CHAPTER 1

Introduction

Normative ethics is in the business of determining what makes right acts right. A wide range of moral theories have been proposed in the literature, each of which specifies what criteria an act has to meet if it is to qualify as morally right. The focus of this book is a particular class of theoretical constructs known as *consequentialist* moral theories. All consequentialists believe that:

C. Whether an act is morally right depends only on consequences.

However, condition C does not capture all central commitments shared by consequentialists.[1] Every so often, consequentialists formulate views about whether an act is morally wrong, and sometimes about whether it is obligatory, permissible, forbidden or supererogatory.[2] Whether an act has any of these deontic properties also depends on consequences. The following general principle captures the central claim to which all consequentialists are committed:

C*. The deontic status of an act depends only on consequences.

Principle C* can be explicated in numerous ways. Is it the actual or expected consequences that matter? Must we always produce the best consequences, or do we just have to bring about sufficiently pleasant ones? Is it conformation to a set of rules (that lead to the appropriate sort of consequences) that makes right acts right, or should the right-making feature rather be identified with the consequences of each individual act?

[1] Condition C is identical to the definition of consequentialism proposed by Walter Sinnott-Armstrong in *The Stanford Encyclopedia of Philosophy* (2011). Note that it is not merely the consequences of the performed act that determine whether it is right or wrong. It is widely agreed that the consequences of all alternative acts also matter.

[2] The notion of moral wrongness is essential in discussions of moral dilemmas. It would be insufficient to just say that some acts are 'not right', since many entities that are not right have no deontic status whatsoever. Also note that claims about what is obligatory, permissible or forbidden are common in deontic logic. In the appendix, the relation between these concepts and the concepts of rightness and wrongness is discussed in more detail.

This book will make no attempt to address these familiar questions. The focus is on a different, so far unexplored, issue.

The hypothesis researched in this book concerns the structure of consequentialist theories. Briefly put, I conjecture that the best version of consequentialism is *multi-dimensional* rather than *one-dimensional*. One-dimensional consequentialists believe that an act's deontic status depends on a single aspect, such as the sum total of wellbeing it produces, or the sum total of priority- or equality-adjusted wellbeing. Multi-dimensional consequentialists believe that an act's deontic status depends on more than one aspect. They may, for instance, believe that the sum total of wellbeing *and* the degree of equality produced by the act directly influence the act's deontic status. Or they may believe, as I will eventually suggest, that the best version of multi-dimensional consequentialism is more complex and takes several further aspects into account.

To start with, it is helpful to consider a basic example with only two aspects. Imagine that you could bring about either a world in which Alice gets 100 units of wellbeing and Bob gets 50, or a world in which Alice and Bob get 60 units each. Which option would be best from a moral point of view?

Alice at 100	*Bob at 50*
or	
Alice at 60	*Bob at 60*

Clearly, the first option would bring about more wellbeing (suppose that this is the first aspect), while the second would bring about more equality (suppose that this is the second aspect). Utilitarians would clearly rank the first option above the second, whereas many egalitarians would favour the second. The preference of the prioritarian depends on the shape of the priority-function used for weighing wellbeing.

However, in the situation sketched above, multi-dimensional theorists believe that we need to make a genuine moral compromise between two conflicting aspects. In this very simple example, we cannot maximise both equality and the sum total of wellbeing, but both aspects directly influence each option's deontic status. Intuitively, the first option is right with respect to one aspect (wellbeing) while the other is right with respect to the other aspect (equality). Therefore, all things considered, no option is entirely right or wrong. Since none of the aspects trumps the other, each option's deontic status lies somewhere between the two endpoints of the deontic spectrum. The best way to render this intuition sharp is to maintain that both options are, literally speaking, a little bit right and a little bit wrong. Or so I shall argue.

