

Chapter 9

Some Ethical Theory

Introduction

If you reflect upon what has occurred so far, we have had a content focus upon ethical issues arising in one variation or another across a range of professions. We have also had a procedural focus. Conceiving of professional ethical problems as ones within applied ethics, we have sought to track connections from the problem at hand to whatever underlying ethical values are being applied. We have written the latter in as various moral premises in the web of argumentation that articulates such connections. So, our professional ethical enquiries are exercises in applied ethics that use our tools of argument and extended reasoning. Of course, apart from the ethical principles forming the moral premises in various arguments constituting an enquiry, there will be premises of other types; in particular, such premises will form the ‘connective tissue’ between such underlying moral principles and the judgements made in their application. Moreover, some enquiries will focus on descriptive or conceptual issues as little sub-problems in their own right at times. Nonetheless, the *main* underlying issues to get sorted out are whatever the moral values are that you are applying to some professional ethical issue of concern to you.

In this chapter, I wish to focus on those moral values, or ethical principles, in their own right. The chapter will have two main sections: one of these sketches some elements of what is normally called ‘*Normative Ethical Theory*’ and the other introduces you to some of the issues in what philosophers call ‘*Meta-ethical Theory*’. I’ll explain these labels in due course. For now, I just wish to emphasize that the issues involved in each area are quite complicated and their treatment in, say, a philosophy major in a BA programme would be more sophisticated than I have space for. None of what I say below is original and if you wish other slants on it or to pursue matters further, then enrolling in such a programme (or just in particular moral philosophy units if your college/university permits that) is recommended. There are also many good texts on ethics that are pitched at undergraduates and I would recommend browsing the bookshop on your campus and seeking your tutor’s or instructor’s guidance.

Although limited in their treatment, I judge the ideas and issues that I portray in this chapter to be worth you putting in the effort to wrestle with in a serious way. Your professional activity is governed by your moral values and the more that you understand what they are and what confidence it is reasonable for you to have in them, the better.

Three Types of Ethical Theory

There are several sorts of enquiry that one might engage in concerning moral values. They are too frequently muddled together and, to avoid such confusion, a preliminary task will be to distinguish them clearly. I wish to distinguish three types: Descriptive Ethics, Normative Ethics and Meta-ethics.

Descriptive Ethics

First, one might wish to know what moral values are held by the members (or a majority of them) of the society of which one is part (or by some individual or group of individuals). Note also that one might find that this society mostly holds some value that some other society tends to reject. Or it might be that, although most members of a society tend to endorse some value, some individuals within that society reject it. Moreover, some individuals might reject the values that are held in their *own* society but find themselves in agreement with those of another society, past or present. Answers to such questions would be descriptive propositions, ones *about* the values of others – as opposed to ones advancing any moral value stance themselves.

A feature of these sorts of descriptive propositions about moral values is that there is no inherent problem in noting that variation exists among humans and across (and within) the societies they form. If the task is to *describe* those values, then one simply does that. Saying that various people do or don't hold this or that value is *neutral* as to whether you should agree with them or not (a point we will return to below).

I say 'neutral' because finding out what ethical views are held by various people, including societies or sub-societies, is rather like finding out other facts about them – like their income, taste in clothes, social status and so forth. They are probably to be seen as lying within the domain of the social sciences and presumably could be researched by use of survey instruments or whatever. My point here is that *descriptive* ethics, being solely concerned with what *is* the case, is silent as to what *should be* the case. Describing what some person's or group's moral values *are* is not the same thing as saying what moral values they *should* live their lives by. You might recall that, very early in the piece, I made a point of distinguishing propositions of this descriptive sort (like: 'Most Australians *think that* it isn't wrong to cheat on an income tax return') from moral propositions (like: 'Cheating on an income tax return isn't wrong'), pointing out that the simple presence of 'wrong' (even in its moral use) did not signal a moral proposition to be present. So, our first sort of enquiry (*about* the values held by various moral agents) is a social sciences descriptive style of enquiry. Having outlined it, and being at pains to distinguish it from what follows, I will pay it no more attention.

