

Going Algorithmic

Consequentializing and Deontologizing

Considering its sheer length and depth, Parfit's *On What Matters* is as valuable and insightful for the study of moral underdetermination as any case study might arguably be. However, as impressive as it may be, it is just that: one case study.¹ There are many philosophers who don't accept Parfit's results about the convergence of the moral traditions and hence won't accept it as an instance of moral underdetermination. It is thus natural to inquire whether there are more general considerations, beyond single examples, that indicate that moral underdetermination might be a more pervasive phenomenon. As our discussion of the different strategies in science in Chapter 1 has shown, this is what some defenders of underdetermination in science have done when they started construing algorithms that would prove underdetermination for any scientific theory.

Interestingly, we find something strikingly similar going on in two recent projects in normative ethics, so-called *consequentializing* and *deontologizing*. Proponents of consequentializing, or *consequentializers* for short, claim that for any non-consequentialist theory, they can come up with a consequentialist counterpart that is deontically equivalent, that is, leads to the same set of verdicts about particular cases.² This has brought another group to the fore, so-called *deontologizers*, who promise the same for their tradition, that is, construing deontically equivalent non-consequentialist counterparts for any consequentialist theory.

¹ I will argue in Chapter 5 that considerations having to do with the *completeness* of moral theories potentially make Parfit's argument even more valuable than one might expect given what has been said so far.

² A small note on terminology. I will use the predicates *original* or *target* when I refer to the non-consequentialist theory, which consequentializers attempt to copy. I speak of the *new* or the *counterpart* theory to refer to the consequentialist theory that is produced in order to copy the original theory's verdicts. I will also sometimes speak of the *consequentialized* theory to refer to the original non-consequentialist one and of the *consequentializing* theory to refer to its consequentialist counterpart. The same terminology – *mutatis mutandis* – will be used when discussing deontologizing.

As with Parfit, these philosophers do not themselves think of their projects in terms of underdetermination. If their projects are to serve as further examples of underdetermination, I will thus once more have to proceed in two separate steps. First, I will outline what we can call the *technical* side of the two projects.³ This concerns the *recipes* that have been proposed in order to construe deontically equivalent counter theories. In both cases, I will rely mostly on what proponents of these projects have suggested, although, since deontologizing is less of a unified project, I will have to flesh out some of those suggestions in more detail myself. Second, I will argue that *pace* how consequentializers and deontologizers interpret their own results, we should best understand these results in terms of moral underdetermination. If this is correct, we are presented with a case of moral underdetermination that is much more far-reaching than Parfit's.⁴

4.1 Consequentializing: A Simple Recipe and a Two-Step Procedure

Let us start with the older of the two projects, consequentializing. Its point of departure lies in an observation by Dreier (1993). While analyzing the structural difference between consequentialist and non-consequentialist theories, Dreier mentions the possibility that this difference might not be as deep as is commonly thought. The reasons for this, according to Dreier, is that we can copy the content of any (plausible) non-consequentialist theory by following a simple recipe:

We merely take the features of an action that the [non-consequentialist] theory considers to be relevant, and build them into the consequences. For example, if a theory says that promises are not to be broken, then we restate this requirement: that a promise has been broken is a bad consequence. (Dreier, 1993, p. 23)

Dreier calls this move *consequentializing*. He does not, in that article, provide more details as to how we might render the figurative speech about the building of features of actions into consequences into a more precise thesis. Yet others have taken up this task, and an especially promising way of doing so has turned out to be in terms of sets of particular deontic verdicts. Such sets can be thought of as the output, or, in another, much used term, the extension of a theory. Using this terminology, we can formulate Dreier's

³ Schroeder (2017, p. 1476) also aptly speaks of the *mechanics* of consequentializing.

⁴ Some of the material in this chapter is drawn from Baumann (2019).

recipe in the form of a more precise thesis. Here is Douglas Portmore's phrasing:

Deontic Equivalence Thesis (short: *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)

Compare this to Dreier's more recent formulation:

[...] [T]he *Extensional Equivalence Thesis* states that each plausible moral view has a consequentialist extensional equivalent: the two views agree on the deontic status of every act. (Dreier, 2011, p. 98)

Consequentializing can thus be understood as a procedure that aims to prove deontic equivalence, defined in terms of the set of deontic verdicts that is yielded by a theory, between any plausible non-consequentialist theory and some corresponding version of consequentialism.⁵ Note that this does not mean that equivalence must be proven between all non-consequentialist theories and *one* version of consequentialism. Instead, we could, in principle, come up with as many different consequentialist counterparts as there are versions of non-consequentialism, and, indeed, quite different versions have been proposed. However, it should also be noted from the start that all of the consequentializers work within an *act-consequentialist* framework. Unless specified otherwise, we should thus always understand consequentialism as act-consequentialism in the context of consequentializing.⁶

If proven true, the Deontic Equivalence Thesis amounts to a direct challenge to the first tenet of the Textbook View. If the verdicts of any version of non-consequentialism can indeed be accounted for by some version of consequentialism, this contradicts the assumption that those alternative traditions disagree regarding their extension.⁷ Yet how are we to turn the simple recipe that Dreier mentions into a more detailed strategy that could prove the highly ambitious consequentializing thesis? This is the *technical* side of consequentializing.

⁵ We will shortly see what the caveat "plausible" theory amounts to. Consequentializers mostly have in mind structural features, such as a theory allowing for dilemmas. In addition, there is presumably also a connection between the plausibility of a theory and the plausibility of the verdicts it yields. However, consequentializers will be anxious not to let the class of implausible theories become too big, lest their thesis becomes uninteresting.

⁶ Indeed, Portmore (2009, p. 341) goes so far as to state that on his understanding, rule-consequentialism counts as a form of non-consequentialism.

⁷ Most consequentializers think that it also challenges the second tenet, as we'll see in the second part of this chapter.

The Technical Side of Consequentializing

How does one consequentialize? As I understand it, there are two main moves.

Move I: Building the Non-consequentialist Features into the Consequences

First, we need to somehow *build* the non-consequentialist features into the consequences of acts. How can this be achieved? Basically, there are two options.

The first option involves a reinterpretation of what count as the *causal* consequences of acts. This option had been discussed before consequentializers came along. An early advocate is Oldenquist (1966). Oldenquist thinks that we cannot distinguish acts from their consequences in a way that would make this distinction morally relevant. His reason for this is that:

Any attempt to specify what a person did involves reference to what happened, in other words, to what could be called consequences, however direct and immediate they might be. (Oldenquist, 1966, p. 181)

For example, Oldenquist insists that lying can be seen as the consequence of speaking in a certain way. Speaking, in turn, can be viewed as the consequence of a set of intentions and bodily activities. We can see how it goes on. Oldenquist thinks that we can, in this way, carve any event into ever more fine-grained events:

In general, almost any occurrence or set of occurrences, having a conscious human source, may, at the sole cost of stretching certain conventional ways of speaking, be called either an action or the consequence of an action. (Oldenquist, 1966, pp. 181–182)

If Oldenquist is right, any morally relevant feature that non-consequentialists identify can be reinterpreted as a causal consequence of some preceding act, allowing consequentialists to account for it. The suggestion should thus look appealing to consequentializers.

Or so it would seem. Widening the causal conception of consequences in this way has clear downsides. First, Atwell (1969) objects that Oldenquist's reasoning is based on a wrong model of moral action. For instance, the act of cheating, on this model, would be carved up into a sequence of events, such as writing down anticipated answers on cards and looking at them during an exam. However, this analysis seems inaccurate. An act may well be an instance of cheating because it is comprised of these events. But it is not identical to any of these events, nor to any sequence of them. Second, as Oldenquist himself acknowledges, this model comes at the price

of stretching the conventional ways of speaking. Yet this is putting it too mildly. Oldenquist's suggestion is closer to a redefinition of the concept of a consequence than a stretching of it. We would therefore be changing the subject instead of giving an answer to the original question about the differences between consequentialism and non-consequentialism.

