

# Consequences, and Moral Mathematics

## 2.1 Introduction

We have met briefly with the idea of a moral theory based upon consequences in our first chapter (see section 1.3). It may be supposed that no moral theory can regard consequences as entirely irrelevant, though few have gone as far as Kant (chapter 4) in detaching consequences entirely from the issue of moral goodness and moral action. Consequentialist theories, of which utilitarianism in its various forms is the most common and most discussed, form an extreme case of 'taking consequences seriously'. A consequentialist theory will have at its core a value - the promotion of pleasure, say - and will see all moral judgements in the light of this value. The most common forms of consequentialist theory are those which are labelled, by proponents and opponents, 'utilitarian'. This is quite a good label, because what the utilitarian takes as the measure of the rightness or wrongness of an action is its usefulness or utility in bringing about some change, either promoting or inhibiting a desired consequence. The right action will typically be the one which maximizes the value aimed at, with other actions being less right in accordance as they cause less of the value to be achieved. Some theories speak of 'happiness' as the desired outcome, but what they mean by happiness may be different. For some it may be a subjective state or feeling (pleasure and the absence of pain); for others it may be the bringing about of a situation where people's desires or preferences are satisfied, which may not be the same as their feeling the most pleasure or the least pain. There is some controversy about where the great 19th century utilitarians stood on such questions, though it seems plausible to hold that John Stuart Mill viewed preference or desire satisfaction as most important. Amongst modern utilitarians Peter Singer, whose book *Practical Ethics* has become a widely used text in moral philosophy and medical ethics, adopts a self-consciously 'preference utilitarian' approach.

It becomes clear very quickly that whatever the value seen as central, there will always be cases, invented or real, which identify clashes between the theory and our moral intuitions. As we suggested in the first chapter, it is an open question how far such clashes ought to be taken as indications of shortcomings in the theory, and how far they should be seen as indications of inadequacies in our everyday moral perceptions. Examples abound where it appears that an action may be necessary to promote pleasure in a population, but where the promotion of pleasure comes at the expense of an action which we would normally call unjust. The centrality of the language of justice to our moral perceptions may make us wonder whether there are actually a number of values at stake in morality (at the very least, pleasure and justice). After all, it does seem a bit odd to say something like 'Well, I agree that this would be unjust, but is it wrong? It seems, to many people at least, that justice is an absolutely basic moral requirement. What we need to do, perhaps, is to examine the relationship of pleasure to justice. If injustice could serve as a reason to reject an action which we agree would maximize pleasure, then clearly 'maximizing pleasure' and 'justice' are distinct values. Perhaps not doing an unjust action always takes precedence over doing a pleasure-maximizing action.

Typical consequentialist theories do not operate with a plurality of values, but see all values as aspects of one. If that one value is pleasure, or happiness, then our commitment to justice will need to be justified by reference to the value of pleasure, or happiness.<sup>1</sup> If there were more than one value, we could not judge an action's moral value by consequences alone. We would need to insert an explanation of the ways in which the values fitted together, and show why pursuing some values could not be justified if it were at the cost of violating others.<sup>2</sup> This explanation could not be expressed in terms of a higher value which arbitrated between the various competing values, because then the higher value would be the only value which really counted.

## 2.2 What value do we seek?

One of the advantages of a moral theory based around consequences is its simplicity. Moral decisions are seen as judgements about outcomes. Bad action is, from the perspective of the agent, action aiming at less than maximum utility, and from a more "objective" perspective it is action which fails to maximise utility (whatever the agent thought, or aimed at). There are, of course, a range of rather less than simple questions which come in the wake of this summary account.

One of the first is the judgement regarding what value is to be promoted. It may be, in fact, that the theory is not a theory promoting anything at all. Instead it may aim to work against something. The more classical utilitarian theories thought of pleasure and pain, or happiness and suffering, as indicative labels for the directions on a single 'scale' or "measure" of value. A unit of pain reduced was a unit of pleasure added. It is actually very far from obvious that such a single scale exists. One challenge lies in the apparent hierarchy of pleasures which some philosophers recognized. It is not that a 'higher' pleasure is simply 'more' pleasure. Recognition of, and desire for, a higher pleasure might plausibly consign one to a life of disappointment and frustration, so that believing in the existence of possibilities for higher pleasures, and struggling unsuccessfully to get them, may go hand in hand with much suffering, where being satisfied with lower pleasures would not. However, as John Stuart Mill noted, few of those who believe that there are higher and lower pleasures would sacrifice the pursuit of higher pleasures for a life of satisfaction with the more mundane.

