agree with expressivists about how to think of moral explanation, are thus not expressivists.

What about the other main anti-realist camp? Just as scientific anti-realists fall into two camps, those who accept a literal, factualist understanding of the claims of science and those who do not, so too do metaethical anti-realists. Constructive deonticism clearly does not fall into the latter category, that of expressivism. Instead, it agrees semantically with the other major branch of moral anti-realism, error theory, in that both think that we should construe ethical discourse in a literal, factualist, way. Does this mean that constructive deonticism is a kind of error theory?

This issue is more complicated. Van Fraassen's view of scientific explanation is often described as a kind of fictionalism.¹⁶ Fictionalism about a domain of inquiry can be characterized as follows:

- (a) As against the instrumentalist or the non-cognitivist, the fictionalist maintains that claims made within the discourse are genuine representations of how things stand, and that they are therefore normally capable of truth and falsity.
- (b) As against one sort of reductionist, the fictionalist maintains that the language of the discourse is to be interpreted at 'face value'. Claims within the discourse genuinely imply what they are most naturally taken to imply. [...]
- (c) As against one sort of realist, the fictionalist maintains that the ultimate aim of the discourse in the area is not (or need not be) to produce a true account of the domain, but rather to produce theories with certain 'virtues' – virtues a theory may possess without being true. (Rosen, 2005, p. 14)

This description fits well with how I have described constructive deonticism, and unsurprisingly so, since it is the ethical mirror image of van Fraassen's position.¹⁷

Now, fictionalism is often taken to be a kind of error theory. But this is a mistake. It depends on the kind of fictionalism. *Revolutionary* fictionalism makes a prescription to the effect that we *should* revise the aims of the

¹⁶ Even more, van Fraassen is often credited as one of the pioneers of modern fictionalism. Compare Kalderon (2005, p. 1) and Rosen (2005, p. 17).

¹⁷ The description also makes explicit that constructive deonticism comes out as a *factualist* position about moral discourse. As Kroon (2011, p. 791) observes, this is a strange misnomer for a position with "fictionalist" elements.

discourse at hand, and it does so because it takes these aims to be mistaken. This renders it a kind of error theory. In contrast, *hermeneutical* fictionalism is a description of the actual aims of a discourse. The discourse, it is claimed, is already understood in terms of a fiction, and hence there is no mistake involved. This is not an error theory.¹⁸

What kind of fictionalism is van Fraassen's view? Again, the answer is that it is complicated. In the minimal definition of constructive empiricism referred to above, van Fraassen claims that "science *aims* to give us theories which are empirically adequate [...]" not "science *should* aim to give us theories which are empirically adequate [...]." This suggests that van Fraassen's view should be classified as a form of hermeneutical fictionalism. No mistake has been made, and hence van Fraassen's view should not count as an error theory.¹⁹ Constructive deonticism, by copying this formulation, would seem to constitute a hermeneutical form of fictionalism, too, and, hence, is not an error theory. Yet van Fraassen's view isn't that easy to classify. As Rosen (1994) has argued, it is very implausible to assume that constructive empiricism is simply a view about the attitudes of (most) scientists. This would make it an overtly sociological thesis, for which van Fraassen does not offer any concrete evidence. In response, van Fraassen has clarified his view and denied that there is any obvious connection between what *science* aims at and what *scientists* aim at, his view only pertaining to the former. The mirror view in ethics would thus hold that ethics aims at deontically equivalent theories, while leaving open the possibility that ethicists are often mistaken in what they take the aim of ethics to be.

This is an important distinction, because it helps us make better sense of the relation of constructive deonticism to the grounding model of explanation. The pragmatic view of moral explanation that constructive deonticism puts forward is *not* what most ethicists have in mind. Instead, it is a prescription for how ethicists should think about moral explanation once they have (hopefully) come to see the futility of explanation on the grounding model. In this regard, constructive deonticism constitutes a form of revolutionary moral fictionalism and thus a form of error theory.

That is, insofar as ethical explanation is concerned. Yet the moral version of fictionalism (both in its error theoretic and non-error theoretic forms) does not only pertain to moral explanation. It pertains to (the discourse of) morality in general. In contrast, constructive deonticism claims that while we cannot have knowledge of moral explanation, and moral explanation

¹⁸ Compare Kalderon (2005, pp. 5–7) and Kroon (2011, p. 792).