It is essential to realise that multi-dimensional consequentialists do not believe that conflicting aspects can be aggregated into some *composite* aspect by assigning weights to each aspect and then maximising the weighted sum (or aggregate aspects using some other aggregation principle). Theorists who believe that conflicting aspects can be aggregated into some composite aspect do not defend a truly multi-dimensional position. If it were possible to aggregate two or more conflicting aspects into a new – possibly very complex – aspect, then those aspects could obviously be reduced into a single aspect, which would entail that the theory in question was actually one-dimensional. By definition, multi-dimensional consequentialists believe that an act's deontic status depends on two or more *irreducible* aspects.[3]

1.1 ONE-DIMENSIONAL VS. MULTI-DIMENSIONAL CONSEQUENTIALISM

Let me try to render the distinction between one- and multi-dimensional consequentialism more precise. To start with, something counts as a moral aspect if and only if it directly influences an act's deontic status, irrespective of how other aspects are altered. That something directly influences the deontic status of an act should be understood as a claim about functional relationships: an aspect, a, directly influences the deontic status, d, of an act if and only if d is a function of a.[4] The distinction between one- and multi-dimensional accounts of C* can thus, without loss of generality, be couched in terms of one- and many-place functions.[5] Let the set of C*-aspects be the set of all properties that can affect an act's deontic status according to consequentialist theories, e.g., the wellbeing produced by the

[3] Expressed in the terminology introduced in section 1.1, two or more aspects can be reduced to a single aspect if and only if there is a one-place function of some aspect that assigns the same deontic status to all alternative acts, under all possible circumstances, as that assigned by the many-place function that describes the relation between the initial set of aspects and an act's deontic status.

[4] The identity function is excluded from this definition. An act's deontic status is not an aspect that determines its deontic status. Also note that the notion of 'influence' captured by the notion of a function is very broad and could be further specified in numerous ways. The distinction between one- and multi-dimensional theories outlined here is intended to be compatible with all reasonable specifications.

[5] A function is an ordered triple of sets, $\langle A, B, F \rangle$, where A is the set of arguments of the function, B is the set of output-values, and F is a set of ordered pairs of arguments and output-values, such that every element in A is the first element in one and only one ordered pair. Many-place functions are functions in which the elements of A are vectors that contain more than one element. The functions considered in this book are not curried functions.

act, or its degree of equality, and so on. The key distinction researched in this book can then be stated as follows:

One-dimensional consequentialism $=_{\text{def}}$ the view that an act's deontic status can be characterised by a one-place function of some C*-aspect.

Multi-dimensional consequentialism $=_{\text{def}}$ the view that an act's deontic status can only be characterised by a function of several C*-aspects.

The word 'can' is key to this distinction. For something to count as a one-dimensional theory it is sufficient that it is *possible* to characterise an act's deontic status by a one-place function of some aspect. Whether it can also be characterised by a function of more than one aspect is irrelevant. To illustrate this point, consider a simple analogy with geometry. The area of a square can be characterised by a one-place function (of its side), but it can also be characterised by a function of two variables (of its length and width). The relationship between the area of the square and its four sides is nevertheless one-dimensional. In contrast, the relationship between the area and the sides of a rectangle is not one-dimensional: both its length and width are essential for determining the area. The general principle is that if it is possible to reduce a set of aspects into some smaller set of aspects, then we should do so – at least when pursuing theoretical inquiries.

The notion of a *moral aspect* is closely related to that of a *moral dimension*, but the two terms are not synonymous. In its most general form, a dimension can be conceived of as the conceptual space in which an aspect can be altered. Consider again the analogy with geometry. The area of the circle depends on only one aspect (its radius) whereas the area of the triangle depends on two aspects (its base and height).[6] All three aspects are elements of the same dimension (length). This is not always the case, however, as can be seen by considering an analogy with physics: mass and time are different aspects, but they are also elements of different dimensions. Strictly speaking, multi-dimensional consequentialism need thus not be a multi-dimensional theory. In extreme cases, all aspects that determine an act's deontic status could belong to the same dimension. In order to keep the terminology simple, however, I will keep the term 'multi-dimensional consequentialism', but otherwise respect the distinction between dimensions and aspects throughout the book. It is also worth

[6] Someone might object that the area of the triangle depends on just one aspect: the base times the height divided by two. My response to this objection is that the base times the height divided by two *is* the area, not the aspect that determines it. Other, similar objections that seek to challenge the individuation of moral aspects can be dealt with in analogous ways.

pointing out that the theory to be spelled out here is, as a matter of fact, truly multi-dimensional. Not all aspects that determine an act's deontic status belong to the same dimension.

Utilitarianism, prioritarianism and egalitarianism are the most well-known examples of one-dimensional consequentialism. Advocates of all these theories agree that an act's deontic status is determined by a single aspect. Utilitarians believe that an act is morally right if and only if the sum total of wellbeing it produces is at least as high as that produced by every alternative act. That is, the act's deontic status is a function of the relative order between the arithmetic sum of the wellbeing produced by the act and the arithmetic sum of the wellbeing produced by each alternative act.