This is one aspect of a problem we could term 'the commensurability problem'. Two things will be commensurable if we can put them onto a common scale and weigh or measure them against each other. I can hope to work out the relative costs of owning one of two cars by making certain assumptions (I won't crash, there will not be serious manufacturing defects) and then working out their likely life, fuel consumption, road tax and other costs, etc. The calculation involved may not be at all easy. If I buy a smaller car I may sometimes need to hire a larger one when relatives visit, but would the extra fuel costs of owning a larger one be more than hiring one occasionally? Yet in all of these judgements I have a single value -

---

<sup>1</sup> That is what John Stuart Mill attempts in his book *Utilitarianism* (1863, many subsequent editions).

<sup>2</sup> This is an aspect of the revised theory of natural law developed by Germain Grisez and John Finnis; see for example John Finnis, *Natural Law and Natural Rights* (Oxford: Oxford University Press, 1980) and John Finnis, Joseph Boyle and Germain Grisez, *Nuclear Deterrence, Morality and Realism* (Oxford: Oxford University Press, 1987), which contains an excellent summary account.

monetary cost of ownership - in view, and I can factor all the costs into my final verdict in a simple way. I just add them up.

In making a choice of career, it can hardly be so simple. Personal fulfilment, financial security, the desire to serve others and perhaps to be well thought of by friends; all of these will be factors I may take into account. How much money is personal fulfilment worth, if a tradeoff is to be necessary? The question may seem at one level absurd; at another it may seem obscene. (Imagine that the choice were between marrying a particular woman and having a certain level of financial security. We might say: if the question arises for you at all, then she certainly shouldn't marry you!) At the level of absurdity, the problem lies in the pretence of a single scale of value to which all our moral judgements can appeal.

Certain forms of utilitarianism recognize this problem. A minimal recognition of it lies in the claim that pain and pleasure are themselves incommensurable, so that a moral imperative to reduce pain and minimize suffering would not be the same as an imperative to promote pleasure. This would set a barrier against forms of utilitarianism which could sensibly ask whether if it gave *enough* pleasure to a sadistic games teacher to see his pupils suffer the suffering might be justified. 'Negative' forms of utilitarianism are those which mark such distinctions as this and seek to promote actions which reduce the negative moral quality, even at the expense of actions which promote some accepted good quality. This is because there is no amount of suffering which could be the acceptable price of a certain amount of pleasure.<sup>3</sup>

## 2.3 How will it be for you?

Another area of difficulty and a distinct one lies in the comparison of pleasures between persons. It was a desire of utilitarianism in some of its forms that each person should count for one and none for more than one, but to count anyone else is to require some measure of the amount of utility given to him or her as the outcome of some action. Knowing how much utility will be derived by giving the donor heart to person X rather than person Y is not just a question of problems in predicting the future, because as long as the measure of morality lies in subjective feelings it lies in something which cannot be compared. How can I say how much pleasure it would give *you* rather than to *me*? Only, perhaps, by an act of imagining the pleasure that it would give to me and then projecting the results of my judgement on to you; by assuming, other words, that you and I are alike. Maybe we are not alike; maybe you would get greater pleasure or get pleasure of greater value.

There are problems with comparisons at a variety of levels. If I have never had children nor serious mobility problems and am trying to allocate resources across a variety of health services which includes infertility treatments as one candidate for funding and hip replacements as another, do I have the resources to hand to help me make such a decision? I can engage in acts of sympathetic imagination which enable me to contemplate the awfulness of immobility quite easily. I recognise the suffering of thwarted desire in those who desire

---

<sup>3</sup> There are connections to be made here with certain approaches to questions of the existence of God, and of evil and suffering, in the philosophy of religion. A 'greater goods defence', one which sought to justify evil as *outweighed* by certain goods, would not be likely to appeal to a negative utilitarian.

parenthood, but if I do not desire children myself I am unlikely to see it simply as ‘thwarted desire’ and not understand the precise nature of its frustration to them. One obvious move we might make accepts the difficulty as a real and intractable one in my own case and seeks help from an ‘expert’. ‘Expertise’ here is found in the persons of those who have had experience both of fertility problems and of mobility problems resulting from joint deterioration. We may hope that such people could give relative weight to the frustrations and guide us in our allocation of resources.

An immediate reaction to this proposal will be to wonder whether every expert will come up with the same weighting, and if not, how damaging that might be for a resolution of our difficulty. One move which utilitarians might make is to accept this as a potential problem, but to say that where the relevant experts, as we are calling them, consistently state an overwhelming preference in one direction rather than another, that is a non-arbitrary reason for regarding this as a utility-maximizing direction of policy. If the frustration of infertility overwhelms that of immobility (or vice versa) in the case of every person we ask who has experienced both, maybe the whole of our budget allocation should go in the direction of the one felt to be most frustrating.