Normative Ethics

Although *descriptive* ethics is not addressing the question of the ethical values we *should* live our lives by, that is just what *normative* ethics does.

I have said earlier that professional ethical questions are ones best conceived of as lying within the domain of applied ethics. As that name suggests, what one is doing in applied ethics is *applying* some ethical values to the issue at hand. But *what* ethical values should one be applying? *Normative* ethics is that area of moral philosophy that tries to answer that question.

Although not every approach to normative ethics is of this sort, the approach that I have taken in this book is to encourage what I will call ‘principled judgements’ concerning professional ethical issues. I have sought to have you support your judgements by quite elaborate and explicit appeal to underlying moral principles that you hold (and which will form various moral premises) yet to realize that those underlying moral principles might conflict, even within one person, and, as a result of exploration of that conflict, some might be revised. One of the central tasks of normative ethics is to ask what, at the deepest level, such principles should be and, if your moral judgements are to be ultimately warranted by appeal to a matrix of moral premises, what can we generally say about the sort of thing that they might offer?

To date, I have concentrated upon the task of tracking the complexities of the connections between professional ethical issues and the moral principles that inform views upon them and upon the task of trying to expose and sort out clashes among your principles at that deeper level. Below, in this chapter’s section on normative ethics, I will discuss some views as to what those principles themselves might look like.

Meta-ethics

If *descriptive* ethics concerns itself with describing what various individuals or groups hold to be good or bad, right or wrong and *normative* ethics asks what sorts of moral principles we should live our lives by (and, for us, use as the basis of various judgements on professional ethical issues), what is *meta-ethics*? I will answer this in a roundabout way.

I have tried to help you to make judgements on particular professional ethical issues by appeal to some more general principles (in role as deep value premises). As you would realize by now, all of this is very difficult to do well but let’s say that you did manage to do that and that your views on some issue – say, the propriety of having any compulsory, or core, curriculum imposed on students by force (if need be) – were worked out to your satisfaction. Let’s assume that you came down against such a core curriculum and defended a curriculum that was entirely voluntary on the part of the student. A colleague disagrees and argues for the 4 ‘R’s (reading, ’riting, ’rithmetic and reasoning) as things that every student should, if capable, be *made to* master to some (explicitly identified) level of competency.

Both of you being, by now, trained minds, you sit down and mutually explore your differences. It becomes apparent that there is no factual dispute between you. You agree as to what the world is like so your disagreement at the level of curriculum cannot be traced to a disagreement about the truth status of some descriptive premise somewhere in someone's supporting argumentation. You meticulously check the logic of each other's reasoning and can find no flaw and you seem to have the same understanding of various concepts and their connections. So, why do you disagree about the curriculum? Well, given the above, there is only one possibility left – disagreement at the level of moral premises (and, ultimately, very deep ones at that – see Chapter 7 on such disagreements). It might be a fairly straightforward dispute at that level (he believes in each individual's duty to be useful in serving the group's purposes; you don't; you believe in the importance of individual control over what one is; he doesn't). Or it might be some more complex form of dispute concerning priorities, or issues of degree, or whatnot (again, see Chapter 7). But, whatever it is, you and he might simply be in irreconcilable deep moral disagreement. And note that, despite such disagreement with another moral agent, each of you is satisfied in your own mind as to the satisfactoriness of your stance.

Where to next?

The dispute is *not* rationally solvable. You are each reasonable, you each agree on the facts (so it is not as if further scientific research will help). If the topic is one of but minor importance, you might 'agree to disagree'; if it is a major issue, you might, in some sense, fight. Is there any sense, though, in thinking that, even if you and he can't manage to resolve your dispute and even if you are both internally satisfied with your own stance, one of you is just *wrong*; that is, in some sense of the word, a moral *error* has been committed? Could one of you be a moral 'flat-earthier', simply making a mistake – much as a flat-earth theorist, however sincere and self-satisfied, is simply making a mistake? And, if so, how could we tell *who* was in error? Or is ethics not like that and moral propositions not able to be thought of as being true or false?