Prioritarians believe that benefits to those who are worse off should count for more than benefits to those who are better off, but, as Derek Parfit explains, 'that is only because these people are at a lower absolute level. It is irrelevant that these people are worse off *than others*.'[7] This suggests that prioritarianism should be equated with the view that wellbeing has a decreasing marginal value in the same way that most people have a decreasing marginal utility for money, as observed by decision theorists and economists from Bernoulli onwards. Note, however, that wellbeing and utility are two distinct concepts: wellbeing denotes some more or less vaguely characterised state (usually a mental state, but sometimes a non-mental one) whereas utility has many other, often technically precise, meanings. In decision theory, for instance, utility commonly refers to a cardinal representation of a set of preferences over lotteries that satisfy certain structural conditions.[8] In the present exposition the focus will be on wellbeing, since wellbeing is more directly relevant in discussions of ethics than utility.

Egalitarianism is a claim about how to aggregate wellbeing enjoyed by different individuals. Egalitarian theories come in many different versions. Briefly put, the key idea is that an act's deontic status depends, at least partly, on relative differences in wellbeing. The mere fact that a certain act leads to someone being worse off than another contributes to the act's moral wrongness.[9] It is, of course, possible to maintain that equality is *all* that matters, but most egalitarians defend a more moderate view according to which relative differences, together with the sum total of wellbeing produced by the act, jointly determine the act's deontic status. The views defended by such moderate egalitarians are best described as composite one-dimensional moral theories.

[7] Parfit (1997: 214), my italics.
[8] See Peterson (2009b) for an overview.
[9] See Broome (in press), Parfit (1997) and Rabinowicz (2002).

Which are the most prominent multi-dimensional theories? A casual survey of the literature shows that many views that *appear* to be multi-dimensional are, on closer inspection, one-dimensional. Consider, for instance, value pluralism. Value pluralism is the axiological hypothesis that there are many different types of values, which cannot be reduced to a single feature of an act or entity, such as happiness or preference satisfaction. Several consequentialists accept pluralist axiologies. For instance, Fred Feldman defends a pluralist axiology according to which 'the intrinsic value of an episode of pleasure or pain is a function of two variables: (i) the amount of pleasure or pain the recipient *receives* in that episode, and (ii) the amount of pleasure he *deserves* in that episode'.[10] Other, less radical, examples of value pluralism can be found in the works of John Stuart Mill and George Edward Moore. In Mill's view, the value of a pleasurable experience depends on whether the pleasure is of a 'higher' or 'lower' kind, and any amount of the former kind is more valuable than every amount of the latter kind. Moore's pluralism is more modest. He argues that there is a plurality of *bearers* of value, but he is a monist at the foundational level. On his view, intrinsic goodness is a simple and non-natural property.[11]

Value pluralists need not, and typically do not, reject one-dimensional explications of C*. The claim that an act's deontic status depends on a single aspect is compatible with the claim that the intrinsic value of a set of consequences cannot be reduced to a single super-value. Mill explicitly acknowledges this point. He writes: 'It is quite compatible with the principle of utility [i.e., Mill's one-dimensional version of C*] to recognise the fact, that some kinds of pleasure are more desirable and more valuable than others.'[12] Mill's point is that his value pluralism is a claim about the intrinsic value of consequences, not a claim about moral rightness, and he can therefore accept the claim that an act is right as long as no other act brings about more intrinsic value. This leads directly to a one-dimensional explication of C*. Briefly put, Mill can be said to claim that an act is right if and only if the sum total of intrinsic value of its consequences, when aggregated in the appropriate way – that is, by acknowledging that any amount of a higher pleasure counts for more than every amount of a lower pleasure – is not exceeded by the intrinsic value brought about by some alternative act. Although Moore rejects Mill's claim about higher and lower pleasures, he could nevertheless accept the same one-dimensional criterion of rightness.

[10] Feldman (1997: 162). [11] Moore (1903/1959). [12] Mill (1861/1963, vol. x: 241).

