

The Big Picture in Normative Ethics

In the previous two chapters, I have examined both an extended case study as well as two projects that promise a more extensive form of deontic equivalence of moral theories. Each time I have argued that, contrary to how proponents of these projects interpret their results, the best interpretation is in terms of moral underdetermination. This is an interesting result in itself and hopefully can help us better understand these projects. However, the time has come to go beyond such examples and consider the big picture. How are we to evaluate the case for moral underdetermination, which forms does it take, and what lessons does it hold? As the discussion of scientific underdetermination has shown, some forms of underdetermination are considered to be quite uninteresting or even trivial, and these charges might transport to ethics. Much thus depends on which theories can be shown to be underdetermined, how plausible they are, and how widespread the phenomenon is. It is to these more systematic questions that we need to turn if we want to assess the importance of moral underdetermination as a general phenomenon.

I start with a critical assessment of how compelling a case for moral underdetermination can be established from what has been said so far. For this, I compare the three projects with the strategies that have been employed in arguing for scientific underdetermination. Next, I try to paint a picture of the various forms moral underdetermination takes in ethics, again applying some of the distinctions that were drawn in the discussion of scientific underdetermination. Finally, I discuss three ways in which thinking in terms of underdetermination might impact the study of normative ethics in the future.

5.1 The Roads to Moral Underdetermination

In Chapter 1, we saw that proponents of scientific underdetermination have employed at least three strategies. Parfit, consequentializers, and

deontologizers, of course, do not intend to argue for moral underdetermination. However, if the underdetermination interpretation of the results of their projects is correct, they can nevertheless be thought of as employing certain argumentative strategies that *de facto* lead to moral underdetermination. It is thus natural to wonder whether their arguments also share some of the observed advantages and shortcomings of the analogous strategy in science. It turns out that, indeed, they do, and comparison to science thus proves highly instructive in evaluating the persuasiveness of the projects in normative ethics.

Algorithms and Gimmickyness

I will start with consequentializing and deontologizing. Both procedures, it can easily be seen, bear a close resemblance to the algorithmic strategy in the philosophy of science. In both cases, we are presented with relatively simple recipes to construe alternatives to a wide variety of theories, scientific or moral. Just as Kukla's algorithm is an all-purpose means to construe empirically equivalent theories to any scientific theory, consequentializing and deontologizing are all-purpose means to prove deontic equivalence between any (minimally plausible) consequentialist or non-consequentialist theory and at least one version of the rival tradition.

Such algorithms are extremely powerful when it comes to generating cases of underdetermination. However, the philosophy of science indicates that this comes at a price. Algorithms have problematic features that threaten their relevance. That there are analogous problems with the algorithms we have considered in the moral domain has not been lost on ethicists. Indeed, such awareness can be traced back to the pre-consequentializing era. Nozick (1974, p. 29) considers a strategy that is similar to consequentializing, when grappling with the question of whether there is a way to include the idea of side constraints into a consequentialist framework. In principle, he agrees that certain modifications might allow consequentialists to attain this goal. However, Nozick (1974, p. 29) remains unimpressed, instead dismissing the whole undertaking as *gimmicky*. That passage has subsequently been quoted frequently.¹ Something about such strategies has struck many as dubious and unserious. However, it is not so simple to put one's finger on what exactly the problem is supposed to be.

¹ Compare Dreier (1993, p. 23) and Smith (2009, p. 257). The charge of gimmickyness is also included in the title of Vallentyne (1988).

What are philosophers objecting to when they call consequentializing and similar procedures “gimmicky”?

Here, it is helpful to look at the objections that have been leveled at the *algorithmic strategy* in science. Philosophers of science, we have seen, have substantiated the claim that algorithms are unserious or gimmicky in two ways.

One way the charge of gimmickyness has been elucidated is by deference to scientists’ judgments. Stanford (2001, pp. 11–12) holds that algorithmically construed theories are not such that scientists would consider them serious alternatives to real scientific theories and that philosophers making use of them are striking a devil’s bargain. Might ethicists be making a similar mistake? A vague impression of this kind might be behind some of the criticism of consequentializing. Consequentializing involves at least two contentious moves (regarding the concepts of goodness and of consequence), and these might look gimmicky. But are these moves really too peculiar to be seriously considered by ethicists? I don’t think so, for three reasons. First, we have already seen that both moves have a respectable historical pedigree in philosophers like Sen and Broome. If something is wrong with the moves, it is at least not something peculiar to consequentializing. Second, the algorithms are much less suspicious than the ones that have been proposed in science. Remember Kukla’s recipe:

Given theory T, construct T₂ which asserts that T holds when somebody is observing something, but that when there’s no observation going on, the universe follows the laws of some other theory T’. (Kukla, 2001, p. 23)

Now compare this to what consequentializers and deontologizers do. Surely, introducing a skeptical streak into our theories as Kukla does is much more controversial? What consequentializers rely on are not such radical moves. Rather, they draw on quite familiar and respectable notions. Granted, they do so in the service of a thesis that is much more expansive than what proto-consequentializers conceived of. Still, the moves do not stray anywhere as far from accepted normative theorizing as Kukla’s algorithms stray from accepted scientific theorizing. This is conceded by the man himself. Kukla (2001) readily acknowledges that the theories his algorithms come up with are *bizarre*. In contrast, many consequentializers consider their theories to be equally as attractive as, if not superior to, traditional theories. Granted, this might be due to bias on the part of consequentializers and deontologizers. Their critics see the case markedly different. However, at least it seems clear that we cannot simply defer to practicing ethicists to dismiss consequentializing theories. The same is

true for deontologizing. The theories we end up with might not satisfy all the desiderata we have for the deontologist tradition, but they are not bizarre in the way Kukla's are. Third, the genesis of consequentializing and deontologizing themselves is less suspicious than in the case of Kukla's algorithm. Kukla's algorithm is introduced specifically to support the non-uniqueness version of underdetermination. The same is not the case in ethics. Consequentializing and deontologizing have sparked interest independently of any considerations of underdetermination. They are thus not as easily written off as gimmicks introduced with the sole purpose of arguing for a preconceived thesis.

The fact that we cannot dismiss consequentializing and deontologizing on this basis, however, does not mean that there is no other way of making the reproach of gimmickyness stick. As we also learned from the discussion in science, there is a second way to substantiate the claim that algorithmically construed theories should not be taken seriously. Like Laudan and Leplin (1991) in the scientific case, critics might try to dismiss the two procedures by identifying more substantive faults that both these procedures share. Laudan and Leplin's two charges are of algorithmic theories being *superfluous* and *parasitic*. Do these charges apply to consequentializing and deontologizing?

Let us start with superfluity. Scientific theories are superfluous, on Laudan and Leplin's account, if what is added by the algorithm in order to guarantee incompatibility with its target theory is from the beginning construed in such a way as to not yield any new empirical insights. Take Kukla's algorithm as an example. The added element in the explanation – that the universe follows different laws when no one watches – has no conceivable empirical consequences. Indeed, this is deliberately so, since if the added element were to have empirical upshots, this would threaten the empirical equivalence with the target theory. The added part in Kukla's algorithm thus seems wholly superfluous because it doesn't even aim at producing new empirical insights. What does the analogous objection amount to in ethics? Here, the charge would be that theories are superfluous if they do not advance on another theory's deontic verdicts. Let us grant that consequentializing and deontologizing produce theories that are superfluous *in this sense*. Still, we need to ask what this means. Superfluity, thus construed, means that moral theories do not make new deontic verdicts about particular cases. Yet, as I have been at pains to emphasize in the previous chapters, the new theories do advance in another way. They do offer different, incompatible, explanations. In addition, as we shall see shortly, they might also exhibit theoretical virtues that the target theories do not have. The

consequentializing and deontologizing theories are thus not superfluous in the *broader sense* of providing nothing substantial beyond the original theories. Hence, the charge of superfluity does not seem to cut deep.

The more important charge, I think, is that the consequentializing and deontologizing theories might be *parasitic*. Laudan and Leplin take theories to be parasitic if they depend entirely on the explanatory and predictive mechanisms of another theory and need to make reference to that theory for whatever they are supposed to explain. This charge is more difficult to refute. It does not take issue with the newly construed theories merely for not advancing on their target theories when it comes to new predictions. Instead, it identifies too close a dependence on the target theory when it comes to explanation as well. If true, this would be a problem for ethical theories and, as it turns out, critics have indeed been quick to note it.

Hurley (2013) gives a good rendering of the charge as it relates to consequentializing. Remember that in order to achieve deontic equivalence, consequentializers have to import non-consequentialist features into the consequences and specify a fitting ranking. This ranking of outcomes will then decide which acts are right or wrong, rendering the theory formally consequentialist. However, as Hurley points out:

[T]he ranking of states of affairs need not play any substantive explanatory role in such a theory at all – it can produce consequentialized theories upon which the ranking of affairs as better or worse is an explanatory fifth wheel. (Hurley, 2013, p. 136)

Does this objection speak decisively against consequentializing theories? I think not. More precisely, it may speak against relying exclusively on algorithmic strategies in our arguments but not against the theories that are produced by these algorithms. Here is what I mean by this. Having as their goal the production of deontically equivalent theories, consequentializers (and deontologizers) need by definition to turn to other theories for their verdicts. The explanatory claim can only be attached after this is done, making it vulnerable to the charge of parasitism. However, this is mainly an epistemic problem. In order to know which verdicts to copy, consequentializers (and deontologizers) need to turn to the target theories. Yet this does not mean that we could not have arrived at the same verdicts had we made the same modifications to consequentialism (and deontology) independently. Portmore (2011, pp. 112 ff.) illustrates how this could work.²

² See also Portmore (2022).