The *third type* of ethical theory, *Meta-ethics*, as it is called, addresses questions such as these (and others of a related sort). Meta-ethics is a major focus of this last chapter. We have met the prefix 'meta' before when we spoke of metacognitive reviews and deliberation. As before, the idea is that you are standing back from particular ethical judgements and principles rather than engaging in the making or crafting of them (much as earlier we were standing back from the stream of substantive arguments rather than engaging in crafting them).

Normative Ethics – More Detail

Introduction

I said above that, having noted the domain of descriptive ethics in order to demarcate it from, in particular, normative ethics, I would cease to discuss it. Our normative ethical interest is not a psycho-sociological one concerning what moral values various individuals and groups *do*, in fact, hold but a moral one concerning what values one *should* hold.

Let's begin by reminding ourselves that, when addressing a professional ethical issue, I have suggested that you work 'bottom-up' in a carefully considered way from some tentative intuitive stance on the issue at hand in order to expose the principles which bear upon that stance. In one example that we used at some length, the issue was the propriety of nurses ever lying to their patients and our tentative stance in the initial argument was that they sometimes should. Why? – because so doing was demanded by a deeper commitment, one to patient welfare.

The point about this process of defence of the endorsement of nurses sometimes lying to patients was that it was done by trying to connect the issue at hand to *another* one. As noted, I will call this 'principled moral decision-making'. In trying to make a principled moral decision, one tries to bring moral principles to bear on the issue at hand. In this case, the initial element in this was to bring to bear the patient-welfare principle. The *connection* of this principle to the tentative stance adopted on the issue was by locating each in an argument. So, an MP outlined the principle appealed to in support of a (tentative) stance outlined in the MC (with the connection usually involving some DP or other). Of course, professional ethical matters are rarely neat enough to be settled by appeal to just one such exercise of appealing to an MP. Mostly, you will find a number of principles that you have some sympathy with that are all bearing upon the issue at hand. Sometimes the connection is direct and sometimes more indirect (via defences and criticisms). And, as we have seen, you will likely find that some of the principles that occur as moral premises in the various arguments that unfold in your enquiry *oppose* each other.

So far, so familiar I hope. *Principled* moral decision-making involves bringing such moral principles to bear, applying them to the issue at hand. That process can be long and involved if done thoroughly. And, of course, sometimes one hasn't the time to do it or the issue is too trivial to bother with sophisticated analysis anyway. However, important matters are worth thorough enquiry, if time allows.

The job of the foregoing chapters has been to explore the complexities of thinking an issue through in some depth and, given that this involves conflicting principles, trying to sort out some priorities such that one's set of moral principles is well enough organized to apply to the problem at hand.

In this section, I wish to talk a little more about the principles themselves. I will put it to you that, when crafting arguments that bring various moral principles to bear through an enquiry, two distinct sorts of exercise might be going on.

One is that you tend to apply principles of a type that is sometimes called ‘deontological’ (from the Greek ‘deon’, meaning binding duty). The other is that you tend to argue as a consequentialist (as in ‘consequences’). I’ll explain these to an extent suitable for the book’s purposes in a moment and then close with a suggestion as to what, as a professional facing ethical issues (as opposed to philosopher of ethics), you could profitably make of these theories.

Consequentialist Ethics

Say that the ethical issue facing you arose when someone is enrolled in a degree course that leads to being certified as qualified to practise some profession. There is an important assignment due but, for whatever reason, she has not managed to prepare adequately for it. Being a student of marginal competence, with a poor academic record, she knows that if she fails the assignment, then she is in grave danger of having her enrolment in the course cancelled. A friend in the same course offers to help with the assignment. There are two main types of principled reasoning that you might engage in when trying to work out whether you morally approve of this or not.

The type that I wish to explore in this section involves arguing in a way that is concerned with what would *result from* accepting such assistance. The act is appraised as right, or wrong, good, or bad, by focusing on the consequences of the action. Hence, obviously, the label ‘consequentialism’ (another common label for this broad type is ‘utilitarianism’).