Similar thoughts have led contemporary consequentializers to choose a second option. Instead of implausibly widening what count as causal consequences, consequentializers opt for what Dreier (2011, p. 99) calls a *constitutive* understanding of consequences. He gives the example of a person aiming to perfect her badminton serve. This person could value the causal consequences that consist of heightened precision and increased fitness. However, the person could also cherish practicing badminton itself as part of what makes her life valuable. She would consider practicing and perfecting her sport as a constitutive part of what makes her life a fulfilled one. The positive consequences thus do not (entirely) consist in the effects of her practice, but in the fact that she has practiced, which makes her life more valuable.

In the same vein, we can build the relevant non-consequentialist features into consequences. Again, Dreier points the way:

[T]hat I have told a lie is a consequence of my lying, but it is not a causal consequence. We might call it a constitutive consequence, for want of a better term. (Dreier, 2011, p. 99)

This talk of constitutive consequences might still sound very technical. Yet there are intuitive ways of understanding the relevant notion. Here are three:

An act's outcome is [also] to be construed broadly so as to include everything that would be the case were the act to be performed. (Portmore, 2007, p. 39)

Broadly construed, an act's outcome is the possible world that would be actual if it were to be performed. (Portmore, 2009, p. 330)

The consequences of an act can be conceived of as the entire state of affairs brought about by the act. (Peterson, 2013, p. 172)

What is common to all these formulations is a very broad understanding of consequences that can include the causal as well as the constitutive notion.⁸ This has the advantage that we do not have to buy into Oldenquist's highly

⁸ Dreier (2011, pp. 99–100) opines that whereas the distinction between the causal and constitutive understanding is of theoretical interest, it has not been important in the history of ethics.

artificial sequential model of actions that construes lying as the consequence of my uttering some words.

The notion is also not radically new. As Dreier (2011, p. 100) points out, it has a precursor in G. E. Moore:

One natural way, and perhaps the most natural way, of understanding the expression “the total consequences of the action, A”, is one in which among the consequences of A nothing is included but what is the case *subsequently* to the occurrence of A, so that the “total consequences of A” means everything which is the case *subsequently* to A’s occurrence, which is also such that it would not have been the case if A had not occurred. (Moore, 1942, p. 559)

Taking up the temporal aspect of Moore’s definition, Shaver (2011, p. 132) dubs it the *backward-looking* account of consequences, going so far as to claim that it is actually the most common understanding of consequences. Broome (1991, pp. 3–4) agrees, arguing that whenever an act is committed, one of the consequences is that this act has been committed and this entails that the badness of, for example, the breaking of a promise resurfaces as a bad feature of the consequences of said act.

The first move of consequentializing thus seems to have some intuitive appeal as well as some historical pedigree. However, it is not enough to get consequentializing off and running.

Move II: An Agent-Relative Conception of Goodness

By incorporating the morally relevant aspects of non-consequentialist theories into the consequences, consequentializers go some way toward getting the same conclusions as non-consequentialists. For example, they can now value the keeping of a promise in itself and not only with regard to its contribution to a state of affairs that is otherwise valuable. But merely reframing non-consequentialist features as (values of) consequences does not necessarily yield the correct verdicts. Consequentializers also need to weigh the (values of) consequences according to a tailor-made theory of the good, so as to yield the same deontic verdicts for a consequentialist counterpart theory. This involves delicate fine-tuning. For example, if a non-consequentialist theory includes an absolute ban on, for example, murder, we need to stipulate that the negative value of a possible outcome where a murder has been committed is infinite.

This kind of tweaking and fine-tuning of one’s theory of the good is probably not an idiosyncratic move on the part of consequentializers but

rather the daily business of any consequentialist.⁹ There is another problem, however, which prompts the second move. The problem is best brought out by an example. Consider a person P who has promised to conduct act A. Consequentializers, as just outlined, can make room for this by conceiving of the keeping of the promise as a valuable constitutive part of the consequences. However, what if by not keeping her promise, P would allow another person R to keep two of his own promises? Assuming that the keeping of all promises in this scenario is of approximately equal value, P ought to break her promise according to the consequentializing theory. For the value of two promises outweighs the value of P's one promise. Yet many non-consequentialist theories would not allow this. They require P to keep her promise, without regard to whether R can keep his.

One might think that this can easily be fixed by adjusting the weight of the value of the original promise. If P's promise should not be broken, we have to simply ascribe it a value that is higher than those of (all of) R's promises. However, the problem with this solution is quickly revealed if we consider what the new theory asks of R. If the value of P's promise is higher than that of (all of) his own, R should, in a similar situation, break his promises if that would allow P to keep hers. Yet this is (again) not what many non-consequentialist theories demand. Instead, they require that both P and R keep their own promises, irrespective of the realization of the other person's promise.¹⁰ The theory would have to specify both that the worth of P's promise-keeping is higher than R's promise-keeping *and* that R's promise-keeping is more valuable than P's. This obviously poses a problem for theories that make the permissibility of acts dependent on which has the better outcome.

The culprit is quickly identified: It is the *agent-relativity* of some elements of non-consequentialist theories. Typically, duties or prohibitions express a moral claim to an agent that is relative to her in the sense that she herself is to make sure that she refrains from, or follow through, a specific action without considering the consequences for the general prevention, or realization, of these acts by others. Thus, I am to refrain from breaking a promise even when, through this, I would prevent others from conducting more such acts.¹¹

⁹ For this reason, I do not classify the weighing itself as a separate step, as do Betzler and Schroth (2019, p. 5).

¹⁰ For more detailed discussions of the problem, see Emet (2010, pp. 2–3), Hurley (2013, pp. 127–131), and Betzler and Schroth (2019, pp. 5–6).

¹¹ Formally, agent-relative elements can be defined as those that are not explicable without a pronominal back-reference to the agent. Compare, for example, Pettit (1997, pp. 124 ff.).

The problem is that consequentialist theories are traditionally conceived of as being *agent-neutral*. They judge the rightness of an action by its contribution to a state of affairs that is evaluated without referring to the role that a specific agent played in realizing it. If I can make sure that other people are in a better situation to keep their promises in the future by breaking a promise of my own, then I am required to do so. More precisely, the problem lies with the concept of good that is prevalent in most versions of consequentialism, namely an agent-neutral concept of the good. The notion of goodness is neutral in the sense that if a state of affairs is good, it is so simpliciter and not with regard to agents realizing it. If keeping promises is valuable, then keeping more promises is more valuable, irrespective of who keeps them. All states of affairs that result from moral actions are ordered according to one scale of goodness that pertains to all persons. What is the right thing for me to do and what is the right thing for you to do are both functions of a single ordering of outcomes. Even if consequentializers can thus extend the good to include something like the keeping of promises, they fall short of copying all the non-consequentialists' intuition as long as they are not able to distinguish between who commits those acts.

What can be done about this? The solution is near to hand. If the problem lies with the agent-neutral theory of the good, we have to modify that theory.¹² Specifically, it means that consequentializers have to adopt a relativized notion of the good. Since they cannot claim that both P's promises are more valuable than R's and R's are more valuable than P's *full stop*, they need to claim that the situation in which P keeps her promise is *better-relative-to-her*, while the situation in which R keeps his promise is *better-relative-to-him*. Consequentializers thus do not have to claim that P's promise both outranks R's and is outranked by it on the same scale. Instead, there are two scales, one ascribing a higher value to P's promise, the other to R's. Importantly, being *good-relative-to-x* is not the same as being *good-for-x*. Keeping her promises must not be good for P in the sense that it contributes to her own welfare or furthers other of her goals. Instead, the state in which she has kept a promise is good from her standpoint in a way in which it is not special from someone else's standpoint.

The notion of such relativized goodness is by no means uncontroversial.¹³ However, just as the first move, it is also not radically new. Sen (1982, pp. 19 ff.) makes use of this idea when arguing for his goal rights systems, that

¹² At least this is the conclusion that consequentializers have drawn. Setiya (2018) provides the outlines of another solution, but see also Howard (2021).