¹⁹ Compare Kroon (2011, p. 791).

might thus best be thought of in terms of a fiction, we can have knowledge about which acts should be done. Hence constructive deonticism is at best *half* error theory. And even that might be granting too much. As van Fraassen (1980, p. 11) quips, the fundamentalist theist, the atheist, and the agnostic all agree when it comes to what religious texts entail. That doesn't make the atheist half fundamental theist. The same seems true in the present case. The defining idea of error theory is that our moral discourse rests on a fundamental mistake that renders all our positive moral statements wrong. Constructive deonticism does not buy into this since it allows for knowledge about what we should morally do. As I will shortly explain, this difference is all important since there is a crucial disanalogy in the way that realism is defined in ethics in contrast to science. Because of this difference, constructive deonticism might not even count as an anti-realist position in ethics, rendering the difference to error theory even more fundamental.

But before I attend to this, let me, as a brief side-note, mention one way in which constructive deonticism seems to have an advantage over error theory. Many people have wondered why Mackie, after setting out his metaethical view in the first part of *Ethics: Inventing Right and Wrong*, continues to do normative ethics in the second part in a seemingly unfazed way. If we don't believe that there are any truths in ethics, why bother about normative ethics? Mackie (1977, p. 239) alludes to a fictionalist answer when, on the last page of the book, he talks about a "useful fiction." Still, we might also wonder why we should continue to use a fiction once we know that it is just that. In contrast, according to constructive deonticism, we can be right or wrong about what we should do, and we can, at least in principle, even know when we are acting rightly or wrongly. We only start making a mistake when we assume that we can find out what ultimately makes these acts right or wrong. But this, I contend, does not impact most of normative ethics in any serious way. Now to be clear, I am not claiming that fictionalists have *no* answer to this challenge.²⁰ But constructive deonticists have a disarmingly simple one. We are allowed to continue with most of our everyday moral discourse (and normative moral theorizing) because nothing is wrong with it. Hence, constructive deonticism can accommodate some of the error theorist's worries, those about moral explanation and deep metaphysical truths. Yet it does so at a much lower price, because it upholds the validity of much of normative ethics.

²⁰ Compare Joyce (2005) for a promising account.

Constructive Deonticism as a Form of Realism

Constructive deonticism is clearly not an a form of expressivism and at most half error theory. It thus does not fall (neatly) under either of the two main strands of metaethical anti-realism. Does that make constructive deonticism a realist position? This would be highly surprising, considering that van Fraassen's structurally analogous position is being put forward with the explicit goal of providing an alternative to scientific realism. However, the philosophy of science need not be authoritative in this matter. There is a significant dialectical difference between the debates about realism in science and in ethics. In the philosophy of science, the issue of realism versus anti-realism (at least as pertains to the epistemological component) is generally framed roughly like this: Realists and anti-realists agree that we can have knowledge about observable facts. Where they differ is whether this also holds for unobservables. Realism entails the claim that there are truths to *all* (or close to all) scientific claims, and that we can in principle know these. Anti-realism, in turn, claims that we have no access to truths about *some* of our scientific claims (the ones about unobservables). This is what makes van Fraassen's view an anti-realist one even though he accepts that we can have knowledge about some scientific claims (the ones about observables). The dialectical situation in metaethics is different. Here, realists are traditionally only taken to claim that there are *some* claims about morality that we can know to be true, and anti-realists in turn claim that there are *no* truths (and hence no knowledge) to be found in ethics.²¹

If this is the present dialectical situation, constructive deonticism is a realist position since it allows for knowledge of truths about the deontic. By the same token, the skeptical argument from underdetermination does not even begin to challenge moral realism, for it only challenges our knowledge about the explanatory, not the extensional, claims of our theories. What is more, one might even think that the realist's case of coming closer to finding truths in ethics has been advanced by the assumption of convergence between the traditions. After all, the skeptical argument acknowledges that the best moral theories agree regarding extension. Even if there is a problem of underdetermination when it comes to the theories' explanatory claims, at least there are no more disagreements remaining about the deontic side. Realists might thus be more confident that they have found at least some truths in ethics.