Consider the cheating example. One consequence perhaps is that such a cheat would get a passing grade instead of fail. A related consequence (in our earlier scenario) is that the cheat will be able to continue enrolment in the course, thus saving time and money. Another, more indirect, consequence is that prospective employers will have the false belief that the cheat is competent as certified by her university. Even more indirect as a consequence is that future clients might have incompetent provision of professional services by the cheat. Another type of consequence, dependent upon how well known it is that cheating occurred, is a decline in the reputation of the university’s degrees. And, if known to occur yet be unpunished, another will be a rise in the number of undergraduate cheats. And so on.

As is obvious, a given act has all sorts of consequences, both direct and indirect, and if they were appealed to in ethically judging the rightness or wrongness of the action, then pointing out those consequences would form descriptive premises in the arguments (what we call ‘means/end’ arguments) deployed in coming to a judgement. However, mere descriptive *noting* of consequences is hardly enough to drive a judgement, one needs to know whether the consequence portrayed is a *good* one or a *bad* one. And, as should be familiar to you by now, whatever criteria you were appealing to in order to rate some consequence as good or bad would be put in as the MPs of various means/end arguments. So, in illustration, one might have:

MP Jane should be able to continue enrolment in her law course.

DP Unless Jane cheats on her assignment for 'legal ethics' in her law course, she won't be able to continue enrolment in that course.

So,

MC Jane should cheat on her assignment for 'legal ethics' in her law course.

Recall, though, that you were also introduced to the business of premise *defence*. Say we carried that out for the above MP in response to the query: 'What is so important about Jane's continuing enrolment?'

Whatever the argument offered in its defence, there will have to be some further moral premise appealed to. No matter what this might be, that further moral premise might itself face the same sort of query demanding that it be defended. So, say that, in advancing a case for Jane's continuing enrolment, appeal is made to the importance of her being available for professional employment as soon as possible. And we could ask why *that* is so important. (We touched on such matters in Chapter 7.)

This begins to look like an endless business. No matter what is offered as a deeper value in defence at any given point, it seems that we could keep probing for the yet deeper story, for the deeper value underlying the one at hand. Fortunately, it is not like that and this sort of chain of defences ends. Indeed, it has to end, on pain of what logicians call: 'a vicious infinite regress'. It is impossible, in a real moral value system, one held by real and finite moral agents, to have such an infinite chain of values. What happens instead is that one gets down to one or more fundamental values, ones that act as a sort of moral bedrock, as the ultimate ends in our means/end chains of moral justification.

As to what such final values might be and how many of them there might be, disagreement reigns among consequentialists. Two popular ones are these:

'Always act so as to maximize the greatest sum total of human happiness'; and

'Always act so as to maximize the greatest sum total of human desire satisfaction'.

A few things to note: The first is that these are *not* the same end – desire satisfaction is *not* automatically connected with being happy, a lot depends upon the desire; and happiness might occur without it constituting the satisfaction of any preceding desire.

Another point is that there are variations within these broad concerns for human happiness or desire satisfaction. So, instead of largest *sum total* happiness (which might be achieved by quite uneven distribution of that happiness), one might favour the most *even spread* of happiness. One might also be concerned with the happiness of species other than humans and trade that off against human happiness (Peter Singer's name looms large here). Or, instead of the greatest amount of desire satisfaction, one might similarly want some sort of *evenness* of spread of desire satisfaction. Further clarifications and more sophisticated and

complicated variations have been offered on these two themes but it is beyond the scope of this book to pursue matters further.

Also, ends other than some variation of the ‘happiness’ or ‘desire satisfaction’ ones have been advanced. So, one might be concerned to maximize true belief, or freedom of thought, or religious conviction, or Again, pursuing such possibilities in detail is beyond us and you should consider reading further in the philosophical literature.