Alternative methods of carving up allocations of health budgets have been considered, including that of asking the potential recipients of the services to allocate priorities across the range which may be made available. The potential recipients will not (all) be ill, will not (all) be even potential sufferers from the afflictions which may be treated, or needful recipients of the services provided. There will be allocations between men and women in services such as testicular cancer screening and treatment, and care in pregnancy. Because they will not all be ill or pregnant, we might wonder whether the act of imagination involved in setting priorities based around the negative utility of suffering from an untreated condition is one which they can realistically make. Because not all people are even potential bearers of some of the conditions under consideration, we discover there are situations where the concept of a ‘relevant expert’, suitably experienced in the relevant set of frustrations and able to decide between them, just breaks down.

This need not be the end of the project. In some circumstances it may well be that, viewed from a utilitarian perspective, and allowing that not all frustrations are of the same order of significance, it is still a good way of allocating relative significance. Even in a simple situation where there are only two possible recipients of funds, though, we would need to decide whether (say) *all* of the funds should go to infertility treatment and *none* to hip and knee replacements (or at least none until all infertility cases have been adequately treated and the most potentially fruitful infertility research funded), or whether there is a presumption that both should get some money because both have some value (though hips and knees get far less).

The discussion so far may make us skeptical about the possibility of applying utilitarian principles to real-life dilemmas. Interpersonal comparisons seem difficult to make, and the one available criterion for making them in a non-arbitrary fashion (finding people who have the capacity to make the relevant comparison from their own subjective experience), even if it were possible in some cases, is ruled out absolutely in others.<sup>4</sup>

---

<sup>4</sup> It is also worth remarking that certain aspects of one of the specific issues mentioned, infertility, are likely to be under-played or ignored by this method. Infertility might be seen as a medical problem for specific individuals, but often it will be a frustration for a couple, and will be experienced very differently as a phenomenon in one couple rather than another. The subjective assessment of relative

One possible response here would be to deny the impossibility of comparing utility values interpersonally; another would be to admit the impossibility, but to assert that even though it is in practice impossible we still have to do it if we are to make moral judgements at all. Perhaps the considered view of a utilitarian theorist should be that it is indeed impossible, on sober reflection, to do what we are being asked to do, but that it is what we believe ourselves committed to anyway in any sort of recognition of the value of happiness and disvalue of suffering in others. If we are committed to this, we will just have to make the best job of it we can, recognizing the possibility of challenges and the necessity of revisions if we find our own valuations wildly out of line with most other peoples' ones.

## 2.4 Local and remote consequences

We noted earlier two ways in which actions might be assessed; by the 'spirit' of the action (the desire and intention of the agent to promote the best consequences), and by the actual consequences which that action results in. The distinction is important, because the actual consequences of our actions are impossible to predict. I save a drowning man, and in later life he becomes a tyrant - am I responsible? I take the car to work when I could have cycled, and it is therefore unavailable to you when you need it in an emergency later in the day. Someone has a baby; the difference this will make to the future depends in part on all the other children (many as yet unborn) who may share his life in different ways.

The cases mentioned are ones where it is impossible to predict. There are others where we might want to change this to 'difficult', but agree that the consequences are in some measure predictable. Were we to agree that consequences are the moral measure of our actions, we do not need to hear this as an invitation to disregard responsible reflection altogether, along the lines of "One can never tell what the outcome of our actions will be, so one might as well just do whatever one fancies." Where the outcome of an action depends upon the behaviour of the stock market, or the character of a friend, there are people who are better and worse judges of share movements, and there is more and less superficial reflection on character. It is natural to say that, where personal responsibility in a consequence-driven moral framework is concerned, we should be responsible for the best assessment of consequences we can manage, and for acting in accordance with the best judgement we can make.<sup>5</sup>

---

weighting to be given to fertility treatment over mobility assistance by one individual does not take in to account the impact of the treatments on the partner, or on other dependent relatives. They may well not be asked for their views at all. Utilitarians will also need to account for the utility accruing to any as-yet-unborn person who might be born as a result of the fertility treatment, which introduces a further, and very significant, complication.

<sup>5</sup> This may mean judging who is likely to be the most reliable judge. Choosing a pension fund will often have this character, where there has really been proper choice and not the exercise of that 'method' often known as 'the calibrated pin'.

These thoughts present a radical divergence between the assessment of an actor (an agent) and of an action. An action is right or wrong in accordance with its consequences; an agent is praised or blamed according to his or her having done the best possible in the circumstances to promote the right consequences. Now, the consequences of our action could plausibly be said to be captured by the ugly formula: 'the total outcome in terms of utility - value throughout world history from this point, given that I did *this*'. There will be a variety of alternative histories where I did *that* or *the other* or *nothing at all* (which, because this too has a utility outcome, has to count as an action for our purposes). The right action is (objectively) the one with the maximal aggregate. If two possible actions yield the same total utility value, and they are each actions which maximize possible utility, there is presumably no moral basis on which one might decide in one way rather than the other.