¹³ See especially Schroeder (2007) for a critique.

is, systems which incorporate the fulfillment of rights in the evaluation of states of affairs. Consequentialist theories traditionally suffer from an inability to accurately account for rights. Just as in the discussed example of promises, rights seem to be in tension with an agent-neutral conception of goodness. However, Sen argues that consequentialist theories can account not only for agent-neutral values but also for agent-relative ones. Broome (1991, pp. 3–16) picks up the same point when he holds that it is a mistake to tie teleology down to agent-neutrality. Instead, he defines teleology as any theory that gives us instructions that can be characterized by a better-than relation (together with certain formal constraints). Since many agent-relative theories can be characterized in such a way, they should be considered teleological as well.

In sum, consequentializing, at its core, employs two main moves, adopting a constitutive understanding of consequences and applying an agent-relative notion of the good. Both moves, as we have seen, are not new. We might thus speak of authors such as Broome, Oldenquist, and Sen as *proto-consequentializers*.¹⁴ Moreover, the two steps do not yet suffice to consequentialize *any* plausible theory. In order to account for specific features of specific non-consequentialist theories, consequentializers have to introduce ever greater technical sophistication. Examples include the accommodation of moral dilemmas and of supererogation, as well as the move to time-relativity, which is needed to account for the fact that deontological theories typically don't allow an agent to break a promise in order to minimize future promise-breaking by themselves. Different proposals have been made for how to deal with these additional non-consequentialist features.¹⁵ As of now, it remains to be seen which solutions fit which problems best and whether all technical hurdles can be overcome in the end.¹⁶ However, this would take us too far afield. As I explained in the Introduction, when investigating the phenomenon of moral underdetermination, I have to rely to some degree on the preliminary results of normative ethical theorizing. This is in principle no different from the philosopher of science who, to a certain

¹⁴ Dreier (1993, p. 23) accordingly does not take full credit for consequentializing but credits Broome (1991), Sen (1983), and Vallentyne (1988). The difference he identifies between Broome's treatment and his own is in the scope of application only. In the same vein, Peterson (2013, pp. 168–169) classifies Sen as a *restricted* consequentializer.

¹⁵ For a discussion of moral dilemmas, see Portmore (2011, pp. 88–89), Peterson (2013, pp. 173–174), and Dreier (2011, p. 105); for supererogation, see Portmore (2007, p. 54) and Peterson (2013, p. 174); and for time-relativity, see Louise (2004, p. 534).

¹⁶ For skeptical views, see Brown (2011), Dietrich and List (2017), and Howard (2021).

degree, has to rely on the (preliminary) results of the special sciences. When it comes to consequentializing, the last word has yet to be uttered. However, it is safe to say that thanks to consequentializers (and their opponents), we have a much better understanding now of what technical conditions would have to be met for consequentialism to be able to copy the verdicts that its rival traditions yield.

4.2 Deontologizing: Leveling the Playing Field

Consequentializers' promise sounds great. If successful, it would ensure consequentialism's immunity to any kind of extensional challenge: Whatever ultimately turn out to be the right or wrong acts, consequentialism can account for them. It should come as no surprise, then, that it has been asked whether something comparable can be achieved for non-consequentialist theories, too. If we can construe consequentialist counterparts to any non-consequentialist theory, may not the converse be true as well? This is indeed what some philosophers have proposed in the aftermath of the first attempts to consequentialize. Their proposal is to *deontologize* non-deontological theories.¹⁷

Two Moves, Yet Again

In contrast to consequentializing, deontologizing has been conceived not as an integrated project but more as a direct reaction to consequentializing. Possibly as an effect of this, deontologizing has received much less attention and is characterized by much less of the technical sophistication typical of consequentializing. As a consequence, there is also less consensus regarding the actual implementation. Nevertheless, I propose that, from what has been said, we can come up with a similar two-step procedure.

¹⁷ Since most of the literature is concerned with deontologizing consequentialist theories, I will sometimes speak as if this were the whole purpose of deontologizing. However, it is more accurate to say that just as consequentializers aim to consequentialize any non-consequentialist theory, deontologizers aim to deontologize any non-deontological theory. Let me further mention that Suikkanen (2021) has recently argued that we can also *contractualize*. I will omit discussion of his proposal, not because I take it to be less plausible than consequentializing and deontologizing but simply because it would considerably complicate the exposition of my overall argument. For the record, I do think that most of the general points I make about consequentializing and deontologizing also apply to contractualizing, and the overall argument could thus be generalized to include the latter as well.

Move I: Switching from Consequences to Acts

Consequentializers, when trying to copy the deontic content of a non-consequentialist theory, have to find a way to build the features that the latter theory considers to be morally relevant into the consequences. Deontologizers have to do the opposite. Instead of building features of acts into consequences, they have to find a way to somehow copy the evaluative ranking of outcomes without referring to these outcomes. We can put this idea into a simple form *À la* Dreier's strategy:

Move I: For every outcome of an act that a consequentialist theory prefers, insist that the act leading to that outcome is to be done, its alternatives not to be done, independently of the outcome.

How are deontologizers supposed to achieve this? Two proposals can be distinguished. The first one comes from Hurley. The crucial point, as Hurley (2013, pp. 141–149) sees it, is that consequentialism is not the only theory that employs an evaluative framework. Instead, we can also evaluate actions directly. As an example, Hurley points to Herman, who thinks that:

[...] [T]he chief source of misunderstanding of Kantian ethics is the almost universal commitment to treating it as a moral theory without a notion of value as its fundamental theoretical concept. (Herman, 1993, p. 208)

Against this tradition, Herman argues that value does indeed play the central role in Kant's moral philosophy. Instead of simply proscribing duties that tell us what to do, actions do have to be evaluated as better or worse. It is just not necessary that this evaluation mentions the action's consequences. Instead, for Herman, one can evaluate actions directly. Hurley (2013, pp. 141 ff.) proposes that deontologizers make use of this idea. In that way, they can simply copy the ranking of outcomes that the target consequentialist theory yields and recreate it for the non-consequentialist alternative, but this time as a ranking of acts. In this way, we arrive at a deontically equivalent theory that does not derive the deontic status from the consequences and is thus not consequentialist.¹⁸

How promising is this strategy? I think that it is fine as far as it goes. Hurley's goal is to set out an alternative framework that could do to consequentialism what consequentializing has done to non-consequentialism. Consequentialism has at its core the claim that the rightness of acts is a function of their contribution to outcomes. Hurley's alternative states that it is the goodness of an act itself that makes it right or wrong. We thus

¹⁸ Compare also Hurley (2017) for more on the idea of telic evaluations of actions.

arrive at deontically equivalent alternatives that at the same time include incompatible explanatory claims as to what makes acts right or wrong. Still, Hurley's proposal will go against many philosophers' initial expectation of what a procedure called *deontologizing* is supposed to look like. The reason is that deontologist theories are not classically affiliated with the use of an evaluative framework. Indeed, many non-consequentialists explicitly deny that the deontic status of acts should be (only) a function of an evaluative ranking. Instead, the deontic status of acts is given by duties or rules making no mention of what would be better or worse.¹⁹

I therefore think that a second proposal, hinted at by Sachs (2010), is more promising. When considering how deontologists could copy the consequentializing strategy, Sachs suggests the following:

There will usually be a number of ways to do this, but at least one method will always be available: the construction of a ten commandments-style list of DOs and DON'Ts, one for every member of [the class of case-specific verdicts]. For instance, if "it is impermissible to kill Jones" is a member of [the class of case-specific verdicts], just put "DON'T kill Jones" on the list of principles in [the non-consequentialist theory]. (Sachs, 2010, p. 265)

Like Hurley, Sachs proposes that in order to deontologize, we should focus on acts instead of consequences. However, he does not introduce an additional evaluative framework that ranks these acts. Indeed, his commandment-style theories make no use at all of a connection between evaluative and normative facts. Instead, they consist of primitive directives to do some acts and shun others. The idea is very simple. Take the set of deontic verdicts that a consequentialist theory yields for every specific act and insist that the deontic status is no function of the outcomes of the acts but is given by equally many primitive directives of a DO and DON'T character. This yields a list that is extensionally equivalent to the original consequentialist theory. What is more, while Hurley's theory will probably look too idiosyncratic to most non-consequentialists, commandment-style theories are arguably the prototype of deontological theories.²⁰

¹⁹ Hurley (2013, p. 138) is of course aware of this. Indeed, as Hurley (2017, pp. 50–51) points out himself elsewhere, there is a chapter by Herman titled "Leaving Deontology Behind." In it, Herman (1993, p. 210) proclaims: "[i]n this sense, Kantian ethics is not a deontology." Indeed, Herman (1993, pp. 210 ff.) mounts a comprehensive challenge to *any* moral theory that does not have a conception of value at its core.