Finally, if one has more than one such fundamental principle (concerning the ultimate ends that one will be wanting to be served by human action) then there is no guarantee that they won’t clash. Indeed, there is a guarantee that they will. We touched on this in Chapter 7 and I will return to it in a section below.

Whichever ultimate end(s) is(are) chosen, the idea is that this is where justification runs out. Such values are what are ultimately appealed to in warranting other moral judgements and values. They form the deepest MPs of a web of means/ end argumentation.

Deontological Values

Consider again our cheating example. Instead of appealing to what will *result from* the action, to its *consequences*, one might just look to the action itself, to the *sort* of action it is, in and of itself. So, try this argument.

Jane’s action is wrong because it is an act of dishonesty and acting dishonestly is always wrong.

Upon clarification being sought from the speaker by asking: ‘What’s so wrong about it? – after all, sometimes no harm comes of dishonesty’, we might be told: ‘Whether it does any harm or not is not the point; acting dishonestly is just the wrong sort of thing for someone to do’.

This rationale for condemnation of the cheating is *quite different* to our previous consequentialist ones. In those arguments, the DP claimed that some consequence, or outcome, of the cheating would occur and the MP stated some moral principle that covered that sort of consequence. By reference to that principle, the consequence was rated morally as a good or bad one. In effect, the action of cheating would be, say, morally condemned by forming an appraisal of what the consequences would be and morally judging those *consequences*. The act of cheating is, in a sense, judged *indirectly* by way of judging its consequences.

In contrast, someone inclined to think about ethical issues in a *deontological* way focuses *directly* on the act itself and morally judges that act by categorizing it as an instance of this or that type – where the types, or classes, of actions are ones covered by moral principles. So, in this case, the act is condemned because it is *an instance of the type*: dishonesty.

Of course, just as the consequentialist might become conflicted by an action having both good and bad consequences and have to sort out some priorities,

a deontological theorist might become conflicted by an action being able to be categorized in more than one way, some good, some bad. So, consider the person helping Jane. She might be acting dishonestly by doing Jane's assignment (a bad thing, *prima facie*, we will assume) but she is also satisfying a friend's request (a good thing, say). Further, much as the consequentialist can track defences of MPs in a chain of means/end arguments that ultimately end in something at a fundamental level like, say, some version of the 'greatest happiness' principle, so, too, the deontologist might have a chain of arguments defending some initial MP. The *style* of that chain is quite different, though.

Keep in mind that deontologists are *not* interested in what *results from* an action, just in what *sort* of action it is, in and of itself. Thus, such a defensive chain of reasoning will *not* contain means/end arguments. Rather, our other main type, set relationship arguments, will be deployed. A common deontological pattern in defence of a particular act's rightness (or wrongness) is to locate the act as an instance of a class of actions upon which we have some morally principled position. So, as an example, say that Sarah, whose profession is that of politician, is contemplating how she should vote concerning a proposed war. The war in question concerns a country (Eastland, say) with which the politician's country (Northland, say) has a formal treaty. Eastland has formally sought assistance in accordance with the clauses of the treaty. Although no doubt unfamiliar with normative ethical theory, the politician counts as a deontologist in her manner of thinking about ethical issues and, in effect, thinks as follows:

MP1 All treaties should always be honoured.

DP1 Our treaty with Eastland stipulates that if Eastland requests it, Northland should declare war on any country Eastland is at war with.

DP2 Eastland is at war with Westland.

DP3 Eastland has requested that Northland declare war on Westland.

So,

MC1 Northland should declare war on Westland.

A colleague of Sarah's, Tom, who has a consequentialist bent, suggests that going to war against Westland will have all sorts of bad consequences and Northland should seek some way of wriggling out of its treaty obligations. He challenges Sarah to justify her view that Northland should declare war on Westland. Sarah replies with a fairly feral and abbreviated version of the above structure: 'We have to; we signed a treaty with Eastland and we can't break that'. Tom responds: 'Why shouldn't treaties sometimes be broken?'