## 2.5 Acts and Rules

It will be obvious that one difficulty with taking utilitarianism seriously lies in the expectations which it places upon the moral agent. One way in which these expectations are difficult concerns the pervasiveness of utilitarian moral obligation. Always, at every time and in every circumstance, what I do or fail to do may have an impact on utility. Thus, I find myself seeing the time I set my alarm clock to go off as a matter, not simply of how long I would like to sleep, or when I need to be at my first appointment, but of moral duty. Maybe the sleep would do more good for me than any good which could be done for anyone else by my rising earlier. This is not obvious, though. It would need to be calculated; would need to be thought through.

The other major area of difficult expectations lies just here, in the business of 'thinking through'. If I took the calculation of the consequences of my actions very seriously, I would never act but always be calculating! There must come a point when the cost-benefit analysis of the calculation comes to an end, and that not simply because we can never be 100% sure of the consequences. Thanks to the mathematics of complexity theory (sometimes colloquially known as 'chaos theory', though the mathematical systems studies are hardly 'chaotic') we are now able to see the degree of detail which we would need to know in order to predict the development of weather systems. We would need to know the movements of each bird and butterfly in what remains of the Amazonian forest. Even if the prediction of the consequences of our actions involved no future uncertainties, the level of knowledge we would need, even to be *fairly* sure of the full range of these consequences, is unimaginably vast. At some point calculation must cease, and there be action.

An approach to utilitarianism which seeks to take this problem seriously is called 'Rule Utilitarianism'. The central thought of this form of utilitarianism is that, in most situations, certain rules of behaviour will achieve the utility-maximizing outcome without the disutility of calculation. I have made a promise; should I keep it? On what we might call the 'classic' understanding of utilitarianism, where each action is judged as an individual case, the agent must assess the consequences of keeping or breaking the promise, and decide which course utility demands. Because it judges actions case by case, with no presupposition about what is likely to be the right answer, this classic form is often called 'Act Utilitarianism'. Recognizing that, on the whole, the consequences of breaking promises will usually be worse than keeping them, the Rule Utilitarian proposes a rule: 'keep your promises'. This rule

guides us in deciding what to do where we have made promises, and saves calculation.

This isn't, though, the whole story. Two possibilities present themselves: one, that the rule is to be obeyed always, and that the universal acceptance of the rule is what has utility; another, that the rule is a rule of thumb, generally applicable and useful in avoiding the calculation problem, but having no universal binding force. If I came up against a situation where the principle of utility would obviously be less well served by keeping the promise than by breaking it, ought I to break the promise?

It seems, on utilitarian grounds, that I should indeed break it. So the rules are on such a view *merely* helpful guides for when what to do is not obvious. Such 'rule-of-thumb utilitarians' are actually a species of pragmatic act utilitarian, recognizing the problems faced by utilitarian agents and seeking to offer helpful advice for them. If we could refine the rules of thumb sufficiently, they would get closer and closer to a specification of what ought to be done in some very precisely described situations; closer and closer to the ideal of act utilitarianism itself, in fact.

There are, however, utilitarians who believe that we can never do without rules, and that the rules have utility in a variety of indispensable ways. They structure our expectations of the behaviour of others, and may even have some formative effect on our desires and satisfactions. To this extent, there can be forms of Rule Utilitarianism which commend absolute obedience to rules as a means of promoting utility. The irony then would be that moral agents might obey rules without any clue that the reason why the rules are the way they are comes down to what promotes utility. A thought which is sometimes expressed in the literature on utilitarianism is this: if the most plausible utilitarian theories end up with people obeying rules (about promise-keeping, truth-telling, justice and so forth), how far is this really utilitarianism? And if people are best able to obey the rules when they are innocent of or ignorant of the moral philosophy of utilitarianism - because once they hear of it, they will be tempted to begin to look for exceptions to rules, such as situations where it is morally justified or obligatory to lie, and may thus undermine utility in a variety of ways - it seems that there is a curious self-subversion about utilitarianism, at least in this form.

These criticisms would not apply to act utilitarianism. It would remain a proper question, though, whether the extreme understanding of duty and responsibility promoted by utilitarianism could be said to undermine any increase in utility which belief in the theory would otherwise produce. A final irony might be found in the possibility that living in the light of some other moral theory would actually have better results for utility than act utilitarianism itself. The utilitarian might admit this point, and say that nonetheless utilitarianism makes a vital claim: that each agent is maximally responsible for the total effects of his or her actions (and inactions). We will return to an assessment of that claim in the final chapter of this course.