²⁰ To be clear: It would be strange to outright reject Hurley's proposal on the basis that it employs a non-standard version of non-consequentialism, considering that consequentializers themselves employ a non-standard version of consequentialism. However, if we can come up with a strategy that involves a more classical form of non-consequentialism, this is an advantage.

Sachs's proposal has a problem of its own, however. Sachs likens the list to the ten commandments. Yet the sheer number of directives that would be needed also evokes a different comparison. Kagan (1989, p. 11), discussing a similar procedure, dismissively speaks of a *shopping list*. The derogative moniker points to the main problem of the proposal: Such a list is barely recognizable as a moral theory. Moral principles are not case-specific verdicts, nor do they involve rigid denominators, such as "Jones" in Sachs's example.²¹ Even though principles can be very specific, they need to be applicable in more than one situation to qualify as such. Kagan (1989, p. 11) points toward a technical solution to this problem. Instead of directly referencing particular persons or situations, we can instead use formulations such as "in situations of type 1 do A, of type 2 do B, of type 3 do C." We then describe those situations in such detail that they only fit one particular situation, but without using rigid denominators.

Although maximally specific, these formulations do, strictly speaking, qualify as principles. However, the underlying problem remains. Moral theories usually just don't consist of principles of such specificity. One way to go forward from here is by bundling together some of the most specific directives. Let us assume that there are some types of actions that the consequentialist theory that needs to be copied never requires. An example might be the torturing of innocents for the sake of torture itself. Since the consequentialist theory evaluates all instances of this type of action the same way, there is no need to include specific directives for every particular case in the deontologist counterpart. Instead, deontologizers can pool all the specific directives and substitute the principle "Thou shalt not torture for the sake of torture" for them. We thus arrive at principles of a higher level of generality. The limits of how far we can take these improvements should be obvious, though. "Thou shalt not torture people for the sake of torture" is still a very specific principle that only bears upon a very specific and also quite rare type of action. The pooling move will not work, however, for most types of action. Take the more general act-type of torturing. Although it is one of the most despicable types of actions, some versions of consequentialist theories will allow for instances of it as long as the alternatives are sufficiently bad. This is even more likely for less despicable types of action such as lying or breaking a promise. It is thus a

²¹ As Louise (2004, pp. 521–522) points out with regard to consequentializing, the pronominal back-reference to the agent that agent-relativity requires is not the same as referring to the person by using a *rigid denominator*. There is a moral duty that I look out for my own children, but that moral duty does not refer to me, MB. Instead, the moral duty applies to anyone who occupies my place and stands in the same relation to some person or act as I do.

safe prognosis that the proposed bundling will not take us very far. Instead, deontologizers need to bring in heavier machinery.

Importantly, the history of deontological theorizing itself arguably provides precisely that. One way it does so is hinted at by Portmore (2007, p. 40). He observes that non-consequentialists might sometimes permit actions they usually forbid. They do so when they permit an act that saves five people but kills one as a foreseen but unintended consequence. The idea Portmore alludes to here is what is called the *doctrine of double effect* (DDE). The doctrine states, very roughly, that there is a morally relevant difference between the intended effects of an action and those we merely foresee but do not intend.²² An action, such as a killing, which would be forbidden if intended, can be allowed if it is merely foreseen but neither intended as an end nor as a means.²³ How does the DDE help deontologizers? In the following way. On the one hand, they can uphold general prohibitions on lying or killing, subsuming all actions of lying or killing that the consequentialist theory forbids under these principles. On the other hand, in those cases where the consequentialist theory requires us to lie or kill, deontologizers need to assume that these actions are only an unintended side-effect of a more general action. Say the consequentialist theory requires us to kill one person in order to save five. In order to uphold the prohibition on killing, deontologizers cannot simply allow this. But they can reframe the action as one of saving five people, while accepting, though not actively intending, one person as collateral. This allows deontologizers to add general principles like a prohibition on killing to their theory while at the same time making room for exceptions when the original consequentialist theory requires this.²⁴

The other option consists in the introduction of what Ross (1930) calls *prima facie* duties. *Prima facie* duties, roughly, are duties that would prescribe an action, and thereby become a *duty proper*, as long as they are

²² Compare Woodward (2001, pp. 1 ff.). For a more detailed discussion of the complications of defining the DDE, see Boyle Jr. (1980).

²³ It is important to note that the DDE it is not just an exculpatory principle but a justificatory one. It does not claim that killing as an unintended side-effect is still a wrong act but one we can exculpate someone from. Rather, it is the justified and therefore right action. Only if we understand the principle this way does it help us in construing deontically equivalent counterparts to consequentialist theories, which do justify (and not just excuse) killing in the relevant situations. Compare Boyle Jr. (1980, pp. 529 ff.).

²⁴ It is no coincidence that the DDE would lend itself to deontologizers' purposes. As Woodward (2001, p. 2) puts it: "[t]here are many versions of the DDE, but perhaps most clearly it is interpreted as limiting the application of certain prohibitions." One way we can understand the problem that deontologizers face when trying to cope with the target theory's verdicts is that their theory includes proscriptions that are too rigid. Any means that allows them to make exceptions is therefore welcome, and that is exactly what the DDE renders possible.

the only morally relevant consideration that is applicable to a specific case.²⁵ If they are not, which Ross thinks is almost always the case, one *prima facie* duty must be weighed against another. Ultimately, one outweighs the other(s), thereby becoming the *duty proper*. However, and this is where it gets interesting for deontologizers, *prima facie* duties do not cease to exist when they are outweighed by another *prima facie* duty. This fact is crucial to deontologizers. It allows them to include in their theory general principles, like prohibitions on lying, even though the original consequentialist theory, which they are trying to copy, sometimes requires people to lie. Deontologizers simply need to claim that other *prima facie* duties outweigh these *prima facie* principles when the consequentialist theory yields the respective results. What if a consequentialist theory requires us to keep our promise in one case but to break it in another? No problem. Deontologizers must merely posit that the *prima facie* duty to keep one's promise outweighs all other duties in the former case but not in the latter. Granted, this opens up some rather complicated issues regarding the (relative) weighing of different *prima facie* duties. Still, the concept of *prima facie* duties offers at least some more wiggle room for deontologizers, while having a quite impeccable reputation as a non-consequentialist means.²⁶

These remarks are very off-hand. However, they should give an impression of how the history of deontological theorizing itself holds some promising tools for deontologizers. Still, there is a second move required.

Move II: Adding an Explanatory Principle

If the aforementioned tools are working, we arrive at principles that are of a more general kind than the primitive shopping list that Kagan comments on. However, do they count as theories yet? At least according to the model proposed in Chapter 2, they do not. On that model, moral theories should at least fulfill two functions: an extensional and an explanatory one. The resulting list of principles does encode the extensional part, since acts are

²⁵ As I mentioned above, some authors prefer the term *pro tanto* to avoid some unhappy epistemic connotations, but I will stick with Ross's original label to highlight the historical continuity of the idea.