²¹ Again, I am, for the sake of simplicity, setting aside the various *minimalist* and *quasi-realist* views that attempt to make room for the notion of moral truth, although a less fully fledged one than realists aspire to.

This is Parfit's way of looking at the situation. After outlining in great detail how distorting influences might be responsible for many of our disagreements about deontic verdicts, he makes the following claim:

Some other moral disagreements are not about *which* acts are wrong, but about *why* these acts are wrong, or what *makes* them wrong. Different answers are given by different systematic theories, such as those developed by Kantians, Contractualists, and Consequentialists. Such disagreements do not directly challenge the view that we are able to recognize some moral truths. In defending this view, it is enough to defend the claim that, in ideal conditions, there would be sufficient agreement about which acts are wrong. Though we also have intuitive beliefs about why many acts are wrong, and about the plausibility of different systematic theories, we would expect there to be more disagreement about these other questions. (Parfit, 2011b, p. 554)

For what it's worth, I do not think that Parfit offers a good explanation for why we should expect there to be more disagreement when it comes to explanatory questions than deontic ones. But the idea is clear, nevertheless. Some disagreements, Parfit thinks, are less of a threat to realism than others – in this specific case, explanatory disagreements. Since, in a case of underdetermination, there are only disagreements of this kind left, realists, or so the reasoning might go, need not be concerned. They can simply give up the idea of being able to find truths about explanation and restrict their beliefs to the extensional level. This way, they avoid the skeptical threat.²²

Should realists thus be reassured that moral underdetermination does not threaten their view after all and welcome constructive deonticism as one of their own? I think this would be premature. There are two considerations, arising out of two asymmetries, that realists should consider first. My hunch is that after consideration, realists would be very hesitant to accept constructive deonticism as a genuine realist position. Instead, the existence of a position like constructive deonticism might lead us to reconsider the way the metaethical debate has traditionally been framed.

The first asymmetry that realists have to explain is obvious. Constructive empiricism is *designed* by van Fraassen (1980, p. 6) as an alternative to scientific realism. What he denies is that the theoretical claims, featuring unobservable entities, that we put forward in order to explain the observable evidence have the same entitlement to be taken at face value. Most scientific realists accept this challenge. They seem to think that a realist position worthy of the name cannot restrict its realism to the claims that theories

²² Compare also Hooker (2020, p. 17) for the observation that Parfit does not seem to be bothered by these kinds of disagreements.

make about what we can or could all readily observe; rather, it also has to apply to more theoretical claims.²³ The pressing realist questions are in the end not about the (observable) facts that we already agree about but about the further claims that scientists make to account for those facts. What van Fraassen offers, only knowledge about observables, is too meager for scientific realists because it does not match their much stronger faith in science and its ability to uncover the hidden structure of the world. If the philosophy of science is any measure, denying the theoretical claims of our best theories should thus count as an anti-realist position.

Of course, realists in ethics might just thump the table and insist that according to the way the debate has been defined so far, constructive deonticism is a realist position and that's it. Realists could accept that all that counts in ethics is deontic adequacy, meaning that, contrary to the structurally analogous position in science, constructive deonticism should count as a realist position. However, to me, that seems to at least betray a lack of curiosity. If, as ethicists, we are confronted with a so far overlooked position – constructive deonticism – and we learn that the analogous position in another domain is considered squarely anti-realist, we should wonder why it would count as a realist position in our own domain. This should lead realists to inquire into what attracted them to moral realism in the first place and whether what we are left with according to constructive deonticism – knowledge of what is right without any explanation of why this is so – is really all they want from a position that deserves to be called realism. My hunch is that most realists would not be satisfied with this. The whole thrust of the skeptical argument and of constructive deonticism is deeply anti-metaphysical and pessimistic about our abilities to arrive at what are arguably the deepest truths in ethics. Constructive deonticism seems at most a second-class realism for disillusioned metaethicists.

So, to be perfectly clear, I grant that science does not dictate terms to ethics. Realism might come in domain-specific variations. However, we are at least entitled to an explanation for this asymmetry. Without that, it seems ad hoc for the moral realist to simply claim victory. Why, we can ask, is the glass taken to be half full in ethics whereas it is, in the structurally analogous case in science, considered to be half empty?

