

Introduction to some basic ethical orientations

Rick Grush

[1] One of the major branches of philosophy is *ethics*, which involves, among other things, the attempt to articulate ethical principles that underlie our moral intuitions, as well as serve as a guide for morally relevant decisions we might face. We will follow convention in using 'moral' and its cognates to describe first-order situations and attitudes. Moral situations arise all the time, and moral behavior (or immoral behavior) is behavior that is assessable as good or bad. Ethics is *the study of* morally relevant situations and morally assessable behavior.

[2] The way I described the use of ethical principles just now had two components: a descriptive and a prescriptive component. A *descriptive* enterprise is one that simply observes and describes what happens. For example I might observe people who are speaking some language, and simply record what they say and try to determine the rules, if any, that determine what they say; perhaps they sometimes use double-negatives, like 'I don't want no eggs'. In such a case, I would simply observe what they do, and try to account for why it is they are doing that. This would be a descriptive grammar. In the context of ethics, I might undertake a descriptive ethical enquiry, trying to determine the principles that some person or group of people use when making their morally relevant decisions: perhaps they do whatever they feel like, or perhaps they do what they think God wants them to do, or perhaps they do what they think will lead to the greatest overall happiness of people. Whatever the details, the descriptive task would simply seek to determine what ethical principles, if any, are in fact being used by some person or group in their morally relevant behavior.

[3] By contrast, a *prescriptive* enterprise is one that attempts to articulate the 'correct' principles that people ought to apply. For example we are all familiar with prescriptive grammars that articulate what somebody thinks is the 'correct' way to speak a language. Someone involved in a prescriptive enterprise might try to correct the speakers mentioned above, by telling them that double negatives are incorrect. In the domain of ethics, a prescriptive ethics tries to determine what ethical principles we ought to adopt and employ when making morally relevant decisions.

[4] By and large ethical theory involves both components. The problem can be set up more or less as follows: there are roughly two kinds of morally relevant situation: those in which moral intuitions are very clear and practically universally shared. For example 'If Smith offers Jones \$5 to kill Smith's family because Smith is tired of hanging around them, should Jones accept?'. In this sort of case, almost everyone's intuitions converge on the same answer. The second and more problematic area are cases where moral intuitions within a given individual are less clear, and opinions among different people may diverge: is abortion permissible? is euthanasia? should those convicted of violent crimes be given the same rights of access to donated organs as normal citizens when such organs are in short supply? is it right to steal money if you need it to do something that, considered in itself, is a good action (e.g. stealing money from a hard working single mother in order to buy Christmas gifts for orphans)? The typical tactic in ethical theory is to use the cases on which there is a high degree of agreement in order to try to figure out what the ethical principles are that we are implicitly employing to arrive at those judgments, and then using those descriptively arrived-at principles in order to make a prescriptive claim about the less obvious cases.

[5] This discussion has been a bit abstract, but as we go on it should become more clear. In the next few sections, we will cover a few of the most influential ethical orientations.

1. Utilitarian views

[6] The basic utilitarian view is that when presented with a morally relevant decision between a number of options, the right one, *the one you ought to do*, is the one that does the greatest good for the greatest number of people. Before we look more carefully at what this means, let's look first at what it does not say. First off, it does not give anyone, including the person making the decision and that persons friends or family, any special status; it is thus, like most of the views we will discuss, *impartial*. Second, it quite avoids a number of other ethical outlooks, such as divine command (which would maintain that you do what God tells you, regardless of what good or bad comes from it), or deontological views that prescribe certain duties (e.g. respect the dignity of all people) regardless of the results. Third, it makes no immediate distinction between acting and failure to prevent an action: and so accordingly pushing someone off a building is exactly as bad as failing to help someone from falling off (assuming that one is in a position to help without putting oneself at risk).