²⁶ Readers might be confused as to how the inclusion of *prima facie* duties fits with my previous claim, in Chapter 2, that particular moral verdicts are primary due to their determinateness. For Ross thought that while we do know what *prima facie* duties there are, we might actually never know our duties proper. This view is indeed in conflict with my claim about determinateness. However, I do not suggest that deontologizers accept Ross's full view, including his moral epistemology. Ross's skepticism regarding our knowledge of duties proper does not follow directly from his account of *prima facie* duties. Instead, we can combine the notion of a *prima facie* duty with the idea that we often know our duties proper. Indeed, I think that this is the more common way in which ethicists include the notion of *prima facie* duties in their theorizing.

right or wrong if and only if they feature in one of the principles. What is missing, however, is an explanatory relation. Moral theories, on the model we are assuming, entail claims as to what makes acts right or wrong. If deontologizers intend to come up with counterpart *theories* (not just lists), they therefore need to introduce such an explanatory relation with their second move.

Here, Portmore (2007) and his idea of *Kantianizing* suggests a way forward:

To illustrate, suppose that a Kantian takes the principle “an act is morally permissible if and only if it involves treating humanity as an end-in-itself and never simply as a means” to be the one and only fundamental moral principle. The recipe for Kantianizing is, then, as follows. Take whatever considerations [the non-Kantian theory] regards as relevant to determining the deontic status of an action and insist that those considerations are relevant to determining whether humanity has been treated as an end-in-itself. (Portmore, 2007, p. 59)

How would this work? As an example, Portmore considers Traditional Act-Utilitarianism (TAU):

The theorist who wants to Kantianize TAU need only insist that we treat humanity as an end-in-itself if and only if we give equal consideration to everyone’s interests in maximizing aggregate utility. (Portmore, 2007, pp. 59–60)

Portmore’s proposed move looks very straightforward. Moreover, it is also easily adaptable. Portmore (2007, p. 60) suggests that a large number of fundamental moral principles could figure in place of the Kantian principle. We could, for example, *contractualize* a theory by claiming that whatever considerations it claims to be relevant are only relevant insofar as they are the only ones that could figure in principles that “no one could reasonably reject as a basis for informed, unforced general agreement.”²⁷ Generalizing from these examples, the second move thus looks something like this:

Move II: Claim that what makes an act right or wrong is not its contribution to the value of a state of affairs but that it constitutes [some other explanatory principle] being fulfilled.

Portmore (2007, p. 60) goes so far as to say that this could be done using almost any explanatory principle. That seems overly optimistic. Consider a principle that specifies that acts are right because they benefit God’s

²⁷ The formulation is from Scanlon (1998, p. 153).

chosen people. Let us also stipulate that the theory clearly specifies who these chosen people are. Such a theory is very restricted when it comes to the range of verdicts it can yield. For instance, it cannot arrive at the same verdicts as standard utilitarianism since that theory would sometimes require one of the chosen people to suffer for the greater benefit of the not chosen. Portmore's strategy thus only works for theories that contain a fundamental explanatory principle that is highly unspecific when it comes to substantive moral propositions.²⁸ Still, there is no reason to assume that the move will not work employing at least some principles, although which ones these are might have to be decided on a case-by-case basis.

What we thus get is a much more piecemeal strategy than consequentializing. In some sense, it calls for every non-consequentialist to go at it alone. Kantians Kantianize, contractualists contractualize, and so on. However, this does not strike me as too big of a problem. There is no obvious reason why the deontologizing strategy should not allow for some variability. Consequentializing also involves some variability when it comes to the treatment of some features, such as supererogation or moral dilemmas, which have to be accounted for. Similarly, deontologizers need not commit themselves to which explanatory principle to use. Non-consequentialist theories provide a wide array of explanatory principles. In fact, it might be an advantage of the proposed strategy that it allows for different explanatory principles to be plugged in. It might well be the case that different principles fit better with some sets of verdicts, but not others. Deontologizers thus need not insist on a one-size-fits-all strategy, as long as there is a common framework on which improvements can be made.

This framework, or at least the rough outlines of it, is given by the two moves. Taking them together, we can state the recipe for deontologizing as follows:

For every outcome of an act that a consequentialist theory prefers, insist that the act leading to that outcome is to be done, its alternatives not to be done, independently of the outcome. Then claim that what makes an act right or wrong is not its contribution to the value of a state of affairs but that it constitutes [some other explanatory principle] being fulfilled.²⁹

²⁸ Kant's and Scanlon's principles are cases in point: It is not immediately clear which acts are compatible with treating humanity as an end in itself or which could not reasonably be rejected. Rather there is arguably quite some leeway when interpreting what these propositions entail regarding specific actions.

²⁹ I am deliberately trying to imitate Dreier's recipe-style remarks from the original formulation of the consequentializing strategy here. If people find this *gimmicky*, I ask them to postpone their worries until Chapter 5.

The theories that are produced by this strategy qualify as moral theories under our definition of a moral theory. They outline which acts are right and wrong and they put forward an explanatory claim as to what makes these acts right or wrong.

The result of the strategy, if it proves successful, can next be expressed by copying consequentializers' Deontic Equivalence Thesis:

(DET*): [...] [F]or 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.

Granted, this has been very quick. Before we can say with any certainty that deontologizing succeeds, several issues would have to be investigated in more detail. First and foremost, it is not clear that the proposed strategy allows us to arrive at anything resembling enough of a real moral theory. The original list theory, amended with an explanatory principle, only formally qualifies. The additional tools, such as the DDE or the notion of *prima facie* duties, promise some progress in this regard.³⁰ But it remains to be seen how the end products of the deontologizing strategy will look. Personally, I am quite positive about the prospects going forward.³¹ However, we need to at least be open to the possibility that deontologizing might end up with something that is much less systematic than traditional theories. If so, we would have to delve deeper into what the identity criteria are for moral theories. People might argue that systematicity is not a necessary criterion and that even primitive list theories qualify as acceptable.³² Nevertheless, such theories remain unattractive. We expect from a moral theory that it reveals connections between the verdicts it yields. Consequentialist theories do so by means of their theory of the good. Deontologists lack this means. Instead, they can do so by including principles about general action types, such as stealing or promising. It remains to be seen whether the proposed additional tools will yield such principles for deontologizers.

³⁰ Again, readers might be confused how the notion of *prima facie* duties fits in since Ross was also skeptical about fundamental explanatory principles. Yet, once more, I will insist that we don't have to buy Ross's whole view in order to make use of his notion of *prima facie* duties. Compare, for example, Hooker (2000) for the view that we can combine Rossian mid-level duties with a rule-consequentialist explanatory claim.

³¹ For reasons that will be outlined in more detail in Chapter 5. Basically, the idea is that merely making room for the deontic verdicts of one theory in the framework of another is not that difficult. The thorny issues concern the question of how plausible the copying theories turn out to be.

³² I have had several people make this argument in conversations.

This, once more, leaves us in a situation where the success of the overall argument is dependent on contentious issues in normative ethics. However, as I have pointed out repeatedly, this is not in principle different from how the arguments of philosophers of science are dependent on contentious issues in the special sciences. We might still be able to draw interesting conclusions, even if they remain, to some degree, conditional.