In our terms, Tom is seeking from Sarah a defence of her MP1 claim that all treaties should always be honoured (although she may not have explicitly formulated the principle like that to herself – remember all of the problems that we had in Chapter 3 concerning the teasing out of missing moral premises).

Say that Sarah responded to Tom's challenge *cum* request as follows: 'They shouldn't be broken because they are a sort of promise'. Laid out as a structure this feral defence amounts to this argument:

MP2 All promises should always be kept.

CP1 Not honouring a treaty is always an instance of not keeping a promise.

So,

MP1 All treaties should always be honoured.

In effect, when combined with the original structure that we attributed to Sarah, we now have a two-link chain of reasoning. Put in a 'why? – because' way, it goes (somewhat abbreviated) like this: why should we go to war? – because we should honour our treaties; why should we honour our treaties? – because all promises should be kept. Going to war is warranted by noting that such an action is an instance of the type: 'honouring treaties'. And the whole type (or set) 'honouring treaties' is warranted by noting that it is a subset of the set: 'keeping promises'.

We have generally spoken of defences of premises as constituting a *deepening* of a given argument but note that, in the case of such deontological defence exercises (using set relationship arguments), it is better thought of as *widening*, rather than deepening, Sarah's case. The warrant proceeds by appeal to more general principles.

As with consequentialism, this process of MP defence can't go on forever and eventually appeal is made to (usually very general) principles that are 'bedrock' or held to be capturing intrinsic values. These principles specify types of action that moral agents should perform and which cannot be subsumed under some yet broader classification.

Also, as with the consequentialist, it is quite possible that various values that lay down *prima facie* duties (even if they are at this fundamental, or bedrock, level) will conflict with each other in various scenario situations. For instance, one might, as with the case of Sarah above, hold (*prima facie*) that one should always keep one's promises yet also hold that one should always be truthful. Clearly, sometimes keeping one's promise would involve lying. (For instance, one might promise a dying parent that one will keep their daughter from harm and this might only be achievable in some scenario by lying to a drug-crazed and violent house intruder about her presence.)

We will revisit the matter of value clashes below (although, as noted, it was addressed in Chapter 7).

Just to re-emphasize the distinction using the 'cheating' case: that particular deontological judgement that it was wrong might be made even if it is clear that *none* of the consequences that would disturb a consequentialist actually obtain. A consequentialist, on the other hand, would see nothing wrong with an episode of cheating *provided that* no bad consequences ensued from it.

Most people have a tendency to incline towards one or other of these broad theoretical orientations when approaching ethical issues and problems in a

principled way, one that rests their judgements upon *principles* of some sort – but not everyone does. I turn to such hybrid normative ethical views next.

Hybrid Normative Ethical Views

As just mentioned, most people incline fairly much towards one or another of these two broad types of normative ethical principle. However, many people, especially when first trying to explicitly understand their approach to principled ethical decision-making, find themselves thinking both as a consequentialist and as a deontologist, even on the one issue. For instance, in our ‘lying nurse’ example, we had one argument in which the lying was supported by appealing to a good *consequence* of the lying, namely the increased health of the patient who was lied to. In a counter-argument against this argument’s ‘patient-welfare’ MP, appeal was made to treating patients with respect. We didn’t bother to defend this ‘respect patients’ CMP but, had we done so, I doubt that the story would have involved advancing a means/end argument pointing out the good consequences that would flow from respecting patients. Rather, it seems likely that the construal of the ‘respecting patients’ value would be *deontological*, such respecting would be just seen as the right sort of way to act and were it to be defended, then I would imagine that it would be by appeal to a broader commitment to respect for people generally.

So conceived of, and thinking of the enquiry to that point as one in which the counter-argument has been mounted by the author as an exercise in probing self-criticism, we seem to have a moral agent who is a normative hybrid of consequentialist and deontologist. He is inclined to want to act in a respectful way as just the right way to act in and of itself (and not with some good consequences in the back of his mind as the justification for so acting) but concerned that (sometimes) doing that will have consequences that might be bad enough to warrant disrespectful treatment upon those occasions.