The second asymmetry that realists would have to explain to us is the differential importance they accord to disagreements about extension in comparison to explanation. If disagreements about the deontic are indeed a

²³ Compare Kosso (1992, pp. 102 ff.) for the significance of observability in the realism versus anti-realism debate.

threat, as most realists acknowledge, what makes explanatory disagreements so special that they are not? To be sure, realists are not wedded to the notion that we can know the answer to *every* moral question. However, can they just concede that a whole class of statements (those about moral explanation) are beyond our reach? It would seem that they cannot, at least not without further argument. Usually, realists argue that when we are not able to agree about a moral question, this can be explained by phenomena such as vagueness, lack of evidence, or bias. But it would still have to be shown why explanatory statements summarily suffer from such an impairment. At least at first sight, the explanations we put forward every day for why certain acts are right or wrong are to be taken at face value just like our claims about which acts are right or wrong. This presumably also holds for our moral theories. It thus seems arbitrary to claim that we need not worry about knowing the truth about explanatory claims. Thus, the specific problem here is not that we are excluded from knowing some moral truths (although I do think that that is a problem, too). What makes constructive deonticism especially threatening to realists is that it entails that we are excluded in principle from ever having knowledge about what is a whole class of moral claims, and a very important class at that: the class of explanatory moral claims. The fact that it is a whole class of claims that we don't have knowledge about calls for a systematic explanation.

Perhaps realists can come up with satisfying explanations for both asymmetries. Personally, I doubt it. This then leaves us with two options when it comes to classifying constructive deonticism. One is to stick with the old classificatory scheme, according to which constructive deonticism formally classifies a form of realism, because it allows for knowledge of some moral truths. Yet this, if I am right, would be an uneasy peace, and constructive deonticism would remain a position that few realists would whole-heartedly want to accept. The other option is to rethink the dialectics of the realism debate in metaethics. Constructive deonticism, at least in spirit, is very much on the anti-realist side even though the traditional classification doesn't put it there. This could be a reason to amend the classificatory scheme. Perhaps allowing for knowledge of some truths in ethics is not enough to classify it as realist. Instead, it does depend on which truths a theory claims to be able to account for. Not allowing for any explanatory knowledge in ethics is just not enough for a realist position, and constructive deonticism thus classifies as anti-realist.

My sympathies, as should have become clear by now, are with the latter move. At the very least, the burden of proof is on realists to show us why

the new position should be considered a realist position. Barring further arguments, the idea of restricting the aim of our moral theories to deontic adequacy looks both ad hoc from a realist point of view and suspiciously similar to that of restricting the aim of our scientific theories to empirical adequacy, which is, after all, an anti-realist suggestion. The new position would yet have to earn its right to be called a realist position. Whether this can be done, I propose, is one of the most interesting further issues that moral underdetermination brings to light, holding the potential to impact the whole dialectics of the realism debate in metaethics.

7.3 Two Worries

To wrap up the discussion of constructive deonticism, I want to attend to two worries. The first one concerns the distinction between the deontic and the explanatory, or, more specifically, why we should think that this distinction carries as much weight as I have so far assumed. The second one concerns the question of whether constructive deonticism is a stable position. These are real worries insofar as I do not have definitive answers to offer. Instead, I think that these are the two most pressing issues that a constructive deonticist needs to answer in the future.

The Deontic vs. the Explanatory

My argument in this book has relied very heavily on the distinction between deontic and explanatory claims in ethics. Indeed, as the discussion of constructive deonticism's status as a metaethical theory in the previous section has shown, focus on the systematic difference between our epistemic status with regard to, on the one hand, knowledge of deontic truths and, on the other, knowledge of explanatory truths is perhaps *the* distinctive aspect of the position. Against this background, one might wonder whether this isn't putting too much emphasis on the distinction.

There is a parallel worry in the philosophy of science. One of the earliest criticisms leveled at defenders of scientific underdetermination concerns the selective skepticism that underlies their arguments. Proponents of underdetermination think that it is only a theory's claims about unobservables that are threatened by underdetermination, whereas the claims about observables are left untouched. This presupposes a relevant epistemic difference between the two classes of statements.