[7] The basic idea of utilitarianism is easy and plausible enough: the morally right thing to do is to make the world a better place. Thus what makes some act right or wrong is its utility in bringing about good consequences, and not, for example, the intentions behind the act. So when presented with some decision one does one's best to calculate the expected utility of the different options, and then implements the option that has the highest expected utility.

[8] So what is expected utility? The first idea is that utility is cashed out in terms of something like happiness, and so we suppose that we can assign various quantitative values to various kinds of happiness (hedons) and unhappiness (dolors). We then, for each possible action we might perform, we determine what all the possible outcomes are, and how likely that outcome is. We then multiply the outcome times the likelihood of that outcome, and add up the result for each outcome. The final result is the expected utility of that action.

[9] For example, suppose I have \$100 dollars and someone gives me the opportunity to make a bet. I role a die, and if it comes up '6', I get \$1000, and if it comes up any other number, I loose my \$100. My two options are: take the bet, or don't take the bet. What is the expected utility of each? The expected utility of not taking the bet is easy. There is only one possible outcome (since I am not doing anything), and on that outcome, neither loose nor gain anything. So the expected utility of not taking the bet is \$0. (Note that I am now using dollars as the units of utility.)

[10] What about taking the bet? There are six possible outcome, one for each number that could come up on the die. One a 1 through 5, I loose \$100; and on a 6 I win \$1000. So first I determine how likely each option is, and then I multiply that likelihood times the loss/gain, and then add them up. So:

- 1: $1/6 \times -\$100 = -\16.67
- 2: $1/6 \times -\$100 = -\16.67
- 3: $1/6 \times -\$100 = -\16.67
- 4: $1/6 \times -\$100 = -\16.67
- 5: $1/6 \times -\$100 = -\16.67
- 6: $1/6 \times +\$1000 = +\166.67

[11] Adding these up, we see that the expected utility of taking the bet is +\$83.34. This means that this is what I can expect, as a sort of average, if I take the bet: a gain of about \$83. So this is a good bet, as far as expected utility goes.

[12] Now for a morally relevant example: we can either increase the number of students at UCSD, or keep the number of students constant. Each choice will have a number of certain and probable effects. Benefits of increasing the number of students is that a certain number of people who would otherwise not be able to attend UCSD would be able to attend. This would make them happy, and might also have beneficial effects for society in that it would make more people better educated, etc. On the other hand, it would lead to greater crowding on the campus and perhaps place larger demands on the facilities and staff. Taxes would have to be raised, which would make a great number of families slightly less well-off than they would have been, and it might also decrease the overall quality of education at UCSD if classes become overcrowded, etc. So what we would need to do is for each course of action (admit more students, or keep the number constant), come up with some idea of what all the costs and benefits would be and then compare the totals for each course of action.

[13] Now that the basics are done, on to some more details: what is the 'good' that we are trying to maximize? Some historical options have been: pleasure, happiness, preferences, and values. Pleasure (Bentham's position) is easy enough to understand. Some things bring people pleasure, such as eating and having sex. And we can also understand physical suffering, such as having toothaches or being beaten up. On this conception, we look for what sorts of pleasures and displeasures might arise from various choices, and maximize the pleasure. While this is easy enough to understand, it has some disadvantages. Most obviously, it seems to place bodily pleasure and enjoyment at too fundamental a place in ethical theory. It might, for example, justify just hooking everyone up to morphine I.V.s if we could get robots to run everything else in the world. This seems wrong -- it seems as though doing this would eliminate a number of things we value, and so perhaps mere pleasure is not the best thing to maximize.

[14] (Note that here we see the interplay between descriptive and prescriptive theories. On the basis of some moral intuitions about making the world a better place, etc., we come up with a tentative ethical proposal: the right thing to do is to maximize pleasure. If this were right, then it might be a principle we could appeal to in unclear cases. But we can see that the principle can't be right because it conflicts with some clear moral intuitions we have: most people would not think that a world run by robots full of morphine addicted humans experiencing great pleasure for their entire lives would be a good thing. So on descriptive grounds we know that this can't really be a principle that we in fact actually endorse, and perhaps nor should it be one we want to endorse. The hope is to find principles that accord with, and perhaps explain, our moral intuitions in all the clear cases, and can thus serve as guidelines in the unclear cases.)