So, not just a moral dilemma, or conflict, but one that involves elements of each of our normative theoretical orientations.

Although just illustrated with a value *conflict*, hybrid thinking can also occur in other ways. One might, for instance, just think, deontologically, that one has a duty to act honestly. Although that commitment is not dependent upon any consideration of consequences, one might *also* favour it on the basis of its good consequences. Consider our ‘cheating student’ scenario. Not only might one decide not to cheat because that counts as a case of acting honestly, one might *also* support it with the consequentialist point that, if one did cheat, then that would run a high risk of lowering confidence in the university’s grade integrity which, in turn, would

The above portrayal of deontological and consequentialist views and their relationship to our processes of principled argumentation is inevitably sketchy but I hope that it gives you some feel for two major approaches to warranting stances on professional ethical matters or, if you like, major ways of thinking about the sorts of fundamental moral principles that you would be appealing to, and applying,

to reach some such stance. As I said earlier, there are many good undergraduate texts on normative ethics and if you wish to pursue this level of theory further, then that is where to turn. For now, if you understand the two approaches, I would just suggest that you reflect on the ways that you think about ethical problems and see which is the approach that you are most comfortable about.

Diagrammatically put, what we have so far is:

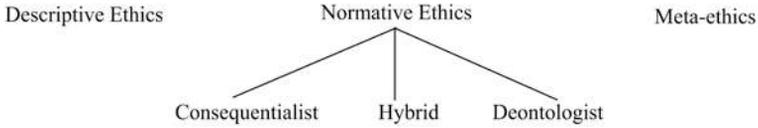


Diagram 9

Meta-ethics – More Detail

Introduction

As noted earlier, we have met the prefix ‘meta’ before in ‘metacognitive’. It is from Latin and means ‘above, or beyond, or after’. With ‘metacognitive’ the idea was that you mentally stood apart from your thinking (or attempts at knowing your values, hence the ‘cognitive’ which is meaning-connected to knowing) and thought about your thinking. (In that case, the particular sort of metacognition involved was planning your enquiry’s direction in the light of progress to date.)

Meta-ethics is thinking about ethics, in particular, thinking about the nature of moral values and value judgements. Humans engage in all sorts of intellectual, or quasi-intellectual, enterprises – science, art, religion and so on – and we’ll be trying to work out what is distinctive about one of them: ethics. These issues are hugely controversial in philosophy and I don’t expect you to end up with a thoroughly worked out meta-ethical view. I certainly won’t be giving answers, just challenges, questions and puzzles. What I hope will occur, however, are *three things*.

First, that one or other position of the spread of views that I am about to introduce to you will capture, or articulate for you, your existing meta-ethical tendencies. *Second*, that the difficulties facing your preferred view will give you pause for thought. And *third*, that you will see the enormous problems caused by these complexities for moral agents in dispute with other moral agents.

In the remainder of the chapter, I’ll be introducing you to some broad sorts of meta-ethical theory and to some of the argumentation concerning them. I don’t have space to do the job thoroughly (that would be a text in itself) but, as with normative ethics, there are loads of good introductory texts around. If you do read further, then be warned that the literature is rather inconsistent when it comes to terminology and the labels I use for the various theories might not be consistent with some other texts. Also, contrary to my focus on argumentative depth in the foregoing, I will just present the initial part of what would be a long enquiry as the

merits of various arguments got teased out. I intend merely to provoke thought and alert you to some issues.

The Descriptive/Moral Distinction Revisited – Moral Objectivism and Moral Subjectivism

Almost the first thing we did (way back in Chapter 2) was to outline the distinction between descriptive propositions and moral propositions. Yet some meta-ethical theorists would call the distinction spurious; for them, moral propositions are just a variety of descriptive propositions. How so? Read on. The view is controversial and, even if it is right, the distinction was still a useful one for our purposes at that time even if it turns out to be too crude. It is such a, if you like, *descriptive* construal of moral propositions that I wish to look at first. The meta-ethical theory in question is *Moral Objectivism* (in contrast to moral *subjectivism*, which we will come to in due course).