However, some commentators have questioned whether we can give a principled defense of this distinction. Maxwell (1962, pp. 7–15) provides a vivid discussion of this point. He starts by asking what difference there

is between observing something directly, through a window, with the help of glasses, through binoculars, or, finally, through a microscope. There is an obvious continuum here, since with each step, some of the immediacy of the former way of perception is lost. However, no one would probably want to draw the line between the observable and the unobservable at the first or second step. Where is it to be drawn then? Proponents of underdetermination need to provide non-arbitrary criteria for drawing the observable-versus-unobservable distinction. However, in Maxwell's assessment, they cannot do so. Instead, it is on the basis of contingent human restrictions, physiological and technological, that they make the distinction. Yet, such contingent factors, Maxwell thinks, do not have the significance proponents of underdetermination accorded to them.

How can defenders of underdetermination react to this challenge? A first step is to notice that Maxwell only establishes that the distinction is a vague one, not that it is one without importance. Many people will probably think that there is a relevant difference between seeing something with our own eyes and using sophisticated instruments to observe microscopic objects. The fact that they cannot point out where exactly the distinction lies does not necessarily speak against the importance of the distinction. Still, we need an argument that establishes at least such a vague distinction. The following is offered by van Fraassen (1980, pp. 13–19). Granting that the distinction is vague and relative to our human capacities, nothing more is called for. Underdetermination asks for our justification, *qua human beings*, in believing in two different classes of claims, those about observables and those about unobservables. Although these claims may not be logically distinct, whether we are justified in believing them depends on our epistemic capacities, which are contingent on our biological human nature. As humans, we are thus justifiably more assured when it comes to what is observable *by us* than what is not. No more principled distinction is therefore called for.²⁴

What can we learn from this discussion with regard to moral underdetermination? I think that the challenge has an obvious analog in ethics. The skeptical argument puts a lot of emphasis on the distinction between deontic and explanatory claims. The class of the former claims underdetermines the class of the latter. But why should this distinction carry that much weight? To make things harder, a van Fraassen-style answer is not

²⁴ Compare also the discussion of Cartwright (2007, pp. 37 ff.), which is, to a large degree, seconded by van Fraassen (2007, pp. 342–344).

available to the ethicist. The distinction cannot be explained by referring to our epistemic limitations *qua human beings*. Whatever one's theory of how we can attain moral knowledge (if one thinks that this is possible at all), it presumably does not include reference to some sense organ that contingently makes a difference between what is right and why it is right. So this way of answering the challenge seems blocked for the proponent of moral underdetermination.

What other options remain? At first sight, it seems that skeptics in ethics have a very straightforward answer at their disposal. In contrast to the distinction between the observable and the unobservable, the distinction between the deontic and the explanatory is one of *kind*. The reason why van Fraassen has to emphasize our experience *qua human beings* to such a great degree is that there is no logical distinction between claims about observables and about unobservables. However, there *is* a crucial difference between deontic claims (which actions are right) and explanatory claims (why these actions are right). What is more, this is a point that many realists in ethics, who are likely to be the ones to push this challenge, would have to accept because they accept the grounding model. On the grounding model, as Berker (2018) makes clear, the explanatory claims of moral theories turn out to be metaethical claims. If we accept this, the distinction between the two classes turns out to be one of kind, and we thus have a good basis for treating them differently epistemologically.

However, the issue is more complex for the constructive deonticist. Since constructive deonticists do not accept the metaphysically strong version of the grounding model, they cannot distinguish between deontic and explanatory claims on that basis. Moreover, there is at least one understanding of pragmatic explanation, Baker's unifying generalizations, on which a principled distinction seemingly cannot be drawn at all. On Baker's understanding, moral explanatory claims turn out to be nothing more than generalized deontic claims. Yet if this is true, then the distinction between explanations and particular deontic claims cannot be one of kind.

Where does that leave us? I think that we have some kind of draw. Attacking the distinction between observables and unobservables is one of the most promising strategies for scientific realists to counter the skeptical argument from underdetermination. If what I have said so far is true, then at least one prominent kind of moral realist, the one that accepts a grounding model of explanation, cannot pursue this strategy because in their view, moral explanation – understood as picking out grounding relations – is fundamentally different from deontic claims about what we should do. In this regard, the skeptical argument once more seems to be on firmer

ground in ethics than in science. However, when we consider the prospects for constructive deonticism to answer this challenge, the issue becomes less clear. By likening explanatory claims to generalizations of deontic claims, Baker's pragmatic account muddies the waters. Constructive deonticists thus have to insist, *pace* Baker, that moral explanatory claims go beyond mere generalization. Whether this succeeds, I contend, is one of the most important questions constructive deonticists need to answer.