[15] Another principle has been happiness (Mill's position), where happiness might include bodily pleasures, but also includes 'higher' pleasures such as appreciation of art, enjoying the company of family and friends; seeing one's children succeed; and so forth. While this seems better than pleasure, it has the problem that it is much more variable than pleasure: very different things make different people happy, and so determining what will lead to the greatest happiness may be no easy matter.

[16] A similar approach is to maximize preferences. We assume that each person has some set of preferences, and what we need to do is to maximize these preferences. An advantage of this is that in some contexts it might be possible to merge this with economic theorizing. For example, assuming that people spend money as a function of their preferences, the free market will evolve to maximize preferences. Thus we don't have to figure out if people prefer McDonalds or Burger King: we just let the market evolve and it will, almost by definition, evolve to maximize preferences. (At least that is the hope.)

[17] A final option would be to maximize the prevalence of some sort of value, such as freedom or knowledge. This has obvious advantages, but it is subject to a number of problems, including the fact that different people and different cultures may have different ideas about what values are the right ones, and if this is the case, our ethical theory will be powerless to make any decisions.

[18] In addition to having different choices for the good that we might want to maximize, there are two versions of utilitarianism that differ over what we are calculating the utility of: acts or rules. *Act utilitarianism* says that for each individual action we try to assess its consequences and then perform the act with the best expected utility. On an act utilitarian view, it is actions rather than rules that we are assessing. According to *rule utilitarianism*, what we are interested in assessing are *rules*, not individual acts. The two approaches might give conflicting results. So for example, a rule utilitarian might come up with convincing data that show that a law prohibiting smoking in public buildings is a rule that will lead to the best overall outcome. And so we might adopt that as the rule, and on its basis kick someone out of a building for smoking on a given occasion. On the other hand, an act utilitarian might reason that on this particular occasion, since only the smoker is present and the ventilation system is such as to remove all smoke before anyone else arrives, the best thing would be for the smoker to light up.

[19] There are some common objections to utilitarianism. The first two are based on the problem of determining consequences. First, it is often very difficult if not impossible to determine what the expected utilities of a given set of acts or rules will be to any degree of certainty, even when we are using simple goods such as money and bodily harm. Things are only worse when we move to other more tenuous and variable things like happiness and values. To this the utilitarian might respond that in some cases we can either come up with a decent guess, and in others the relative goods and bads will be quite clear. The opponent will respond to this that this might be right, but we need ethical guidance exactly in a great many areas that are unclear, and it looks as though utilitarianism, either rule or act, may not be able to help us in many of these cases.

[20] A second objection is that because of the difficulty or impossibility of predicting all of the consequences of an act, we might end up doing acts that are, by utilitarian principles, quite bad, even though our intentions are good and we were as rational and as responsible as we could have been. The man who swerved his car in order to not hit the little 5 year old boy in the street might, in some sense, be morally culpable if that boy is Adolph Hitler. In such a case, the 'right' thing to do would have been to run the boy down. The utilitarian might respond by saying that they are not in the business of assigning moral responsibility to anyone other than the responsibility to do the best they can with the utilitarian principles. While they might agree that, as it happens, running a five-year-old Adolph Hitler down in a car might have been a morally better act than swerving to miss him, they will refrain for saying that the agent is morally culpable, because the agent could not have known the future consequences. (But the utilitarian, especially the act utilitarian, would say that if the person knew that this boy would grow up to do the things he was going to do, and there was no less harmful way to prevent him from doing them, then the driver should have run the boy down.)