He proposes a coherentist method, where we revise our pre-theoretical judgments about what is right and our judgments about the ranking of acts in light of each other until we reach reflective equilibrium. In this way, the evaluative ranking of outcomes need not be decided by independent considerations about what is right or wrong. Again, consequentializers' and deontologizers' goal of copying non-consequentialist theories renders it necessary to start with a given set of deontic verdicts. But this is a feature of the strategy. It does not mean that the resulting theories could not have been conceived of in a different way, rendering the charge that they are parasitic empty.

I thus think that charges to the effect that algorithms are a non-starter are too strong. These charges are overstated, and there is no knock-down argument against such algorithms in ethics. Nevertheless, I also think that if we want a better understanding of moral underdetermination, we should heed the call that Stanford (2006) directs at his philosopher-of-science peers for more case studies. Just as it is an advantage for proponents of underdetermination in science to be able to point to real-world cases, so it is preferable for proponents of moral underdetermination to be able to point to non-algorithmically construed cases of underdetermination in ethics. Before we see such cases, it will be difficult to counter the impression that anything so simple and at the same time so powerful as consequentializing or deontologizing must be too good to be true.

Reevaluating the Significance of Generic Examples

Fortunately, we have also seen a case study that is of a very different kind. If consequentializing and deontologizing can be interpreted in terms of the algorithmic strategy, Parfit's case study should be read in light of what we have called the *inductive* strategy in science. Parfit's project provides the kind of hard-earned example of underdetermination that Stanford (2001, 2006) urges philosophers to look for in the scientific realm, instead of relying on the algorithmic strategy. Indeed, Parfit's book-length, in-depth study makes for an especially impressive case study. Whereas Mill and Hare do little more than *gesture in the direction* of an agreement between Kant and consequentialists, Parfit attempts to *show* it at great length. He takes pains to develop his preferred versions of the traditions by engaging with countless problems and objections. *On What Matters* might well include the most extensive argument for the convergence of moral theories to date. Thus, when it comes to its contribution to the analysis of moral

underdetermination, Parfit's project delivers what has been missing in the other two, and it does so in an especially impressive way.

To be clear, Parfit cannot serve as a case study *for* consequentializing or deontologizing. Both of these projects make use of specific argumentative moves that are not the ones Parfit makes.³ Parfit is neither a consequentializer nor a deontologizer. Indeed, Parfit (2017a, p. 403) makes short work of the consequentializing strategy, echoing Schroeder's criticism that it does not appeal to any common-sense understanding of goodness and opining that it therefore does not lead to genuine convergence. Nevertheless, Parfit's project is instructive for the study of moral underdetermination precisely because it is immune to the charge of gimmickyness. The fact that Parfit does not argue via some algorithm means that he is neither vulnerable to the objection that no practicing ethicists would take his proposed theories seriously, nor to the objection that his preferred theories have a substantial flaw that can be established merely by looking at how he argues for them. Granted, many normative ethicists have taken issue with parts of Parfit's arguments. Yet, whereas the suspicion of gimmickyness is at least not beyond the pale in the case of consequentializing and deontologizing, there is no way that one can dismiss Parfit's arguments on the same basis. Parfit's preferred theories are clearly not *superfluous* or *parasitic*. They include different explanatory claims even though they do not go beyond their rivals extensionally. Nor, for its sheer complexity, is Parfit's argument subject to the charge by Norton (2008, p. 40) that any procedure that is at the same time very simple and very powerful must be too good to be true.

However, the fact that Parfit does not rely on algorithms also means that his project shares some of the downsides that have been associated with the inductive strategy. It means, for one, that we cannot simply duplicate Parfit's reasoning and transport it to other theories. The extensional equivalence is due to the very specific shapes Parfit's preferred theories have. If one of them had a different shape, the Convergence Argument would probably not succeed. For example, Parfit cannot simply adjust the theory of the good of his preferred consequentialist theory in order to account for the verdicts of the other theories, as consequentializers can. Our insights into moral underdetermination, as far as they are gained from the analysis of Parfit's projects, are thus far more reliant on the specific arguments he makes. I don't think that this invalidates the philosophical insights we have gained. Just as in science, there is always a possibility that some theory

³ In addition, consequentializers rely on an act-consequentialist framework, whereas Parfit prefers a rule-consequentialist one.

is proven wrong, and if this is the case for Parfit's theories, we shall have to revise his example or even relinquish it and find a new example of underdetermination in ethics. As Ross (2009, p. 145) observes, Parfit is humble enough to acknowledge that he might not have succeeded in his search for the best theories. There is no reason for me to claim more than Parfit does.

Still, one can ask what the real value is of just *one* specific case of moral underdetermination. Philosophers of science would not be too impressed with being presented with only one example of underdetermination. What interests them is whether underdetermination is a more pervasive phenomenon that pertains to most or even all scientific theories. Parfit's example is thus a very meager basis for an inductive case for moral underdetermination. I think that this worry is justified, *if* we understand it in one of two possible ways. Here is the justified worry. Parfit's example comprises just three moral traditions and only three versions of these. But we cannot be sure that these are the correct ones, nor that they are the only correct ones. There are others to be considered.

First, there are other moral traditions that have not been considered at all. Let us take virtue ethics as an example. Is it plausible that theories of the virtue ethicist tradition have deontically equivalent counterparts? Consequentializers would want one to think so. Peterson (2013, p. 170) considers virtue ethics in passing. He approvingly refers to Driver's discussion of the relation between consequentialism and virtue ethics, singling out a passage where Driver (1996, pp. 122–123) states: “[o]n my theory, the value of all these character traits resides in their tendency to produce good consequences [...]” Peterson takes this to show that consequentialism and virtue ethics are not fundamentally different moral views. The only major difference Peterson sees is one of focus: Whereas consequentialism zooms in on the evaluation of acts, virtue ethics asks what kind of person we should be. Apart from this difference in focus, however, the theories need not differ; specifically, they do not need to differ in their extensions. Chappell (2012, pp. 173 ff.) thinks that something similar explains why Parfit does not consider virtue ethics in his quest for the best theories. Whereas one part of virtue ethics is too unsystematic to be considered relevant, the other part is just a version of motive consequentialism. Thus, virtue ethics either fails to be a real alternative or comes down to a form of consequentialism. However, there is also a way of interpreting the issue in terms of underdetermination. As Chappell (2012, pp. 174 ff.) goes on to explain, the main difference between virtue ethicists and consequentialists is that the former, but not the latter, accept that some reasons do not come from the future (in the form

of consequences). Virtue ethicists and consequentialists might thus arrive at the same verdicts but include incompatible claims as to what constitute reasons for those claims.⁴

Second, there are other versions of the traditions that Parfit does consider. The version of consequentialism that has usually been thought to be most strongly opposed to the other moral traditions is act-consequentialism. Parfit's preferred version of consequentialism is of the rule-consequentialist variety, however. Parfit does go to great lengths in Volume Three to show that some of the deepest disagreements between act-consequentialism and common-sense morality as well as his other preferred theories might be resolved. Still, the reconciliation is only of a partial nature, and even though Parfit (2017a, p. 435) expresses hope that even more reconciliation might be possible, he ends up not endorsing act-consequentialism.⁵ For these reasons, it would certainly be welcome to see a more detailed case of moral underdetermination involving a theory of the act-consequentialist variety.⁶ One promising candidate might be *Kantian Consequentialism*, a theory proposed by Cummiskey (1990, 1996). Cummiskey, in contrast to Mill and Hare, not only claims that some of Kant's principles might be open to reinterpretation in a consequentialist vein. Instead, he thinks that we can give a genuinely Kantian grounding for consequentialist deontic principles.⁷ At the same time, Cummiskey is no deontologizer, since he does not attempt to copy the content of *any* consequentialist theory. Instead, he thinks that the Kantian foundational theory leads to a distinct set of deontic principles based on a distinct (two-tier) theory of the good. If, as would have to be shown, this distinct set of principles can also be grounded on act-consequentialist foundational principles, Cummiskey might provide another example of moral underdetermination, this time between a Kantian and an act-consequentialist theory.

Finally, there is also the possibility that the best versions would be members of the same tradition. Both consequentializers and deontologizers attempt to copy the verdicts of theories from rival traditions, and they thus provide examples for underdetermination between theories from rival

⁴ This fits well with the assessment by Hooker (2020, p. 7) that Parfit should have aimed for a *quintuple* theory, including virtue ethics as well as common-sense morality. Compare also Crisp (2020, p. 275).

⁵ Compare Hooker (2020).

⁶ Compare Crisp (2020, p. 275).