Feldman's value pluralism is also compatible with one-dimensional explications of C*. Although Feldman believes that the intrinsic value of a set of consequences depends on both the amount of pleasure at stake and desert, this axiology is consistent with the one-dimensional claim that an act is right if and only if the sum total of intrinsic value brought about by the act is not exceeded by what could be attained by some alternative act.

In this context, it is also worth considering the relation between multi-dimensional consequentialism and John Broome's proposal for how to aggregate wellbeing.[13] In *Weighing Lives*, Broome proposes a theory of how to aggregate wellbeing across people, times and possible states of the world. He explicitly refers to these three entities – people, times and states – as 'dimensions'. However, Broome's aim is not to determine an act's deontic status, but to measure its 'goodness'. Strictly speaking, Broome's view is thus neither multi-dimensional nor one-dimensional. However, by adding suitable claims about how goodness is related to an act's deontic status, his theory could be rendered compatible with either type of consequentialism.[14]

Multi-dimensional consequentialism should furthermore be distinguished from dual-ranking theories. According to Doug Portmore's dual-ranking account of consequentialism,

> the permissibility of an act is a function of how its outcome ranks relative to those of its alternatives on a ranking of outcomes (i.e., the principal ranking) that is in turn a function of two auxiliary rankings: one being a ranking in terms of how much *moral* reason the agent has to want each of these outcomes to obtain, and the other being a ranking in terms of how much reason, *all things considered*, the agent has to want each of them to obtain.[15]

On the face of it, this may appear to be a multi-dimensional theory, since an act's deontic status seems to depend on two separate aspects. However, note that Portmore claims that there is such a thing as a *principal ranking*: the two auxiliary rankings do not directly influence an act's deontic status but rather determine its deontic status indirectly. Portmore's theory should thus be classified as a one-dimensional explication of C*, in which the two

[13] Broome (2004: 12–18, 104–6). See also his (1991).

[14] The criticism I offer against one-dimensional consequentialism in chapter 7 may also apply to Broome's account of goodness. Broome would, for instance, not accept what I call the trade-off principle. It is also worth mentioning that, while Broome might be tempted to accept principles C1 and C2 (stated later on in this chapter), his view seems to be incompatible with principle C3 (also stated later in this chapter).

[15] Portmore (2011: 118).

underlying aspects are merged into a new, composite aspect, which ultimately determines the act's deontic status.[16]

In the same way that one-dimensional explications of C* are compatible with several different axiologies, multi-dimensional consequentialism is compatible with a number of different axiologies. Multi-dimensional consequentialism is a claim about the *structure* of the features that determine an act's deontic status, not a claim about the intrinsic value of consequences. In many discussions of consequentialism, it is tacitly assumed that only one single feature of a set of consequences is valuable for its own sake, but this is a controversial axiological standpoint that need not be accepted by advocates of C*.

To sum up, multi-dimensional consequentialism is not equivalent to value pluralism or any other major view proposed in the literature. Multi-dimensional consequentialism addresses a different issue: this theory is not a claim about value, it is a claim about the structure of the entities that determine an act's deontic status. Many prominent views that on first appraisal appear to be multi-dimensional are, upon closer inspection, sophisticated one-dimensional views.

1.2 THREE KEY CLAIMS

The multi-dimensional explication of C* proposed in this book consists of three key claims. As explained above, the most striking difference between this theory and traditional consequentialist theories is the idea that an act's deontic status does not depend on a single (composite or non-composite) aspect. The following principle summarises the first claim:

C1. The deontic status of an act depends on several irreducible moral aspects.

Principle C1 follows from the definition of multi-dimensional consequentialism stated in section 1.1, so this is a claim that all multi-dimensional consequentialists are bound to accept. However, in order to formulate a normatively plausible multi-dimensional theory, which fits well with our considered intuitions, two further non-definitional claims need to be added. Both these claims raise substantial moral issues and are logically independent of C1. Consider first:

C2. The binary relation 'at least as good consequences as' is not a complete ordering.

[16] Portmore formulates his criterion as follows: 'S's performing φ is morally permissible if and only if there is no available alternative, ψ, that S has both more *requiring reason* and more reason, all things considered, to perform, where a requiring reason is just a reason that has some moral requiring strength' (Portmore 2011: 137, his italics).