[21] Another set of objections has to do with determining whose values and goods we are to count in the utility calculation. This issue has a number of aspects. First, whose values and goods do we appeal to? Should we force another culture to change because we think it will increase their happiness even if they do not think so? (Do we pressure cultures that maintain that women are subordinate to men to change, even in cases where all the people, including the women, claim that they like it that way?) Do we count the utility of ourselves, all humans, all adults, all animals? And do we include all future generations? Arguably, if we take all

future human generations into account, tremendous utility could be gained by subjecting a few hundred thousand people now to a wide range of medical experimentation and tests that we would think now are unacceptable; but looking forward to the future, which might have many thousands of generations over perhaps millions of years, the medical gains made now might well outweigh the current suffering.

2. Deontological views

[22] The deontological orientation in ethics had its first major proponent in Immanuel Kant. This approach is broad, but to a large degree can be seen as based on the concepts of *duties* and *rights*. What is meant by duties and rights here must be understood correctly. We often think of duties as things that are imposed upon us by external powers, and rights as things granted by external powers. And surely there are many such duties and rights. I have the duty to pay taxes and to serve jury duty when called upon; and my license gives me the right to drive a car.

[23] But these are not the kinds of duties and rights that deontology is concerned with. When we speak of human rights, for example, we speak of rights that human beings have not because of some external power, this or that government, granted them this right. In such a case this power could take that right away. Rather, we mean rights that a human being has simply in virtue of being a human being, not because of any external license was granted to them for the purpose, but simply because they are people. And the duty to respect human rights is not derived from some law that we have to follow, but rather, the deontologist argues, is derived from within, from the nature of reason and morality itself. To some extent systems of laws written, adopted and enforced by governments are designed, in part, to protect such rights and enforce these kinds of duties. But the deontologist claims that even in these cases, the law is not *creating* the rights and duties. These were there all along. The law is simply codifying these rights and duties in an enforceable way.

[24] For Kant, one of the primary duties we have is to treat people, including oneself, always as ends, and never as mere means. An end is the final purpose for which one does something, and the means are how one achieves that purpose. So Kant is saying that we have the duty to always treat people as ends – that is, always as worthwhile in themselves – and never as merely having a use for achieving some other purpose. This ethical regulation simultaneously defines a duty all people have, and a right all people have. We might, very roughly, put it this way: we all have the duty to treat all people with respect. Another categorical imperative for Kant was to always tell the truth.

[25] Exactly what the rights and duties are is a separate question. But pretty much all versions of deontological theories will recognize the duty to respect human dignity, and the right to such dignity; and perhaps also a right to autonomy, and a duty to respect the autonomy of others. These core duties and rights are duties for all, and rights held by all. But the view also makes room for some special duties and rights that are not impartially distributed. For example, parents may have certain duties with respect to their children that their children do not have for their parents. Parents have the duty to provide for the protection, health and education of their children, say, but children do not have the same duty towards their parents (at least not while they are very young).

[26] Thus while, as we saw, utilitarianism was impartial, treating everyone as on the same footing, deontological theories typically have an impartial, universal core (such as the right to human dignity), but can also accommodate special cases of rights and duties. (As a note,

communitarian theories of ethics seek to show how one's position in a community, which might be a family, a team, a country, etc.) can impose various *prima facie* duties. We won't dwell on communitarianism in this course.)

[27] It is easy to see many cases in which utilitarian views and deontological principles give conflicting answers to ethical questions. For example in the movie *Minority Report*, the three siblings who could see the future were being used merely as means, their human dignity and autonomy was being sacrificed in order to achieve some purpose. Now it is quite arguable that the overall benefit, in terms of happiness, received to society was greater than the harm caused, in terms of unhappiness suffered by the siblings. In fact, probably far far greater. While a utilitarian does not like the fact that three people are being made to suffer in some sense, the fact that potentially many *thousands* of violent crimes were being prevented outweighed this fact. Why, the utilitarian will ask, should we allow *thousands* of people to be violently murdered, rapes, etc, just so three people can be spared a loss of freedom? And so using them in this way would be the right thing to do by utilitarian standards. A deontologist, by contrast, is much less concerned with overall happiness, and would state that in fact doing this is wrong, because the rights of these three people are being violated. They are being treated as means, and not as ends. Their dignity and autonomy are being violated.