4.3 Algorithms and Underdetermination

Let us thus assume, for now, that consequentializing and deontologizing succeed. How should we interpret these results, and what do they entail for our understanding of the status of normative theories and the relation between them? My view, of course, is that the two projects lead to a form of underdetermination. However, there are other views that need to be considered as well. I will start with two prominent interpretations of consequentializing which, although located at opposite ends of the spectrum, suffer from a common, yet instructive, shortcoming. Since the underdetermination interpretation does not share that flaw, we have strong reason to prefer it. However, there is a third interpretation that also does better than the others. I will discuss this version together with the underdetermination interpretation, offering a tentative argument to the effect that, pending further reasons, we should prefer the underdetermination interpretation. However, since a thorough discussion of these further reasons involves complicated matters in metaethics, the rival interpretation will remain with us until Part III. Finally, I will consider deontologizing. Fortunately, this will take up less space since deontologizing was mainly conceived as a counter move to consequentializing, and the interpretations do not vary much. In fact, I will argue that where there is a relevant difference, this difference actually strengthens the case for the underdetermination interpretation.³³

All or Nothing

Perhaps the best-known interpretation of consequentializing is that by showing how one can consequentialize a theory, one has actually proven

³³ There are yet other interpretations, most prominently the one offered by Douglas Portmore. However, since that interpretation involves issues of *ampliative underdetermination* that I can only take up in Chapter 5, discussion of this interpretation will have to wait. For a systematic overview of the different interpretations, see Schroeder (2017) and Portmore (2022).

the theory itself to be a variant of consequentialism. Timmons succinctly summarizes this interpretation:

If consequentializing succeeds, then what has [sic] traditionally been taken to be debates in moral theory between consequentialists and non-consequentialists are really best understood as debates within the consequentialist framework. (Timmons, 2013, p. 6)

Louise (2004, p. 536) similarly concludes that:

[...] we are now all under the consequentialist umbrella [...].³⁴

Apart from being the most notorious interpretation, this is also the one that is most strongly suggested by the name of the procedure itself. The talk of “consequentializing theories” seems to imply that something is being done to non-consequentialist theories.

However, it is not immediately obvious how this should work. Why, one can ask, should producing a deontically equivalent counterpart theory change something about the status of the *original* theory itself? Brown has come up with a nice analogy to illustrate this concern:

Suppose I employ you to paint my red house blue, and you then proceed to build another house which is a perfect replica of mine except it's painted blue rather than red. Then you haven't done your job. The new house may be blue, but it's not my house. You haven't done anything to my house; it's still red. (Brown, 2011, p. 755)

The point Brown is making here is that, considering that the original theory remains unchanged by the consequentializing procedure, adherents of non-consequentialist theories can at least ask for some further reason before accepting their sudden conversion to consequentialism. Perhaps surprisingly, considering how prominent this interpretation has been in the wider philosophical discussion, its proponents provide very little in the way of a complete argument. Instead, the claim that all theories turn out consequentialist tends to be regarded as a description of the results of consequentializing, rather than an additional interpretation.

One could, of course, simply stipulate that any theory that produces the same verdicts as a consequentialist theory is by definition also consequentialist. However, that would just beg the question against non-consequentialism and is probably uncharitable to proponents of the first

³⁴ The picture of the umbrella is foreshadowed in Broome (1991, p. 4). Critics, by contrast, have sometimes preferred the metaphor of a consequentialist vacuum that sucks up all other theories (compare McNaughton and Rawling (1991, p. 168) and Hurley (2013, p. 123)).

interpretation. We can come up with a better reasoning. One clue is given by Peterson (2013, p.170), who suggests that we think of a moral theory as: “[...] a device that attaches deontic predicates to acts.” Understood as such a device, moral theories only have an extensional role: Their sole function is to produce the correct deontic verdicts. A theory is thus consequentialist if and only if (and not because!) it maximizes the good. Still, why would that make the original theory consequentialist? The second clue is given by Brown (2011, pp. 755–757). Proponents of the first interpretation, he thinks, also need to argue that *the good*, as used by consequentialists, is a technical term, introduced to pick out just those properties that feature in their theory of the good.³⁵ This allows them to argue that since the original non-consequentialist theory and its counterpart theory do by hypothesis pick out the same things as right, and those things that are right are also those that maximize the good (according to consequentialism), the non-consequentialist theory also maximizes the good. It is thereby proven to be consequentialist, too.

Brown himself remains unconvinced. Instead of turning us all into consequentialists, he thinks that the consequentializing strategy, if successful, would backfire:

But then, it seems, consequentialism would be empty – trivial, vacuous, without substantive content, a mere tautology. (Brown, 2011, p. 750)

This is the second interpretation. Brown holds that if consequentialism can indeed account for all its own verdicts *and* for those of all alternative theories, it would thereby become too wide. Its malleability would turn out to be to its own detriment. It would no longer provide a distinct perspective and would cease to guide us. To be sure, this only applies to consequentialism as a family of theories and not to the specific consequentialist theories that are construed as counterparts in the consequentializing procedure. For remember that the procedure aims to come up with different, but equally many, consequentialist theories for every non-consequentialist theory. These specific consequentialist theories would still provide us with specific sets of verdicts and would thus escape Brown’s objection. Brown (2011, p. 754) accordingly carefully distinguishes between the family of consequentialism as the *genus* in contrast to the *species* of particular theories. Still, consequentialism as a family of theories would, according to the second interpretation, become trivial and empty.

³⁵ Brown (2011, pp. 755–757) calls this the *semantic thesis*.

So, are we all consequentialists now, or has consequentialism become trivial and empty? All or nothing? I think that neither of the two is a plausible interpretation. They share a common flaw, and one which should not come as too much of a surprise by now. The flaw consists in an impoverished understanding of moral theories and their functions. Both interpretations exclude or downplay any functions of theories that go beyond the mere attaching of deontic predicates to acts, most importantly, explanation.

This point has not been lost in the debate, of course. Portmore has highlighted it in an especially succinct way:

[...] [E]ven if two theories agree as to *which acts are right*, that does not mean that they agree on *what makes those acts right*. It is important to note, then, that moral theories are in the business of specifying what *makes* acts right. And so even two moral theories that are extensionally equivalent in their deontic verdicts can constitute distinct moral theories – that is, distinct theories about what *makes* acts right. (Portmore, 2011, p. 109)³⁶

Portmore's criticism, although only directed at the view that consequentializing leads to theories that are no longer distinct, applies to both interpretations. We are not all consequentialists because some of us offer distinctly non-consequentialist explanations. But nor does the consequentialist tradition turn out to be trivial, simply because it can lead to any verdict. The way it explains these verdicts is surely not tautological and clearly distinguishes it from rival theoretical frameworks. If the second tenet of the Textbook View is accurate, consequentialism as a family of theories retains a distinct explanatory character and therefore is not turned into an empty theory. It might be deontically neutral, but it is not so when it comes to explanation.

To be clear, thinking of moral theories as mere mechanisms to produce deontic verdicts is a perfectly valid thing to do, *if all one intends* is to find out whether those theories lead to the same verdicts.³⁷ Too detailed an account of theories would maybe do more harm than good when trying to answer this question. However, we cannot appeal to the same impoverished view to warrant sweeping claims about the nature of moral theories. To think of theories as mere verdict-producing mechanisms is a reductionist stance. As long as the theoretical differences between non-consequentialist and consequentialist theories have not been convincingly settled, there is

³⁶ Compare also Portmore (2007, p. 60), Dougherty (2013, p. 534), Woodard (2013, p. 262), and Tenenbaum (2014, p. 234). The point seems also to have been clear to some of the *proto*-consequentializers, such as Broome (1991, pp. 8–9).

³⁷ To be fair, Brown (2011, p. 753) explicitly adopts such a view in order to make as charitable a case for consequentializers as possible. However, this does not change the fact that as a general description, his interpretation is untenable.

no reason for non-consequentialists to give up their own theory *on the sole reason* that it leads to the same verdicts as some other theory. Pending further arguments to the contrary, non-consequentialists are perfectly right to reject the reductionist stance and insist that their theory is, after all, substantially different from consequentialism.

The same points hold, *mutatis mutandis*, for deontologizing. Being able to account for the verdicts that consequentialist theories yield neither means that we are all deontologists, nor does it render that tradition empty. Indeed, the possibility of deontologizing provides an additional argument against this line of thinking. As Portmore (2007, p. 59) observes, the first interpretation becomes even less plausible if we take deontologizing into account. For if consequentializing establishes that we are all consequentialists, by the same token, other strategies would establish that we are all Kantians, or contractualists, and so on. Moreover, we would be all of these *at the same time*. Portmore (2007, p. 59) finds this patently absurd because we are talking about rival theories that accept different fundamental moral principles. We cannot be consequentialists and deontologists at the same time, since we would then have to accept incompatible rationales for what makes acts right. Similarly, if being able to account for one's opponents' verdicts renders a family of theories empty, then, if deontologizing succeeds, deontology is an empty tradition as well. That would make the two most prominent traditions of moral theorizing empty, prompting the question of whether there are any non-vacuous theories left in ethics.