⁷ Cummiskey (1996, p. 11–15) holds that the principles of a moral theory have a consequentialist structure if they include no *agent-centered constraints* on the maximizing of the good. Kantianism, on his view, does not entail such constraints and hence leads to consequentialist principles.

traditions. However, nothing speaks against cases of underdetermination between theories from within the same tradition. As it happens, Parfit might after all provide an example of this kind. In my presentation of Parfit's arguments, so as not to complicate matters too much, I have left out one fact about Parfit's interpretation of Kant. As it turns out, Parfit (2011b, pp. 338–342) considers his interpretation of Kant to be a contractualist one. He considers his preferred Kantian principle – “Everyone ought to follow the principles whose universal acceptance everyone could will” – to be contractualist because it includes the viewpoint of not only one agent but everyone. Indeed, Parfit (2011a, p. 342) goes so far as to call it the *Kantian Contractualist Formula*. On the one hand, the fact that two of Parfit's preferred theories turn out to be contractualist might seem disappointing, because the fact that they can agree comes as less of a surprise.⁸ On the other hand, the issue is certainly interesting for the study of moral underdetermination since it showcases how underdetermination can be an issue within and not only between the moral traditions. If the classification of (Parfit's) Kant as a contractualist is accurate, we are presented with a case of underdetermination not between moral traditions but within one. In principle, nothing speaks against this. The main traditions of moral theories are broad churches, and some of their versions might be radically different. Even if we can include Kant under the contractualist umbrella, that does not mean that his version of that tradition is compatible with all the other versions.

The same is probably true for other traditions of moral theorizing as well. As a simple example, take a classical consequentialist who only values the sum of goodness and what might be called an egalitarian consequentialist who appoints some value to the equality of outcomes. Seemingly, these theories would come to different assessments of the rightness of acts that produce an unequal distribution of goodness. However, that of course depends on their theory of value. Non-egalitarian consequentialists might well subscribe to a theory of value that entails a very steep decrease in the marginal value of additional goodness that can be achieved if one person is better off than others. These consequentialists might thus arrive at very egalitarian results themselves. Though the egalitarian and non-egalitarian consequentialists would have disagreements about whether equality of outcome is a value in itself, they might well come to the same conclusions about what the distribution of goodness in the actual world should be.

⁸ Compare Bykvist (2013, p. 347).

Underdetermination between versions of the same theories thus also calls for a more thorough investigation.

In sum, there is a need for more case studies. As long as we don't know whether Parfit's theories are the best variants of the only true traditions, further investigations are imperative. However, there is also another way to doubt the significance of Parfit's case study. One might object that, even if one grants that Parfit provides one example of moral underdetermination between correct moral theories, what's the big deal? In science, to make a convincing case for underdetermination, we would need to see examples from many different fields (physics, biology, chemistry, etc.). Ethicists, one might think, would have to do the same before we'd need to take moral underdetermination seriously. That line of reasoning is misguided, I think. The reason why will become clear when we discuss the issue of *local* vs. *global* theories below. Before we do so, however, I want to say a bit about the fact that we have not seen arguments that are analogous to the third strategy in science, the *holistic* one.

The Road Not Taken

In the philosophy of science, we have seen that authors sometimes connect holism so intimately with underdetermination that they distinguish insufficiently between the two. In our discussion in ethics, however, holism has not played any major role so far. What is the explanation for this? One obvious explanation could just be that of limited sample size: Maybe consequentializers, deontologizers, and Parfit are exceptional in this regard and there are examples of the holistic strategy to prove deontic equivalence in ethics as well. It might also reflect a deeper point, however. Remember that holism alone does not lead to underdetermination. One can simultaneously hold that no hypothesis is ever tested in isolation and that there is only one theory that will withstand a more holistic form of testing. It is rather Duhem's and Quine's conviction that we will indeed find alternative theories that are equally as well supported as ours that underpin their arguments. That conviction might not be as pronounced in ethics. Perhaps ethicists are more convinced that if their principles are tested in a holistic way, those principles are the only ones to withstand.

That does not mean that use of the holistic strategy is impracticable in principle in ethics, though. Recall Quine's analogy about a web of beliefs. According to Quine, revisions to our web of beliefs are possible at basically any place, from the beliefs closest to observation to those at the very center. Quine (1951, p. 43) insists that not even the laws of logics are immune to revision. He proposes that if we want to hold on to some belief in the face

of seemingly recalcitrant observations, we might go as far as to revise the law of the excluded middle. Does this hold for ethics too?

There is a debate at the intersection of ethics and political philosophy that provides some insights: the debate about so-called *Dirty Hands*.⁹ Following Walzer (1973), philosophers have discussed whether morality sometimes demands from us that we dirty our hands. That is, they have asked whether we sometimes have to act wrongly in order to do the right thing. Walzer (1973, pp. 166–167) provides the example of the president of a country on the brink of civil war, who faces the question of whether they should torture a captured terrorist in order to prevent a series of destabilizing future attacks. Two rival positions that have been taken on the issue are of particular interest to us.¹⁰ Proponents of dirty hands, mostly non-consequentialists, answer in the affirmative. They think that although some acts are so bad that we always act wrongly when we commit them, we nevertheless have to commit such acts sometimes. Opponents, mostly consequentialists, answer the question in the negative. They hold that whether an act is right or wrong only depends on its consequences. Thus, if the consequences of the (allegedly) dirty act are the best overall, the act should be done, and it cannot be wrong at the same time. There is thus a clear disagreement here regarding the question of whether we sometimes have to commit a wrong act.

However, and this is where it gets interesting, proponents and opponents of dirty hands can nevertheless agree on what should be done in such situations. They both hold that the act should, in the end, be done. They differ only in the explanations they put forward for why we are supposed to carry out the act in question. A dirty hands proponent will advise you to commit the act, *even though* by doing this you will simultaneously commit a wrong. A consequentialist, by contrast, will explain to you that the act has the best overall consequences and is thus right *full stop*. What the dispute comes down to, at the core, is about some proposition that looks very much like a *rule of logics* for the moral realm. Those who think that dirty hands scenarios are impossible appeal to a claim like this:

An act that should be done can never be wrong.

Those who think that dirty hands scenarios are possible deny that the claim is true. Their differences are thus about a belief that, to adapt Quine's metaphor, constitutes the very center of the web of (moral) beliefs.

⁹ The discussion here follows Baumann (2021b).

¹⁰ I exclude, for simplicity, absolutists who think that a dirty act should never be done, no matter how horrific the consequences.

Nevertheless, proponents and opponents of dirty hands do not necessarily need to disagree about what should be done. The case of torturing someone to avert a major catastrophe might at first only seem to be acceptable to consequentialists. However, if deontologists want to allow for it, they can find a way to make room for it. In the dirty hands framework, they do so by amending one of the moral rules of logic, that is, the claim that acts that should be done can never be wrong.

Of course, this case is not an example of a full underdetermination scenario, since we haven't yet seen whether there are other deontic verdicts on which proponent and opponents of dirty hands scenarios disagree.¹¹ However, it at least indicates that we find in ethics precisely what Quine predicted: the willingness to even change the rules of (moral) logic in order to hold on to beliefs that one finds highly plausible. Although the debate is not described in these terms, it nevertheless provides some evidence that holism might play a role in ethics.

Ethicists, just like philosophers of science, apply many different argumentative strategies. If the arguments result in deontic convergence between incompatible theories, we can describe those ethicists as using argumentative strategies that do *de facto* lead to moral underdetermination. The discussion of different strategies to argue for scientific underdetermination might then prove instructive, and indeed, I have argued in this section that it is. Comparing the algorithms in ethics to those in science shows that the former are not as implausible as the latter, thus strengthening them to some degree. The way in which philosophers of science have urged us to look for more generic case studies should nevertheless be heeded by ethicists as well, in order to alleviate some of the doubts that surround algorithms in both domains. Parfit provides just such a more generic case study. Yet, whether we can add to it enough other examples to make for a convincing inductive case will still have to be proved. This was to be expected, and if we believe philosophers like Stanford, the only way to go from here is by collecting hard-earned examples of underdetermination involving other normative ethical theories, such as virtue ethics and act-consequentialism. Finally, an appeal to holism might work in ethics as well, but that remains to be seen. Since ethicists have not directly argued *for* moral underdetermination so far, it is perhaps no surprise that the picture we see is still very sketchy. But I think that it is nevertheless intriguingly multifaceted.

¹¹ Sorell (2000) offers some indication that this is indeed the case.

The same is true to an even greater degree when it comes to the forms of moral underdetermination that arise from these different argumentative strategies. It is to this topic that we now turn.

5.2 Varieties of the Phenomenon

Just as the phenomenon of scientific underdetermination comes in many disguises, so our discussion has indicated that the same is true for ethics. To get a more systematic overview of the variety of forms that moral underdetermination takes, it will be helpful to go through the distinctions from Chapter 1 once more and see if and how they apply to ethics. When discussing scientific underdetermination, I relayed the assessment by Laudan (1990) that anyone can rally behind the claim that scientific theories are underdetermined *in some sense*. The same is likely true for ethics as well. The truly interesting question is rather in *what sense*.

Existence and Uniqueness, No Egalitarianism

The first issue concerns the scope of the thesis, that is, how many theories are claimed to be underdetermined.

The results of Parfit's project are best understood in terms of the *existence* version. Parfit attempts to show that the best versions of the three most important traditions are deontically equivalent. If, as I have argued, they are nevertheless incompatible, this means that there is underdetermination between at least these three moral theories. Parfit thus provides evidence for

*The existence version**: For some moral theories, there is an alternative theory, which underdetermines theory choice.

Consequentializers and deontologizers are more ambitious.¹² They employ all-purpose means to show deontic equivalence between any

¹² The one exception is Portmore who, at a later stage, moderates his ambitions:

The hope is that by consequentializing such theories we can arrive at a substantive version of consequentialism that, like act-utilitarianism, embodies act-utilitarianism's compelling idea, but that, unlike act-utilitarianism, is compatible with our considered moral convictions. Such a version of consequentialism would be a hybrid that possessed the best characteristics of each breed. (Portmore, 2009, p. 331)

At this point, Portmore's goal is merely to come up with a commonsensical version of consequentialism.