[28] This leads to another difference between deontological and utilitarian views. Deontological views typically must claim that there is a very morally significant difference between causing something, and failing to prevent something. This is similar to the distinction between acts of commission and acts of omission. Utilitarians do not recognize this as a significant difference. Note that in the *Minority Report* case, one could argue that in order to treat the siblings with dignity, one would have to allow thousands of people to be victims of violent crimes such as rape and murder. The utilitarian sees the two options as (a) imposing on the dignity of three people while respecting the dignity of thousands; or (b) respecting the dignity of three while effectively sentencing thousands to be the victims of violent crimes (which is also an affront to their dignity, of course). For the utilitarian, *who* is causing the crimes or indignities is not the a relevant concern. But for the deontologist, this is a crucial concern, for if *I* am in the position of releasing the siblings or keeping them confined, then this is the issue that *my* duties concern. It is true that if I let them free, I will effectively be failing to prevent thousands of violent crimes. But on the deontological view, what actions I perform carries more moral significance than what things I fail to prevent. And so my duty would be to treat these three siblings as ends, respect their dignity, and free them.

3. Contractarianism/contractualism

[29] Finally, a brief mention of contractarianism (my brief discussion will more or less conflate contractarianism and contractualism). The contractualist is concerned to find a justification for specific ethical norms, a justification that doesn't *necessarily* come from expected utility calculations or appeal to rights and duties. For the contractarian, ethical norms are those that are or would be agreed to by rational agents in some society. For example, Hobbes felt that the power of sovereign monarchs was justified because rational agents would realize that without some sort of power with various rights and duties, social life would dissolve back into a state of nature without order, and would be, as he put it, "solitary, poore, nasty, brutish and short." By agreeing to give up certain elements of one's own autonomy in many circumstances and give them to a sovereign, a group of people can thereby create an orderly, lawful and safe social setting. Notice that there are some similarities in this line to a utilitarian, but there are crucial differences. The similarities is that as Hobbes describes it, each agent reaches the conclusion that giving authority to a sovereign is a good choice based

on utilitarian principles: doing that will lead to better consequences than not doing that. The difference is crucial, however, for the contractarian does not claim that the utilitarian calculus is what justifies the rules that are adopted. What justifies them is the explicit or implicit agreement of those in a society, their social contract. Why an agent would agree to such a contract is another matter. Hobbes felt it was by and large a matter of utilitarian reasoning.

[30] More recent versions of contractarianism are similar. Rawls, for example, asks us to imagine being behind what he calls 'the veil of ignorance'. Behind the veil, each agent knows more or less how the world is structured, but does not have any idea who one is, one's age, race, gender, cultural background: all particular information, is unknown. Rawls then asks what ethical principles would one agree to in those circumstances, knowing that once the agreement is made, one will then be placed in some particular situation in the world (particular gender, race, culture, etc.). Rawls claims that the ethical principles rational agents would agree to in such circumstances are then binding.

[31] Exactly what ethical norms would be agreed to by people in such a situation is not really clear. One might conclude that a system of rights and duties would be the centerpiece of a social contract one would want to endorse; or one might opt for a utilitarian ethics; or one might opt for some entirely different set of ethical principles. Various contractarians have of course not only defended the basic idea of what justifies a certain set of norms (the social contract), but have also argued about what specific set of norms rational agents in such a situation would agree to (as Hobbes argues that rational agents would agree to give power to a law-providing and law-enforcing monarch). But the main point is that contractarianism is not primarily an account of what specifically the correct ethical norms are, but it is rather an account of what gives some set of ethical rules their normative force. On this account, what gives a set of ethical rules normative force is the fact that a rational agent would agree to them in order to construct an acceptable social order.