To sum up, consequentializing and deontologizing neither show that consequentialism or deontology are the only game in town, nor do they render those traditions empty. Instead, if these procedures succeed, we are presented with the finding that consequentialist and non-consequentialist theories can be extensionally equivalent while at the same time maintaining their theoretical differences.

The Underdetermination Interpretation and Its Main Rival

This last formulation, of course, is equivalent to claiming that consequentializing and deontologizing lead to underdetermination. Just as in the case of Parfit's project, Dietrich and List (2017) are the first to point out this analogy. Commenting on Dreier's and Portmore's observations about extensionally equivalent theories, they concur that:

It is entirely possible [...] that different normative theories entail the same permissibility verdicts, despite arriving at them in different ways. (Dietrich and List, 2017, p. 424)

However, they go on to comment that:

If we are interested not only in how we ought to act but also in why we ought to act in that way, then we cannot generally consider two extensionally equivalent theories as equivalent *simpliciter*. (Dietrich and List, 2017, p. 425)

This situation, they observe:

[...] is analogous to the case of science, where two or more distinct theories may explain the same observations and thus be observationally equivalent, despite being explanatorily different: W. V. Quine's (1975) famous *empirical underdetermination problem*. (Dietrich and List, 2017, p. 425)

Besides making explicit the analogy to science, Dietrich and List's *reason-based representation* of moral theories sheds light on why the underdetermination interpretation does not suffer from the same flaw as the two we just discussed. Moral theories have a reasons-structure that varies from tradition to tradition (as well as within traditions). This explains why there can be differences between theories even if they are extensionally equivalent. Still, as I observed in Chapter 3, due to their more formal interest, Dietrich and List say relatively little with regard to why the underlying explanations that the reasons structures encode are indeed incompatible.

This issue becomes particularly pertinent when we consider a third interpretation of consequentializing that I have not discussed so far. This is the view argued for by Dreier (2011). Short of claiming that all theories turn out to be consequentialist, Dreier (2011, p. 98, pp. 111–114) instead argues that consequentializing proves the different traditions to be merely notational variants.³⁸ To illustrate, consider the following example from Portmore (2011, pp. 109–110). Two people compare a substance's temperature to that of water boiling at sea level. One person uses Fahrenheit as a scale, the other Celsius. On one person's account, the substance is being measured at a temperature of 100° Celsius; on the other's, it is being measured at 212° Fahrenheit. Are these people disagreeing? Clearly not. Instead, what is at issue is a matter of notation only, since the different expressions pick out the same property. Consequently, we would not want to say that the two people are offering incompatible explanations.³⁹

Dreier (2011, pp. 114–115) argues that the same holds true for moral theories. In support, he offers some interesting considerations that stem

³⁸ The interpretation is foreshadowed by Oldenquist (1966, p. 180), one of the proto-consequentializers.

³⁹ Compare also Schroeder (2017, p. 1480), who calls this the *assimilation argument for consequentializing*.

from a very particular understanding of our concept of goodness or, more precisely, our concept of goodness as it relates to states of affairs, such as in a good outcome. Dreier thinks that when applied to states of affairs, goodness is inextricably connected to the role it plays in guiding choice. The notion of a good outcome, as Dreier puts it, is *moored securely* to what we perceive our proper choices to be. He does not go so far as to claim that we have *no* independent notion of goodness as it relates to states of affairs. Yet he thinks that whatever additional connotations might attach to this notion, they are ultimately too feeble to come apart from action-guidance. Applied to moral theories, this means that when a moral theory says that something is good because it leads to a good outcome, it is not implying anything more substantially than what it conveys through its verdicts about actions anyway. Since consequentializing and consequentialized theories do, per definition, imply the same verdicts about actions and since what the consequentializing theory implies about outcomes has no substantial content *above* what it implies about what the correct actions are, we have to consider the two theories as merely differing in their wording.

Dreier's interpretation, the *notational variants* interpretation, poses a real challenge to the underdetermination interpretation. To see why, recall what I have identified as the main problem of the two other interpretations above. Their common flaw, I have argued, is that they presuppose a reductionist understanding of moral theories – one that excludes the explanatory function – which, *at first glance*, is implausible. Just because theories can agree about what we should do does not make us all consequentialists (or deontologists), and neither does it render those traditions empty. Dreier, it is important to note, presupposes the same reductionist picture. He makes this explicit when he advances the following claim in support of the notational variants interpretation:

Extensionality Thesis (ExT): ExT says that nothing but extension matters in a moral view. (Dreier, 2011, p. 98)

Yet, Dreier does not simply assume ExT. ExT holds, Dreier thinks, because the notion of goodness as it relates to a state of affairs is moored in action-guidance. Dreier thus offers a justification for his reductionist view of moral theories. This justification needs to be considered before we can reject his interpretation. Note further that Dreier need not be concerned about the possibility of deontologizing that sheds additional doubt on the other two interpretations. If we think that consequentializing proves consequentialism and deontology to be merely notationally different, then deontologizing might do the same without this being in any way

contradictory. Aware of this fact, Dreier (2011, p. 115) is happy to state for the record that he has no objection to deontologizing.⁴⁰ The notational variants interpretation hence looks more plausible than the other two. Indeed, I do not, at least at this stage, profess to be able to conclusively counter it. Assessing its (relative) plausibility involves deeper questions about moral semantics and epistemology that I can only take up in Part III.

However, I want to at least make a tentative case that, pending further argument, the underdetermination interpretation should be preferred. To do so, let me start with a simple point. Most people don't buy ExT, *even if* they are willing to accept DET. Instead, Portmore's point from above remains relevant:

As Portmore noted, ethical theories purport to do more than simply identify what acts are right and wrong; they also seek to explain *why* those acts are right and wrong. The Extensionality Thesis therefore seems to be false, because ethical theories that declare the same actions right may nevertheless differ in their explanations of *why* those actions are right. (Schroeder, 2017, p. 1481)

In fact, Dreier (2011, p. 114) partly acknowledges this. He agrees that most people find ExT even less plausible than DET. Yet this is an obvious dialectical problem. ExT is advanced in support of DET. If people find ExT less plausible than DET, it can hardly do that. Moreover, if another interpretation can offer an argument to support DET that is less controversial than Dreier's, this speaks strongly in favor of that alternative interpretation.⁴¹

Can the underdetermination interpretation offer such a superior explanation? One observation might make us suspect that it cannot. That observation is that for two theories to be deontically equivalent, it does not suffice that they merely agree on all past acts in the actual world. Instead, they need to agree about all possible acts. This is captured by the following formulation of the DET:

[...] [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)

⁴⁰ Portmore (2011, p. 110) also thinks that it would not be absurd to think that we are all consequentialists and Kantians and contractualists at the same time if these were merely notational variants of the same view.

⁴¹ As Schroeder (2017, p. 1484) notes, the observations that Dreier offers in support of ExT do not amount to a robust argument. Instead, when advancing his considerations about goodness and action-guidance, Dreier (2011, pp. 111 and 114) only talks about having a *suspicion* and a *hunch*.

As this formulation suggests, consequentializing and consequentialized theories yield the same verdicts in every possible world. In other words, they have to be *necessarily* coextensive in their verdicts. Yet that extends to their explanations as well. For example, if Kantianism and its consequentialist counterpart are deontically equivalent, then it is necessary that, for example, treating someone as an end in itself maximizes outcomes on the consequentialist theory. We treat someone as an end in itself if and only if the outcome of that act is the highest ranked. However, in this case, we might ask where the necessity comes from in such cases. Why is it that treating someone as an end in itself also maximizes outcomes? There is one explanation that seems natural: The two theories might actually, by means of their seemingly different explanations, refer to the same property. If referring to outcomes and referring to means as ends necessarily yield the same verdicts, this might suggest that they are merely different ways of describing one and the same right-maker. And if two theories refer to the same property, albeit through the use of different words, there is arguably no radical difference between them.