(minimally plausible) consequentialist or non-consequentialist theory and at least one version of the rival tradition. Since consequentialism and non-consequentialism are jointly exhaustive categories, this means that we can combine the theories in the following way:

(DET): For any remotely plausible nonconsequentialist theory, there is a consequentialist counterpart theory that is deontically equivalent to it such that the two theories are extensionally equivalent with respect to their deontic verdicts.

and

(DET*): For any remotely plausible consequentialist theory, there is a non-consequentialist counterpart theory that is deontically equivalent to it such that the two theories are extensionally equivalent with respect to their deontic verdicts.

yields:

(DET**): For any remotely plausible theory, there is a counterpart theory that is deontically equivalent to it such that the two theories are extensionally equivalent with respect to their deontic verdicts.

If we further assume that radical explanatory disagreements between these theories remain, we can formulate

*The non-uniqueness version**: For any remotely plausible moral theory, there is an alternative theory that is deontically equivalent to it while at the same time being explanatorily incompatible.

Hence, deontologizing complements consequentializing in a way that yields the moral analog to the more encompassing form of moral underdetermination that has been postulated in the philosophy of science.

What we have not found are instances of the even more ambitious version, the *egalitarian* version, according to which the choice between any two deontically equivalent theories would be underdetermined. What might explain this? One explanation might have to do with the context in which the egalitarian version appeared in the philosophy of science. If Laudan (1990) is right, the egalitarian version only made its real entrance in the scientific literature (after a quick first appearance in Quine) in the context of very broad skeptical arguments. It is not explicitly advocated for but serves as a background for challenging scientific orthodoxy from an

extreme social-constructivist viewpoint.¹³ Since moral underdetermination has not been put to that purpose, one might think, it is not surprising that such a strong version has not yet appeared. However, this explanation falls short. As we will consider in more detail in Part III, at least Dreier's view is likely motivated by broad skeptical intuitions, and the view that there is *no* substantive difference between (deontically equivalent) consequentialist and non-consequentialist theories is hardly less radical than the view that choice between (deontically equivalent versions of) such theories is underdetermined. Instead, I will argue in Part III that an egalitarian version only comes to the fore when we think of the challenge that consequentializing and deontologizing pose as an epistemic, not a semantic, one. But this will have to wait.

So far, we can summarize that, even though the most ambitious version of the underdetermination thesis has not found its way into the ethical debate, the two most common versions – the existence and non-uniqueness versions – can be very naturally construed out of the projects we have surveyed.

The Local Is Global

The issue of scope is broadly analogous between ethics and science. The next issue I will look at is not, and in what I take to be a highly interesting way. Underdetermination theses in science, we saw, come in a local form – pertaining to theories of the special sciences – and in a global form – pertaining to the whole of science. The analogy to the scientific domain might suggest that what we have been dealing with so far in ethics are *local* forms of underdetermination. For example, the existence and non-uniqueness versions of scientific underdetermination pertain to particular scientific theories, from physics, biology, chemistry, and so on. Our three projects in ethics might appear in a similar light. Parfit aims to show deontic equivalence between three specific moral theories. Consequentializers and deontologizers aim to do this for all their rival theories, but these are still specific versions of the other traditions.

However, putting the situation this way is misleading. There is a crucial disanalogy between science and morality that we have not addressed yet. Moral theories do not relate to each other as do, for example, theories from biology and theories from chemistry. Whereas in science, different theories

¹³ Carrier (2011, pp. 189–190) also speaks of the high point of underdetermination's esteem coinciding with social constructivism.

from the special sciences account for different subsets of the evidence in different domains (biological theories for biology, chemical theories for chemistry, and so on), moral theories typically cover the whole realm of the deontic (or something close enough). Here is a representative formulation of this insight by Brandt:

Ethical theory has been interested in finding a set of valid ethical principles, which is *complete* in the sense that all true ethical statements can be deduced from it (given an adequate stock of nonethical or factual premises) [...]. (Brandt, 1959, p. 5)

What Brandt refers to here has subsequently been called the *completeness condition*.¹⁴ The completeness condition states that moral theories are in the business of specifying all correct moral verdicts. This does not presuppose that theories give verdicts of rightness or wrongness for every situation. There might be cases where there is no right or wrong. But it does mean that, insofar as there are correct deontic verdicts for specific cases, moral theories would entail all of them. Comparing the completeness condition to an axiom is a little misleading since, as Brandt's description itself makes clear, it is more of an ideal. There are some ethicists who purposefully restrict their theories to some domain of morality. A well-known case in point is Scanlon (1998), who restricts his version of contractualism to the subset of verdicts that concern *what we owe each other*. However, for a large part of the debate, it is assumed that theories like consequentialism, Kantianism, and contractualism cover the whole realm of the deontic. Consequentialism, Kantianism, and contractualism are not usually thought of as theories that only apply to a very restricted set of situations but as theories that reign over all (or close to all) morally relevant situations.

This means that moral theories, typically, should not be considered local theories. Their scope is certainly closer to that of a complete view of science than to that of a specialized theory in one of the sciences. Even if some theories should fail to produce verdicts for every morally relevant situation, their scope is such as to render the distinction between local and global theories one without a difference in ethics. To put it in a catch-phrase, we might say that, in ethics, the local is global (or something close enough).

This has two very important consequences. First, as I noted above, there is a misguided way of dismissing the significance of Parfit's case study for an inductive case for underdetermination. That objection was that since

¹⁴ Note that this is also another typical formulation of what we have called the *deductive principles model* in ethics. The completeness condition, as well as the deductive model, are further discussed in Lazari-Pawlowska (1991).

Parfit only provides one case of moral underdetermination, that is not very significant, *even if* he is right about it. Scientists would not be bothered by the fact that one correct scientific theory is underdetermined either. The case is different in ethics. Since moral theories are global, this means that if Parfit has indeed identified *the best* versions of the most important traditions and if they are indeed underdetermined, we might well say that our total theory of ethics (at least normative ethics) is underdetermined. That would certainly be of high relevance. Second, as we shall see in Chapter 6, this has dramatic consequences for the metaethical implications of moral underdetermination as well. There are skeptical worries that turn on the fact that our best theories, scientific or moral, are underdetermined. If Parfit has identified those theories, and they are underdetermined, these skeptical worries are thereby strengthened.

Thus, the fact that the local is global in ethics has far-reaching repercussions.

Here to Stay

The third distinction is about temporality. Is the phenomenon of underdetermination in ethics one we face permanently, or is it of a transient nature? Here, interestingly, the three projects diverge.

The case is straightforward when it comes to Parfit. Since his preferred moral theories agree on a set of principles, and all verdicts those theories yield follow from these principles, the underdetermination is permanent. Whatever new situations arise, whatever new thought experiments are conjured up, if we start from the same principles when we assess the scenarios, we will end up with the same verdicts. Parfit does not explicitly address this question. However, some of his remarks on the modal status of moral truths are a strong indication that he considers convergence to be permanent. Thus, Parfit (2011b, p. 489) states that: “[f]undamental normative truths are not about how the actual world happens to be.” If the principles that the three theories come up with are indeed such fundamental moral truths, then they are arguably valid for all possible morally relevant situations.¹⁵

The situation is more interesting regarding consequentializers and deontologizers. The way that the completeness condition is defined above, it

¹⁵ As Nebel (2012, p. 1) points out, this view is a direct upshot of Parfit’s more general metaethical convictions. Non-natural properties cannot be discovered by empirical means, yet if they were merely conditional, that is how they would have to be discovered on Parfit’s view.

entails that moral theories specify the deontic status of *every possible* act of moral relevance. This means that the underdetermination we arrive at is permanent. A quick reflection shows why. The main reason this is so for consequentializing is that, *qua* its recipe, we take over every feature the copied theory considers relevant and integrate it into the consequences of our own consequentialist theory. Yet if the original theory specifies the deontic status for every possible act and the consequentializing theory values *all* morally relevant features in the same way the original non-consequentialist theory does, their evaluation of every act will come out the same in any possible world. Similarly for deontologizing. If the original consequentialist theory entails verdicts for every possible morally relevant case, so will the deontologizing counterpart, since it simply copies the former's verdicts. Underdetermination is thus permanent in these cases, too. At least one consequentializer is abundantly clear that this is how he understands the procedure:

[...] [F]or any plausible nonconsequentialist theory, we can construct a consequentialist theory that yields, in every possible world, the exact same set of deontic verdicts that it yields. (Portmore, 2011, p. 84)¹⁶

Consequentializing, according to Portmore, thus yields permanent forms of underdetermination, and, by the same token, so too would deontologizing.

There is an interesting epistemic issue to be considered here. Even if we assume that moral theories are complete in the sense outlined, it might still be the case that we don't know all the verdicts they yield at the time we try to copy them. In such cases, deontologizing, at least on our proposed recipe, is not guaranteed to yield permanent cases of underdetermination, whereas consequentializing is. The reason for this is as follows. Consequentializing integrates all the morally relevant features of its target theory. By stipulation, it encodes all (and only) the relevant features of a moral theory in its own theory of the good. If a non-consequentialist theory thus applies to acts hitherto nor evaluated, so too does its consequentializing counterpart. Let us say that at some time in the future we are faced with a situation in which we need to evaluate a morally relevant situation of a type that

¹⁶ Compare also Portmore (2022, p. 3). Interestingly, our preferred version of Portmore's thesis above is not conclusive in this matter:

DET: [...] [F]or any remotely plausible nonconsequentialist theory, there is a consequentialist counterpart theory that is deontically equivalent to it such that the two theories are extensionally equivalent with respect to their deontic verdicts. (Portmore, 2011, p.85)

has not been adjudicated so far, perhaps how to treat a form of conscious artificial intelligence, or something of that kind. If the non-consequentialist theory has anything to say about that case, it will be because the case falls under some of its antecedently defined moral principles. Yet if the consequentializing theory indeed encodes all the morally relevant features of that non-consequentialist theory, it will come to the same conclusions about this new case as well.