Skepticism All the Way Down?

The second worry is perhaps even more fundamental and potentially damaging to constructive deonticism. For even if we accept that a principled distinction can be drawn between deontic claims and explanatory claims, the question remains as to whether we can limit the force of the skeptical argument to explanatory claims, as the constructive deonticist wants to do.

To see the problem, consider first an argument that the skeptic might want to make against a realist like Parfit. According to Parfit (and like-minded realists), the fact that our best moral theories arrive at the same deontic verdicts should make us more convinced that we can actually know the truth about particular deontic matters. Presumably this is so because the verdicts (and our intuitions that match these verdicts) are now doubly (or triply) supported by what are all plausible theories.²⁵ However, if this is the reasoning behind the realist's optimism, there is a problem with it. For why are we supposed to believe that the fact that the rival theories agree regarding those verdicts gives us reason to be more confident about them? After all, the verdicts are only supported by theories that we have reason to withhold belief in! At least that is what the argument from moral underdetermination proves, according to the skeptic. Yet if we should withhold belief in the theories, we should not assume that the verdicts that follow from these theories are strengthened. The fact that a theory that we have no reason to believe arrives at some result should not be considered an additional confirmation of that result.

If this reasoning is sound, then the optimism of realists like Parfit is misguided. This, by itself, is no problem for the constructive deonticist. The constructive deonticist does not need to say that convergence between our theories *strengthens* our belief in the truth of our intuitions about particular cases. Our justification in believing them might be strong enough

²⁵ The point is made in a nice way by Ridge (2009, p. 80), who observes that the three traditions that Parfit looks at are at least on the shortlist of the most promising moral theories.

in the first place (without any support by moral theories) so that it is reasonable to think that we can have knowledge about deontic matters even if we don't have knowledge about moral explanation. However, the skeptic might push (or be pushed to push) this line of argument even further. Many people think that our justification in believing specific deontic claims does depend on whether they can be supported by plausible principles. If, as constructive deonticism states, we aren't justified in believing these principles, that support disappears. Indeed, the situation might be even worse because the skeptical argument assumes that not even our best theories can confirm the particular intuitions we have. Shouldn't this also impact our justification in believing that we can have knowledge about the deontic side of morality? Metaphorically speaking, we might say that the pressure that has been put on our explanatory principles by the underdetermination argument trickles down to the particular verdicts. If we should suspend belief in *why* our verdicts are right, we are no longer entitled to believe that we have identified the correct verdicts.

Constructive deonticists need to answer this challenge unless they want to be led to what looks like a more classical error theorist position in metaethics, according to which we can have knowledge neither of what makes acts right nor of what the right acts are. In addition, there is yet another worry. The radicalized skeptical argument is in danger of turning out to be self-contradictory. According to it, the loss of justification in our beliefs in the principles also renders the particular verdicts that follow from them unjustified. We should thus believe neither the explanatory principles nor the deontic verdicts of a theory. Yet, according to the skeptical argument, the reason why we should suspend belief in the principles in the first place is precisely that they are underdetermined by our considered judgments and intuitions about particular cases. This seems problematic. How can we first treat deontic intuitions as evidence for the theories, then make the fact that theory choice cannot be made on the basis of these intuitions the reason why we should suspend belief in the explanatory principles, only to finally make the loss of certainty about these principles the reason to doubt our intuitions? That argument seems to be viciously circular.

How can constructive deonticists reply? I think that the most promising way to go is to argue that we have theory-independent justification to believe judgments about the rightness or wrongness of actions. If we have independent justification to believe that we know what is right or wrong, then the fact that our theories fail to provide any such justification does not undermine this initial justification. It might leave us worse off relatively,

because the hope that our best moral theories might provide such additional justification has to be given up. However, depending on how strong the initial justification is, this need not make us skeptics about the deontic. Still, the constructive deonticist has to make a case that we have sufficient *theory-independent* justification to believe that we can know which acts are right or wrong. This is the second task that needs to be completed in a more thorough defense of constructive deonticism.