This is where our discussion of grounding becomes important. The above reasoning begs the question against the grounding model of moral explanation. A non-consequentialist theory and its consequentializing counterpart do not need to pick out the same grounding relation just because they have the same extension. The grounding relation has a direction – some feature grounds the other – that necessary co-extension does not have. We might treat someone as an end in itself if and only if the outcome of that act is also the highest ranked. But on the Kantian story, an act is right *because* it treats someone as an end in itself, whereas on the consequentialist's story, it is right because it leads to the best outcomes. The property that grounds the rightness of acts is different in these two explanations, even if the two explanations are necessarily co-extensive. We thus have an explanation for why deontically equivalent theories can nevertheless be explanatorily incompatible and, if my observations in Chapter 2 are correct, this explanation is very much in line with how many normative ethicists think of what they are doing.

Dreier has a reply to this. We could just consequentialize *again*, this time at the level of explanations. For example, if a non-consequentialist theory says that an act is wrong because it shows disrespect to persons, the consequentializing counterpart could say that the act is wrong because having shown disrespect to a person is part of its bad consequences. We could thus, as it were, consequentialize explanations. We could reframe the explanations that non-consequentialists adopt in terms of consequences.

And since the notion of a bad consequence, as Dreier has already argued, implies nothing more than what follows from it regarding our choices, the two explanations do not differ substantially either.

This option, consequentializing explanation, is not, I grant, logically impossible. However, I would argue that when we think through what it entails in a concrete case, it loses much of its appeal. To see why, remember that on Dreier's view, there really should not be any substantive disagreements between the consequentializer and their counterpart. As he states:

A moralist who subscribed to some non-consequentialist theory would not really be disagreeing with another moralist who subscribed to an equivalent consequentialist counterpart. (Dreier, 2011, p. 97)

Yet this seems strange when applied to a concrete case. Schroeder offers a good illustration, taking his inspiration from a version of non-consequentialism defended by Kamm (2000):

Kamm [...] argues that constraints are grounded in the moral inviolability of persons. But inviolability – having the status of a creature that may not be permissibly violated – is not something that we can affect with our actions. (We can cause someone to be violated or not violated, but we can't affect whether she is *inviolable*.) Consequentializing Kamm's theory would involve assigning disvalue to violations, which is not the same thing as valuing inviolability. (Schroeder, 2017, p. 1478)

Kamm, on Schroeder's reading, makes a claim about what grounds constraints. Consequentializers, as we have seen, can accommodate constraints in their framework. However, in order to do so, they have to claim that the outcome of someone not being violated is more valuable than all its alternatives. Yet, as Schroeder rightly points out, this is not what Kamm wants to say. Kamm argues that the rightness stems from a feature of the person, their inviolability, not the outcome in which a person has not been violated. Consequentializers, in order to consequentialize Kamm's theory, thus have to make a claim that is very much against what Kamm wants to say.

Dreier, at this point, could still try to raise doubts about whether we do indeed have a good sense of what the difference between these two explanations amounts to. As he says in a reply to a similar challenge from Portmore:

Labeling certain features “good” is not what matters, any more than labeling certain temperature differences as “degrees” matters. What matters are the

“active ingredients” – the features of outcomes that contribute to the verdict about what the agent ought to do. (Dreier, 2011, p. 114)

However, this reply amounts to little more than putting down one’s foot and insisting on the correctness of the Extensionality Thesis. This, I have argued, is not a strong dialectical position since ExT itself is not thought to be plausible by most people. To be clear, this does not amount to a rejection of Dreier’s view. The dialectical situation notwithstanding, he could still be right, and we will come back to the notational variants interpretation in Part III. At this stage, I just want to claim that, pending further arguments, the underdetermination interpretation seems to have a stronger footing.

Deontologizing and the Strong(er) Case for Radical Disagreement

What about deontologizing? Most of what has been said above applies to deontologizing. Just because we can deontologize all non-deontological theories, deontology does not become empty; nor do we suddenly find ourselves under the deontologist umbrella.

Interestingly, the case for the theories being radically different might be even stronger when it comes to deontologizing than when it comes to consequentializing. The reason for this is the following. Dreier, I have argued, might, as a last resort, still claim that even though the explanations that consequentializers and their non-consequentialist opponents advance seem incompatible, appearances might be misleading. And the specifics of consequentializing might lend some support to this claim. Remember that when assessing Parfit’s case study, I said that the default assumption should be that theories of different traditions are explanatorily incompatible. If these traditions are sometimes even explicitly defined against each other, that speaks strongly against the view that they are only notationally different. As a case in point, let’s once more take the temperature example. Nobody would be tempted to define Celsius and Fahrenheit in opposition to each other. In contrast, the new theories that result from consequentializing are, as we observed when discussing the technical side, not truly classical versions of consequentialism. They introduce a couple of modifications that allow them to copy their counterparts’ verdicts. Most importantly, the first move of consequentializing involves acceptance of a constitutive understanding of consequences. Yet, one might wonder, are we really sure that *on this understanding* of the term consequence, there remain any deep disagreements? Couldn’t non-consequentialists simply accept that, *on this understanding*, there really is no disagreement remaining?

This, I admit, is a live option and one I will take up again in Part III. However, it is important to note that this strategy does not work in the case of deontologizing. The difference between deontologizing and deontologized theories is very straightforward. The original consequentialist theories claim that it is outcomes that determine an act's deontic status. The deontologist counterparts claim that it has something to do with the act itself. In contrast to consequentializing, however, the relevant notion of outcome or consequence at stake here is the traditional one, not the more controversial, constitutive one.⁴² We start with a list of DOs and DON'Ts that copies all of the original consequentialist theory's verdicts. We then add an explanatory claim of the non-consequentialist variety. Importantly, that explanatory claim can be of the most paradigmatic form, for example, a standard Kantian explanation. For example, recall the suggestion by Portmore (2007, pp. 59–60) that we can *Kantianize* TAU by copying its verdicts and insisting that we treat humanity as an end in itself if and only if we give equal consideration to everyone's interests in maximizing aggregate utility. On this view, acts are right or wrong because they (fail to) treat humanity as an end in itself, which is a very standard form of a Kantian explanation.

The upshot of this is that if one wanted to claim that there are no substantial differences between the Traditional Act-Utilitarian view and the standard Kantian view, one would have to dispute that there is something relevant at stake between the most standard ways those traditions have been framed. One would have to claim that TAU and standard Kantianism are merely notationally different. Dreier, of course, accept this. Yet for most of us, I would suspect, this claim is even more difficult to defend. The dialectical presumption that at least the classical variants of the main traditions differ radically in the explanations they put forward is very strong, and there are good further reasons to believe it to be the case. Hence the burden of proof, if one wanted to claim that deontologizing and deontologized theories are merely notationally different, is much heavier. Whereas consequentializers' non-traditional understanding of consequences at least leaves room for doubt, it is not clear how deontologizing would do so.

Now some people might suspect that there's something in the deontologizing procedure itself that makes it doubtful that the resulting theories are indeed so different. The deontologizing procedure builds the

⁴² That is, at least, if we follow the recipe outlined above. It is less clear when it comes to the proposal by Hurley (2013), with its alternative evaluative framework.

(evaluation of the) outcomes into the (evaluation of the) action itself. This is how it achieves deontic equivalence with the consequentialist theory. But isn't that just a trick, and how different can such theories *really* be, if this is how one of them arrives at its verdicts? The deontologizing theory seems to be *parasitic* on the original consequentialist one. This, I agree, is a real issue, and we will take it up shortly. However, at this stage it is important to note that these objections do not concern the *radicalness* of the difference. Radicalness, remember, only means that the explanations cannot both be true. This seems to be the case with the utilitarian and the Kantian theories. Thus, when it comes to the question of radicalness, the picture of moral underdetermination that results from deontologizing is even more poignant than the one that results from consequentializing.