Deontologizing, by contrast, does not encode all the underlying information about the original consequentialist theory but directly copies the verdicts. This means that the deontologizing theory can come apart from the deontologized theory if as yet unconsidered situations have to be evaluated. Remember that the first move of deontologizing only yields a shopping list of particular directives that have no application to future cases at all. Such a theory (if it is a theory at all) is not able to produce any new verdicts over the ones it copies from the original consequentialist theory. As we have also seen, we can make some progress in this regard by invoking, for instance, the DDE and *prima facie* duties. We arrive at principles of some level of generality that, in principle, can also yield verdicts in future cases. However, these verdicts might come apart from those of the original consequentialist theory for two reasons. First, *prima facie* duties need to be weighed against each other. In order for the deontologizing theory to arrive at the same verdicts as the theory it copies, there need to be specifications as to which duty outweighs which other(s) in which scenario. These specifications, however, won't be tailored for future unforeseen cases. The weighing specifications might thus yield different results than the consequentialist counterpart does. Second, and more fundamentally, we can imagine new scenarios in the future that are not describable in terms of any of the principles the deontologizing theory includes. This would be the case if the new scenario is not of the same kind – for example, breaking a promise, telling a lie, etc. The original consequentialist theory might still apply to the scenario if its theory of the good can be used to evaluate the consequences. In contrast, the deontologizing theory might be silent when it comes to the novel scenario.¹⁷

¹⁷ This is not necessarily so. A deontologist theory might include the additional instruction to treat all actions that do not fall under one of its principles in a consequentialist way, or to flip a coin, or something of that kind. However, there is no guarantee that the coin flip yields the same verdicts as the deontologized theory arrives at, and the fact that the deontologizing theory needs to include such an instruction does not look like a very attractive feature.

That being said, the idea by Stanford (2006), that underdetermination might be a transient but nevertheless recurrent phenomenon, is relevant here. Using the deontologizing strategy, we can immediately construe new deontically equivalent theories to any consequentialist theory entailing new verdicts. Thus, deontologizing is at least guaranteed to produce recurrent underdetermination. Hence there is a strong case to be made that the moral underdetermination we have encountered so far is of a permanent or at least recurrent variety.

Ampliative Underdetermination and Theoretical Virtues

The final distinction, and the most difficult one to assess, is that between *deductive* and *ampliative* underdetermination. Deductive underdetermination holds whenever two theories logically entail the same predictions, whereas the latter additionally presupposes that the theories are equally well supported. Especially interesting with regard to this notion of being supported are the so-called *theoretical virtues*. Under that term, I will subsume any kind of advantageous feature of a theory that goes beyond a theory being extensionally adequate. We can distinguish two kinds of such virtues. First, there are *substantive* virtues. These consist of features of theories that we have to define by reference to specific desiderata we connect with those theories. For example, a scientific theory should arguably avoid being based on any kind of (controversial) value judgment. The same, of course, need not be a desideratum for a moral theory. A moral theory, although being in some regards like any other theory, fulfills a specific role *qua* being a moral theory and in doing so can do a better or worse job. Second, there are what can be called *formal* virtues. These are advantages that would make any kind of theory preferable, no matter its subject matter. Examples are qualities like simplicity, fruitfulness, non-ad-hoc-ness, and the like.

I start this section by looking at two substantive virtues, one that is often said to attach to consequentializing and one that may hold for deontologizing. I then look at one example of a formal virtue and discuss whether there is something general to be said about these. Finally, I outline how Parfit's theories fare with regard to the distinction.

Consequentialism's Virtue: The Compelling Idea

Let me start with the advantage that has caused the most buzz in the consequentializing debate. For some time now, consequentialists have argued that their tradition has a distinct advantage over its rivals. That advantage consists in being compatible with the so-called

Compelling Idea (CI): It is always permissible for every agent to do what will lead to the outcome that is best. (Schroeder, 2007, p. 279)¹⁸

Why is this idea thought to be compelling? Consider a person in a situation where they do not know what morality asks them to do. That person might reason as follows:

“No matter what morality asks from me, I should at least not be acting wrongly if I do what leads to the overall best results. Maybe I will thereby do more than what is strictly asked from me. But, surely, I cannot be wrong if I do what is best overall. It should thus at least be permissible for me to bring about the outcome that is best.”

This line of reasoning appears sound. However, non-consequentialist theories have difficulties respecting it because they include restrictions on actions that do not take consequences into consideration. Most forms of consequentialism do not have this problem.¹⁹ That it is at least permissible to do what leads to the overall best outcomes follows directly from the claim that we should do what leads to the best outcomes.

Some consequentializers have tried to exploit this fact to their advantage. They argue that their algorithmically construed theories manage to combine a conjunction of two desiderata that was unattainable so far. They can account for the common-sense verdicts of non-consequentialism *and* they can uphold the CI. This, they argue, yields a distinct advantage. The new theories should thus be preferred to their original counterparts, and consequentialism wins the battle of the theories.²⁰ Tenenbaum (2014, p. 233) calls these authors *earnest consequentializers* because they do not dismiss the difference between the different traditions. I have not discussed their view so far because, as Portmore (2022, pp. 13–14) observes, they can (and should, I might add) accept that there is deductive underdetermination between the original and the algorithmically construed theories. What they want to add is just that if we consider virtues of theories that go beyond extension, such as the CI, we will see that the consequentializing theories turn out to be more attractive. But is this really the case?

Not so fast, say their critics. Schroeder (2007, p. 266) acknowledges that the new theories have the right formal structure to account for deontologist features, such as constraints and special obligations. The problem lies

¹⁸ This formulation in terms of *permissibility* is often traced back to Scheffler (1982, p. 4). For similar formulations, compare Portmore (2005, p. 98) and Dreier (2011, p. 100). For a critical perspective see Foot (1985) and Hurley (2017).

¹⁹ An exception are rule-consequentialists. Portmore (2005, p. 98) thinks that this might explain why the move from act- to rule-consequentialism has appeared unattractive to many consequentialists.

²⁰ The main proponent of this line of reasoning is Portmore (2011).

elsewhere, namely with the notion of goodness that consequentializers have to introduce in the process. Schroeder points out that we can always introduce a technical notion, as consequentializers do in the form of the *good-relative-to*. However, nothing guarantees that a theory employing this technical notion is still compatible with the CI. In Schroeder's own words:

But that is where I get lost. *Good* and *good for*, after all, are concepts that I can understand. [...] But since I don't understand what "good-relative-to" talk is all about, I don't understand how it could be appealing to think that you shouldn't do something that will be worse-relative-to-you. I don't even understand what that means! (Schroeder, 2007, p. 291)

On Schroeder's account, we thus have a good understanding of what is appealing about the idea that one should not do something that would make things go worse *impartially*. All things being equal, one should always be allowed to do the act that would make the outcomes impartially better. The same is true for the idea that, all else being equal, one should always be allowed to do what will be better for oneself. *Ceteris paribus* it would be stupid not to do so, or, as Schroeder (2007, p. 291) thinks, irrational. However, it is not clear why the same should be true for relativized goodness. If we don't even have a clear understanding of what *good-relative-to* means, then we can hardly know what should be appealing about the thought that one should always be allowed to do things that are better *relative-to-onself*.

Consequentializers have reacted to Schroeder's challenge in two different ways. The first one is to dig in and insist that there is a perfectly ordinary concept of relative goodness we all understand and that validates the CI. Dreier (2011, pp. 102–103) provides an illustration. Think of a chess player who is contemplating their next move. The player tries to anticipate all the countermoves of his opponent. When they reach the conclusion that none of these moves outweighs the value of their own move, they judge the move to be good. But what sense of *good* is at stake here? Clearly, it is not an agent-neutral sense, since they do not suspect that their opponent will also find the move good (in fact, they hope that their opponent does not find the move good). However, it is also not necessarily good for the chess player in the sense that winning the game is tied to greater well-being for the player. Instead, the move is good relative to the chess player's position in the game. In the same way, an action can be morally good relative to me, without at the same time being either good for me or good simpliciter.²¹

²¹ Compare Forcehimes and Semrau (2020, pp. 261–262) for another example.

The second reaction is to modify the CI. If the theories that consequentializing yields cannot account for the CI, maybe that is because we have misconstrued that idea from the beginning. Portmore (2007, p. 42) argues that this is indeed the case and that what is attractive about consequentialism is instead its compatibility with the following principle:

Permissibility-of-Maximizing View (PMV): It is always permissible [...] for an agent to act so as to bring about the highest ranked available outcome, i.e. the outcome that she has better object-given reasons to prefer above all other available alternatives. (Portmore, 2007, p. 50)²²

The mistake with the old version of the CI, so Portmore argues, was that it had a commitment to an agent-neutral conception of goodness. PMV, in contrast, leaves open the question of what constitutes the goodness of an outcome. Instead, it only states that it is always permissible to maximize *whatever* the highest ranked outcomes turn out to be for an agent.