In order to illustrate principle C2, imagine that the president of the world's only superpower could save the lives of a large number of soldiers by not going to war against the enemy. From a multi-dimensional perspective, not going to war would make the consequences better vis-à-vis one aspect (the actual amount of wellbeing produced by the act), but it does not follow that not going to war would be optimal with respect to all aspects. Presumably, ignoring the aggressive signals from a hostile enemy could be a risky strategy. Therefore, a short but successful war might be better vis-à-vis another aspect: reducing the overall risk faced by the nation, since that would prevent the enemy from developing weapons of mass destruction. Although not going to war may *actually* save many lives, this decision *could* lead to a disaster further down the road.

How should these two aspects, (i) the wellbeing actually generated by the act, and (ii) the risks triggered by the act, be balanced against each other? According to C2, the binary relation 'at least as good consequences as' is not a complete ordering, meaning that it is sometimes false, all things considered, that one act has at least as good consequences as the other. The upshot is that it is impossible to establish a precise exchange rate between all relevant aspects. However, although the all-things-considered ordering of consequences is incomplete, it is nevertheless helpful to assume that it is reflexive and transitive: The consequences of every act are at least as good as themselves; and if the consequences of one act are at least as good as those of a second, and those of the second are at least as good as those of a third, then the consequences of the first act are at least as good as those of the third.

In order to further illustrate why C2 is a plausible claim given C1 and C*, it is helpful to take a fresh look at the example discussed earlier in this chapter, in which an act's deontic status depends on some measure of equality and the sum total of wellbeing generated by the act. Imagine a two-dimensional geometric plane that represents the two aspects separately along two orthogonal axes. Consider figure 1.1. All things considered, which consequence is best: the one that scores 120 on the horizontal axis (wellbeing) and 100 on the vertical (equality), or the one that scores 150 on the horizontal axis (wellbeing) and 50 on the vertical (equality)? It is of course true that the numbers *could*, in a mathematical sense, be aggregated into a one-dimensional ranking by applying various mathematical formulas – and this is what one-dimensional consequentialists think we should do. However, to claim that 120 units of wellbeing and 100 units of equality is exactly as good as, or better or worse than, 150 units of wellbeing and 50

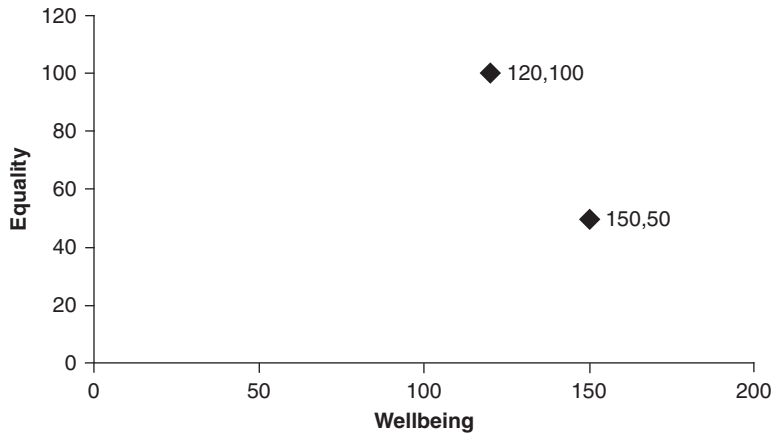

Figure 1.1

units of equality is a substantial moral claim. There is no a priori reason for thinking that consequences can always be aggregated into such an ordering. Sometimes there is, arguably, no complete (one-dimensional) ranking of consequences that captures everything that we intuitively feel to be morally relevant.

Claim C2 will be discussed in detail later on. However, it is worth pointing out here that there are at least two ways in which an ordering can be incomplete. Firstly, the elements in the ordering can be incomparable. This means that, for some consequences, no pair-wise evaluative comparisons can be made. Secondly, the elements in the ordering can be 'on a par'.[17] By definition, two elements are on a par if and only if they are comparable, although it is false that one is at least as good as the other. Parity is thus a unique, positive value relation, one which differs fundamentally from the traditional relations 'at least as good as', 'strictly better than' and 'equally good as'. I will return to C2 and the distinction between parity and incomparability in due course.

The third key claim of multi-dimensional consequentialism, C3, builds on C2 but addresses a different and more fundamental question in moral philosophy:

C3. Moral rightness and wrongness are non-binary entities, meaning that moral rightness and wrongness vary in degrees.

[17] See, e.g., Chang (2002.)