Whether (one of) these two reactions to Schroeder's challenge ultimately succeeds is a question that remains controversial.²³ But it suffices to show that proponents of consequentializing do indeed point to advantages of their theories that go beyond deontic adequacy. This is what the framework of underdetermination would make one expect. If the ability to account for our judgments and intuitions about particular cases does not decide between two theories, it is natural to look for other advantages. This can be framed as a debate as to whether, in addition to deductive underdetermination, we are also facing ampliative underdetermination in ethics. The suggestion of the *Compelling Idea*, if it can be defended, would speak against such a form of ampliative underdetermination, since the CI would give us an additional reason to prefer consequentialism. However, that is just the beginning of the debate.

Deontology's Virtue: Independence from Axiology

Many consequentializers (and consequentialists) take it for granted that consequentialism would be the more attractive theory were it not for its

²² More precisely, Portmore (2011, p. 58) later holds that what's compelling about consequentialism is a specific conception of reasons, the *Teleological Conception of Reasons*. I'll gloss over this since everything I say about the PMV also holds for the TCR.

²³ For a critical assessment, compare Hurley (2014, 2017) and Betzler and Schroth (2019). For a more positive assessment, compare Suikkanen (2009a).

counterintuitive extensional upshots.²⁴ However, couldn't deontologists point to similar theoretical advantages of their theories?

One such advantage comes to the fore in connection with consequentializers' treatment of moral restrictions. Emet (2010, pp. 4 ff.) criticizes consequentializers' account of restrictions as being too egocentric. He provides the following example. If consequentializers want to make sure that person A does not kill person B, even if B would kill several other people, then person A needs to put a vastly greater disvalue on their own killings than on B's (or any other's). In other words, person A needs to claim that *relative-to-them*, their own killings are much worse than other people's. Yet this, as Emet (2010, p. 6) sees it, overestimates A's importance: "[w]hen couched in terms of agent-relative value, agent-relative restrictions seem to represent a kind of moral fastidiousness." The idea is that on a consequentializer's account, person A is so obsessed with their own moral cleanliness that they put an improbably high disvalue on their own misdeeds.²⁵

In contrast, and this is where the positive argument for deontologist theories starts, Emet (2010, p. 6) thinks that non-consequentialists have a better explanation at their disposal. That explanation is inspired by the Kantian idea of inviolability. According to this idea, the force of the moral restriction on killing does not come from the disvalue we attribute to our own breaking of that restriction but instead from the dignity of the other person. There are acts we are simply not allowed to do to another person, even if that would lead to fewer people suffering from the same kinds of acts. We thus avoid having to claim that one kind of killing has much more negative value than the other. Although the deontic result is the same, Emet (2010, p. 7) thinks that the Kantian explanation is therefore to be preferred: "[i]f one seeks to defend restrictions, rather than to merely make formal space for them, an account like the Kantian one mentioned here would be preferable." Emet thus identifies a theoretical advantage of deontological theories that would persist even if we grant deductive underdetermination between his Kantian and a consequentialist counter theory.

This advantage is specific to the Kantian version that Emet prefers. However, the underlying issue is a more general one, and it has been noticed by other commentators as well. Consequentialists often think that it is an advantage of consequentialism that it rests on a theory of the good.

²⁴ Critics of consequentializing sometimes seem to tacitly accept this when they only offer negative arguments against consequentializing instead of making a more positive claim for deontologist theories themselves.

²⁵ For similar points, compare Schroeder (2017, p. 1479) and Howard (2021).

This might, for example, allow for a better empirical assessability of moral questions.²⁶ But some critics have argued that the tables can be turned on consequentializers. Here is how Woodard puts it:

[The consequentializers'] theory of value must distinguish acts according to when and by whom they are performed, and assign greatly different values to such acts even when they are, in other respects, morally very similar. Such a theory would make claims about the value of actions that, taken by themselves, would be very hard to believe. (Woodard, 2013, p. 262)

Non-consequentialists, by contrast, do not have to subscribe to such an implausible theory of value. Since the rightness of an act is not exclusively dependent on its goodness, they can uphold a comparably orthodox theory of value that does not involve the strange weightings. The idea is thus that precisely because non-consequentialists' theory of the good remains at least partly independent of what they prescribe, it does not have to bend for every deontic verdict. We might call this virtue *Independence from Axiology*. Deontological theories can make intuitively plausible claims about the value of acts forbidden by restrictions because they can hold that it is not always the (negative) value that grounds the restriction. Hence, if consequentializers argue that their theory is to be preferred because it can account for all the same verdicts and, on top of that, can uphold the Compelling Idea, deontologists can reply that this is too fast since their theory has theoretical virtues as well.

So far, we have only established that deontologists have a way to counter consequentializers' attack. But couldn't deontologizers attempt something more ambitious? Couldn't they try to argue that deontology is in fact the superior tradition since, if deontologizing succeeds, they can account for all the consequentialist's virtues *and* have the advantage of being independent from axiology? I think that this move would once more be too fast, and the reason is the same as in the consequentializing case. (Partial) Independence from axiology is a virtue that classical deontologist theories have with regard to their consequentializing counterparts. Remember that it is the adoption of an agent-relative notion of goodness that makes the consequentializers' theory of the good (presumably) unattractive. Yet classical consequentialist theories do not buy into the notion of relative goodness. Hence, deontologizers, when they compare their deontologizing theories to such classical consequentialist traditions, do not have that advantage. Indeed, it isn't even clear that deontologizing theories can uphold the virtue of *Independence*

²⁶ Compare Betzler and Schroth (2019) for a discussion of this point as it relates to consequentializing.

from *Axiology*. The reason for this is that a deontologizing theory has to copy the verdicts of a consequentialist theory. That is its purpose. But now assume that the consequentialist theory that needs to be deontologized requests, for example, that we kill somebody in order to save fifty other individuals. The deontologizing theory has to copy this. In other words, the deontologizing theory tells us to do exactly the kinds of things to a person that on Emet's deontological theory are proscribed by the inviolability of persons. The deontologizing theory hence cannot appeal to *this* explanation via the inviolability of persons, and it cannot claim for itself any of the theoretical virtues that come with this explanation.

A pattern emerges. Classical consequentialist and deontologist theories are often thought to have distinct theoretical virtues. However, we may not simply assume that these virtues can be upheld by consequentializing or deontologizing theories as well, since, in the process of copying their rivals' verdicts, these theories are changed in a way that might impact their theoretical virtues. Whatever advantages and disadvantages classical versions of the moral traditions have, we always need to consider whether consequentializers and deontologizers can avail themselves of those as well, and whether they do better with regard to those virtues than the theory they attempt to copy.

Simplicity as a Formal Virtue

Besides substantive virtues, moral theories also exhibit *formal virtues*. These are the kinds of advantages a theory can have independent of what it is a theory of. Since discussing all such virtues in detail would be too much for this section, I instead want to consider one exemplary virtue. My hunch is that most of what applies here can be generalized to other formal virtues, too.

Perhaps the most discussed theoretical virtue in the philosophy of science is simplicity. It has seemed plausible to many philosophers that, all things being equal, it is a positive feature of a theory to be relatively simple.²⁷ The same might be true for ethics. A moral theory is *simple*, according to Kagan, if it:

[...] [Y]ields a body of judgments out of a relatively sparse amount of theory, deriving the numerous complex variations of the phenomena from a smaller number of basic principles. (Kagan, 1989, p. 11)

²⁷ I postpone the question whether simplicity and other formal virtues are indeed indicative of the truth of a theory or merely pragmatically advantageous to Part III.

How do consequentializing and deontologizing theories fare in this regard? One might think that *if* a consequentializing or deontologizing theory can account for all the verdicts of another theory *and* do so in a much simpler way, that might provide a decisive advantage. Yet, once again, it is not at all obvious whether this is the case.

This has not been lost in the consequentializing debate. Some versions of consequentialism might be thought to be especially simple. The prime example, perhaps, is hedonistic act-utilitarianism. That theory rests on a relatively simple explanatory principle, that is, that an act is good if and only if and because it maximizes aggregate utility, combined with a relatively simple theory of the good, that is, that pain and pleasure are all that matter. Yet, as Sachs (2010, p. 265) and Schroeder (2007, pp. 288–290) both argue, consequentializers lose much of this simplicity. By incorporating the verdicts of their target theories, consequentializers also complicate their own theories. In order to account for non-consequentialist verdicts, consequentializers, according to their own declared recipe, need to include *everything* that a non-consequentialist theory considers relevant into their own conception of the good. This implies that their theory of the good will be considerably bloated. How extensive the losses in simplicity are for consequentialist counterpart theories depends on what its target theory is. It depends on how many features the original non-consequentialist theory considers relevant, which then have to be included in the consequentialist theory of the good. Thus, we cannot give an overall estimation of what the losses in simplicity amount to for consequentializing. We can, however, state with certainty that the consequentializing theory will not be as simple as some classical versions of that tradition are.

Something similar is true for the theories resulting from deontologizing. Classical deontologist theories, such as ten-commandment-style theories, have an elegant simplicity in the relatively small body of prescriptions and proscriptions they rely on. Yet this is not guaranteed for the newly created theories that are the result of deontologizing. We have already touched on this issue when discussing the problems of the shopping-list theories in Chapter 4. Deontologizing, at least according to the recipe I have suggested, entails that we construe an extensive list of DOs and DONT's for all the verdicts of the consequentialist counterpart. Such theories would score very low on simplicity.²⁸ I have further suggested that, using such tools as the

²⁸ Sachs (2010, p. 265) himself acknowledges this: “[...] ten commandments-style theories, as Shelly Kagan has pointed out, lack simplicity, power and coherence.”

notion of *prima facie* duties, deontologizers can move some way toward a sparser theory with fewer principles. However, we also saw that there is a need to specify when a *prima facie* duty overrules another. For this, we need additional rules that specify how to weigh different duties. Arriving at a relatively small set of principles might thus come at the price of a more extensive handbook for how to use them. How grave the loss in simplicity will be depends on what theory is being deontologized. If the consequentialist theory that needs to be deontologized leads to more or less commonsensical verdicts, that is, requires us not to break promises etc., the task is arguably easier. A small set of principles and rules on how to weigh them should do the trick for deontologizers. However, if the consequentialist target theory leads to less commonsensical verdicts, we will need a bigger set of rules to specify when one principle outweighs another. This can result in quite byzantine theories with ever more complicated rules on the weighting of principles.

In sum, both consequentializing and deontologizing are subject to losses in simplicity when compared to more typical versions of their traditions, since they need to make room for the verdicts of the theories they are copying. How they fare with regard to these original theories, whether they come out more or less simple, is thus difficult to assess on a wholesale basis.²⁹

Instead, I think that there is no way around judging the relative plausibility on a case-by-case basis. As Norton puts it with respect to scientific theories:

General claims on the relative weight of evidential and other factors in the determination of scientific theories will need to be supported by careful scrutiny of the particular cases at hand. There are no shortcuts. (Norton, 2008, p. 40)

The same is true for moral theories. Arguments to the effect that we can simply copy the verdicts of some other theory *and* retain the advantages of our own theoretical framework are to be regarded with suspicion. The story that we could simply modify our theories to the point where they can adopt all their rival's verdicts and at the same time uphold all the advantages of more classical forms of the respective tradition has struck many critics

²⁹ The picture is further complicated by the fact that there can be different kinds of simplicity at stake, which can be in conflict. What we have just discussed about deontologizing is a case in point. A theory might be simple when it comes to the *prima facie* principles it relies on but complicated as concerns the weighting manual of these principles. Estimating relative simplicity might thus be complicated not only across different theories but even within a theory.

as being too good to be true, and rightly so. Instead, when trying to copy the verdicts of other traditions, we might lose some of the virtues of our tradition, and there will presumably also be trade-offs between different virtues. The question of ampliative underdetermination cannot be decided on a wholesale basis. A more case-by-case investigation is imperative.

A Case Study of Ampliative Underdetermination

This finally brings us back to our more detailed case study: Parfit's three converging theories. How do these theories compare to each other when it comes to criteria that go beyond extension?

Since an exhaustive discussion of potential virtues, both substantive and formal, goes beyond what I can do here, I will only be able to make an informed guess. However, I think that some general observations about Parfit's preferred theories make it highly likely that we are indeed presented with an example of ampliative underdetermination.

The most important point to note is that Parfit's theories aren't only equivalent when it comes to their extensions, that is, the particular deontic verdicts they yield. The three theories also agree on the set of mid-level principles from which these particular verdicts follow. Thus, on these two levels, there can't be any advantage to one of the traditions. The only point where Parfit's three theories do differ is with regard to the fundamental explanatory claim regarding what makes acts right or wrong. Here we have three different claims. The Kantian theory holds that acts are right if and only if and because they follow from principles whose universal acceptance everyone could rationally will. The Scanlonian theory holds that acts are right if and only if and because they are allowed by some principle that no one could reasonably reject. Finally, the consequentialist theory holds that acts are right if and only if and because they follow from optimific principles. Hence, if one of the theories were to have a decisive theoretical advantage that would break the ampliative underdetermination, it would have to be regarding this principle.

Yet that seems highly unlikely. First, consider formal virtues. The one we have been looking at in some detail is *simplicity*. Is one of Parfit's theories more simple? It does not seem so. Parfit's three theories arrive at the same set of deontic principles about which acts are right or wrong. They are thus equally simple in terms of the mid-level principles they require. Furthermore, their main explanatory principles also seem similarly simple. Hence, simplicity doesn't look like a promising tiebreaker. What about other formal virtues? Some of the examples discussed in the philosophy of science are extensional adequacy, consistency, and fruitfulness, that is,

the ability to make novel predictions. Yet these virtues, when transported to ethics, don't seem more pronounced in one of Parfit's theories than another. First, Parfit's theories are extensionally equivalent, hence they are, by definition, equally adequate when it comes to their extensions. Second, there is no indication that the explanatory principle of one of Parfit's theories is either inconsistent in itself or inconsistent with the set of mid-level principles all theories share. Third, since the three theories contain the same set of principles, and all verdicts that follow from the theories follow from these principles, none of them can make novel predictions, that is, tell us which future acts are right or wrong, that the others cannot. Generalizing from this, it seems highly unlikely that, given the structural similarities of Parfit's three preferred theories, one of them has a decisive formal advantage over the others.

What about substantive virtues? Here, the case is less clear. Many proponents of one of these traditions presumably think that their tradition has such an advantage. Indeed, apart from deontic adequacy, this is what arguably attracts philosophers to one or another of the traditions. The two substantive virtues we have discussed above, the *Compelling Idea* and *Independence from Axiology*, are cases in point. Do Parfit's three theories differ with regard to these two virtues? The answer is no. First, Parfit's preferred version of consequentialism is a version of rule-consequentialism. Yet rule-consequentialism is *not* compatible with the Compelling Idea. According to rule-consequentialism, whether an agent is allowed to perform some act depends *not* on whether this act itself leads to the optimific result but on whether this act follows from optimific principles. This, as rule-consequentialists are aware of, entails that we are sometimes forbidden from acting in a way that would, in this very instance, lead to the best outcomes. Yet this means that, according to rule-consequentialism, it is not always permissible for every agent to do what will lead to the outcome that is best, which is what the CI requires. Parfit's preferred consequentialist theory thus does not have the advantage of being compatible with CI.

Second, the two non-consequentialist theories in Parfit's case study don't have the advantage of *Independence from Axiology* either. Or, rather, it is not clear that this even constitutes an advantage relative to the rule-consequentialist theory. The reason for this is as follows. Remember that the disadvantage that consequentializers have when compared to the non-consequentialist theories that they copy is that the agent-relative notion of goodness they have to adopt in order to copy these theories is independently implausible. It states, for example, that one's own breaking of a promise has a much higher disvalue than anyone else's, prompting the charge of moral

fastidiousness. Yet (Parfit's version of) rule-consequentialism does not have to buy into such an implausible theory of the good. It only claims that prohibitions on breaking a promise are part of the optimific rules. Hence, Parfit's two non-consequentialist theories don't have this advantage over the consequentialist theory, either.

Of course, these two virtues are not the only ones we can think of. Since many more substantive virtues might be claimed for one or another of the main traditions, we cannot make any definitive assessments here. Yet, there is one more observation that should make us skeptical that such a decisive advantage could be detected for one of Parfit's preferred theories. This observation is that Parfit's preferred theories are, more or less, classical variants of three of the main traditions of moral theorizing. Yet the advantages of these traditions, perhaps contrary to the potential advantages of less classical variants like agent-relative consequentialism, have been discussed thoroughly for a long time now, and yet none have proved successful in convincing proponents of the other traditions. Indeed, how we should compare different virtues (for example, being compatible with the Compelling Idea or Independence from Axiology) is a question for which, to my knowledge, normative ethical theorizing has not yet found a compelling answer. Proponents of the different traditions remain unconvinced of their rivals' arguments, which is why they often, at some stage, turn to (alleged) extensional differences. However, in the case of Parfit's theories, this is not viable, of course, since the theories are extensionally equivalent.

In sum, we can state that moral underdetermination, just like its cousin in science, is a multifaceted phenomenon. We have encountered examples of most of the versions discussed in the philosophy of science: existence and non-uniqueness, permanent and transient, deductive and ampliative. However, at least as interesting have been the disanalogies. Thus, we have not seen versions of the egalitarian variety and, perhaps most consequential, we have seen that the distinction between local and global theories cuts no ice in ethics. This, as we shall shortly see, is highly significant when it comes to the repercussions of moral underdetermination for the realm of metaethics.

5.3 **Where to Go from Here**

Before we turn to this, I want to leave normative ethics on a more programmatic note. I'll thus finish the chapter with three suggestions where

I hope the study of moral underdetermination might positively impact normative moral theorizing.

Against Falsification and Crucial (Thought) Experiments

The first lesson should be the one for which underdetermination is most prominently known in science (that is, besides its consequences for the realism debate): its upshot against a simplistic form of falsificationism. As Harding (1976, pp. ix ff.) explains, one of the main lessons of both Duhem and Quine is that in addition to the fact that theories cannot be verified, they also cannot be falsified, at least not in a simple way. The defining term in this context is the one of a *crucial experiment*. A simplistic, or *naïve*, form of falsificationism holds that whenever we are unsure which theory is correct, we can construe experiments that deliver new evidence that falsifies all the wrong theories, ultimately leaving only one to be the winner per elimination. The phenomenon of underdetermination speaks against this view. In cases where theories are permanently underdetermined, no forthcoming experiment will be able to decide between the rival theories. The naïve form of falsificationism thus does not bring us any further.

I think that underdetermination in ethics holds a similar lesson. Debates in normative ethics often involve reference to particular cases that are supposed to falsify one or another of the traditions. Known cases are the charge that consequentialism requires the harvesting of organs from one healthy person for the benefit of several sick ones, or the reproach that Kantianism requires us to tell the whereabouts of an innocent person to a henchman at the door. Consequentializers and deontologizers show that this kind of argument is insufficient. There are versions of the different traditions that can account for the verdicts in question, at least if they are minimally plausible.³⁰ Consequentializing and deontologizing thus render the two traditions immune to such kinds of crucial thought experiments. This, we have seen, is indeed what underlies the approach to consequentializing by Portmore (2011). He is explicit that his main goal is to come up with a version of consequentialism that is immune to the kinds of objections that are often raised in order to undermine consequentialism. Of course, this does not end the discussion, as Portmore is well aware. As he correctly predicts, the debate then moves on to other factors like the

³⁰ Which they need to be, or the whole strategy of falsifying theories by proving that they cannot account for the verdicts fails.

Compelling Idea. Yet this complicates matters. We can no longer discard whole traditions of theories on the sole basis that they lead to the wrong verdicts.

It would be uncharitable to charge ethicists with adherence to a naïve form of falsificationism. Most of them are certainly aware of the fact that we almost never falsify a moral tradition with one counter-example. We should therefore be more precise about what the lesson is. It is *not* that discussions in ethics are too often committed to a naïve method of falsification and that bringing in the problem of underdetermination serves to teach ethicists how to better conduct their business. There are highly sophisticated discussions that point precisely to the problems on a theoretical level, as consequentializers' discussion of the Compelling Idea illustrates. The lesson that underdetermination provides comes rather in the form of a systematic rationale for why this more sophisticated way of arguing is not only desirable, but necessary. Consequentializing and deontologizing at a minimum succeed in establishing deductive underdetermination for all plausible moral theories. This means that different traditions of theories can account for all sets of plausible deontic verdicts. We therefore cannot resolve our theoretical quarrels by referring to counter examples alone. Thus, the idea here is not so much to get ethicists to turn away from naïve falsification. This step has mostly been taken already. Instead, I think that underdetermination provides a systematic account for why this naïve falsification cannot work as a method.

A Test Tube for Theoretical Virtues

The second upshot is for our understanding of the role of theoretical virtues in theory choice. The scientific debate has some systematic guidance to offer in this regard. Carrier (2011) suggests that one of the primary functions of underdetermination is to serve as a test tube that can lay open the non-empirical virtues that play a role in theory choice. The idea is simple. When facing a case of underdetermination, scientists cannot make a decision based on the data. Yet, according to Carrier, observation shows that scientists typically do prefer one theory over another. Presumably, this is on the basis of other factors, the non-empirical virtues of theories. Thus, as soon as two theories are underdetermined by the empirical evidence, we can gain insights into what guides scientists' theory choice beyond the data. Carrier (2011, p. 197) thinks that these virtues often operate in a hidden way, only becoming visible when the otherwise dominant factor of empirical adequacy does not do the trick. He likens the case to a solar eclipse, in

which we are suddenly able to see stars close to the sun that were formerly undetectable. Underdetermination, Carrier goes on, plays a fruitful role in helping us investigate what criteria influence scientific inquiry. In doing so, it deepens our understanding of scientific rationality in itself.

The application to ethics is straightforward. Studying cases of moral underdetermination should yield valuable insights into the non-deontic virtues of moral theories. In cases of underdetermination, the deontic verdicts a theory yields do not decide the case for or against it. Suddenly, we therefore have to turn to other features that we ordinarily do not pay attention to. Underdetermination thus works like a magnifying lens to put these features front and center. When all extensional differences are put aside, the most natural next step is to focus more closely on what other differences remain.

Our observations regarding consequentializing and deontologizing have confirmed that this is indeed what happens. When deontic agreement is granted, other features, such as the *Compelling Idea* or *Independence from Axiology*, as well as formal virtues like simplicity and fruitfulness, come to the forefront. We have looked at some of these features in quite some detail already. Of course, the more widespread the phenomenon is and the more ethicists are working on it, the more impressions we can collect. Using moral underdetermination as a test tube will thus presumably tell us a lot about which theoretical virtues ethicists cherish and how they rank them.

One might think that there is a disanalogy to the scientific case that impairs the success of such a study design in ethics. Carrier (2011, pp. 197 ff.) thinks that underdetermination provides an ideal setting for the philosopher of science to study real, working scientists' preferences in theoretical virtues. In an ideal case, this would yield untainted insights into the preferences of scientists. In our case, in contrast, we are observing the decisions of ethicists who are already heavily invested in one theory or another. Consequentializers and deontologizers, because of the stake they have in their moral views, might not be the best informants on what worth to accord different non-deontic virtues. It might therefore be doubtful whether we can expect to get any unbiased results.

I do not think that this difference carries much weight, though. First, it is less than clear to me that working scientists are disinterested pragmatists, ready to switch from one theory to another based solely on how they evaluate theories' non-empirical virtues. Some investment into a long-held theory will probably impede an unbiased choice in science as well. Second, and more importantly, we are not restricted to simply recording what choices are being made by ethicists. We can critically assess these

choices. Thus, if it turns out, for example, that deontologists are forced to deny ever more glaring advantages of their opponents' theories, we do not have to take their word for it that their theories are equally attractive. Underdetermination might reveal a very unlevel playing field when it comes to theoretical virtues. The important point here is simply that underdetermination serves as a highly efficient means to find this out.

For what it's worth, I do think that further investigation into the theoretical virtues in (normative) ethics would be of value. I do not think that they have received the attention that they deserve. In ethics, we generally take for granted that the different theoretical traditions arrive at (very) different conclusions about deontic verdicts. This renders it unnecessary to search for further grounds on which to decide between them. One especially interesting question that moral underdetermination opens up is precisely whether some theoretical virtues might break the tie between the rival traditions.

Importantly for the purposes of investigating the non-deontic preferences in theory choice, it does not matter whether the rival theories we come up with are indeed plausible. As Carrier (2011, p. 198) informs us, the criteria that are made explicit by refuting an extensionally equivalent rival theory may merit attention no matter how implausible that rival theory is. This means that the idea of a test tube to investigate the non-deontic virtues of theories should be welcomed even by ethicists who are not ultimately convinced that there are ampliative forms of underdetermination in ethics. Bracketing out the deontic differences will help us all appreciate the theoretical advantages of our theories more clearly. Thus, even if consequentializing, deontologizing, and Parfit ultimately fail to provide acceptable theories, we will learn a lot about the non-deontic factors that make theories acceptable in the process.

A Demand for an Independent Justification of Moral Theorizing

A general theme of this book is that extension is not all that matters in moral theories. Thinking in terms of underdetermination puts those disagreements that go beyond extension front and center. The first two points seemingly strengthen this turn to the theoretical side of moral theorizing. If our considered judgments and intuitions cannot function in the manner of a crucial experiment to refute theories, that plausibly strengthens the status of explanatory claims. We are not automatically driven to give up our preferred explanatory framework because it might

at first look to have been falsified by some counter example. Furthermore, if underdetermination is used as a test tube for the theoretical virtues, that arguably strengthens their status.

However, and contrary to this general theme, moral underdetermination also holds more critical lessons for the explanatory side of moral theorizing. The main critical import is metaethical and will be the topic of Part III. Yet some of it also pertains to normative ethics. More precisely, moral underdetermination poses a challenge to ethicists to come up with an independent reason to justify their interest in the explanatory side, and in ethical theorizing more generally. One of the main strategies to justify engaging with ethics in an in-depth way, I take it, is by referring to its practical upshots. Ethics, it is sometimes thought, has a unique relevance in (academic) philosophizing because it comes to bear very directly on everyday issues. We all want to know which acts are right or wrong, how to live better lives, and so on. This practical importance, or so it would seem, lends justification also to the more theoretical parts of ethical theorizing. If the different traditions of moral theorizing arrive at different answers as to what is right or wrong, surely we are justified in studying them very closely. To return to two of our examples, if one tradition leads us to harvest innocent people for their organs, or if another one leads us to betray people because we are morally unable to tell a lie, that makes it seem quite important that we deal with these theories. The theoretical side, and with it the explanatory aspect, thus gains some of its justification via the practical side. Because doing the right thing is important, and because the different traditions provide different answers to what doing the right thing is, investigating the latter is important too.

Underdetermination challenges this simple connection. If it turns out that those explanatory frameworks can arrive at the same verdicts, one might wonder why we should even care about them. If underdetermination holds for our best theories, then from a purely practical point of view, it might not matter which theory is the correct one. It might still be that one theoretical framework, though no more justified than its rivals, is more effective in getting beings like us to do the right thing. Thus, from this practical perspective, one theory might still be preferable to others. However, that is not because it yields better verdicts than its rivals but just because it does so in a way that speaks more effectively to us. We would thus only have to conduct some empirical studies about which moral theories better serve the goals of morality. Yet that means that we cannot justify our engaging with those theories by referring to their practical upshots. Ethicists

might thus need to look for other grounds to justify their occupation with questions regarding moral explanation and the like.

Note that I am not claiming that ethicists do not have a way to answer these demands for justification. Moral explanation plays a big role in our everyday moral practice. Analyzing this practice on a philosophical level surely has some justification of its own. We can better our understanding of what constitutes good and bad reasons, better or worse arguments, and so on. Yet, in the face of underdetermination, I think that there is some pressure on ethicists to refer more explicitly to these issues when staking out their contribution. The easiest justification, according to which theoretical issues (typically) have some practical consequence, is blocked by the phenomenon of underdetermination.