Moral Objectivism The name ‘moral *objectivism*’ is appropriate in that such theorists contend that the correct understanding of what is being said when someone issues a moral judgement is that a claim is being made about the *objectively present moral facts* of the case. We are, then, to understand moral propositions to be in much the same line of business as ordinary descriptive ones; each is trying to tell us the facts about what the world is like – it is just that a *moral* judgement is aimed at bringing to our notice a *particular sort* of fact. So, much as: ‘Grass is usually green’ asserts that a certain property, greenness, is typically found in a certain sort of thing, grass, so: ‘Stealing is usually wrong’ tells us that a certain property, wrongness, is typically found in a certain type of event, or action, stealing. According to the objectivist, each purports to tell us something about *the way the world objectively is* and each is true or false depending upon the factual accuracy of the proposition. In development of the point, consider the following as illustrations.

P1 The earth is round.

P2 The earth is flat.

Whether P1 or P2 (or neither) is true is a matter, not of our whim, or fancy, or of current scientific fashion, but of the cold, hard facts of reality. P2 was at one time the dominant view among theorists but we would take them to have had *false* beliefs, to have been simply misguided as to what the facts *really* were. P1, on the other hand, we take to correspond to reality, to the objectively existing facts of the universe. Now, consider:

P3 Stealing is always wrong.

P4 Stealing is the right thing to do if it is the only way to stay alive but wrong otherwise.

Although P3 and P4 differ in their judgement in only one circumstance – when one thives to maintain one’s life, they do differ. Which view is morally superior? According to the *objectivist*, one (or both – maybe stealing is *always* right or right in other circumstances) of these propositions is simply *false*. That is, in an *exactly* analogous way to the dispute between P1 and P2, at least one of P3 or P4 has just got its *facts* wrong, simply fails to describe the world as it really is, as a matter of objective fact.

Subjectivism Contrasted with this is the view of the moral *subjectivist*. You have probably heard the expression: ‘Beauty is in the eye of the beholder’.

Contrast these two propositions:

P5 That painting is beautiful.

P6 That painting is ugly.

We would not usually think that the dispute between these two propositions is one about what the facts are, about what the painting is *really* like. Rather, we allow that such judgements are a personal, or subjective, thing. It is not that the painting is, as a matter of fact, in itself, beautiful; it is just that the speaker of P5 favours it, responds in a certain way to it. And, even if we agree with P5, it is not as if the speaker of P6, like that of P2, is in *error*; it is just that she responds in a different way to the painting.

Crudely put, much as the meta-ethical *objectivist* aligns P3 and P4 with the scientific or descriptive P1 and P2, the meta-ethical *subjectivist* aligns them with P5 and P6. Rightness and wrongness, like beauty and ugliness, are held to be ‘in the eye of the beholder’.

So, in contrast to the *objectivist*, the claim here is that to have a moral principle is *not* to have a theory about the moral property present in a certain class of situation. It is just to have a certain *attitude about* such situations. And, to make a particular moral judgement is *not* to claim that a particular (moral) factual state of affairs obtains but merely to express an attitude, or stance, about something.

For the moral *subjectivist*, someone’s proposition that something is wrong means no more than that that person is opposed to it happening (and similar remarks for other pieces of moral language could be made). This is actually a bit too ‘broad brush’ and there are sub-varieties of *subjectivist*, but, for the moment, we’ll ignore such sub-varieties and focus on distinguishing generic *subjectivism* from *objectivism*.

So, in summary, for *objectivism* there is some sort of moral fact of the matter, whereas for *subjectivism* there is nothing more going on in morality but the preferences of moral agents. (Not that *objectivists* don’t prefer ‘the good’ as well, it is just that they think that there is *more* to morality than that, whereas *subjectivists* don’t.)

Our first major meta-ethical distinction, then, is that between meta-ethical *objectivism* and meta-ethical *subjectivism*. Put diagrammatically, so far we have:

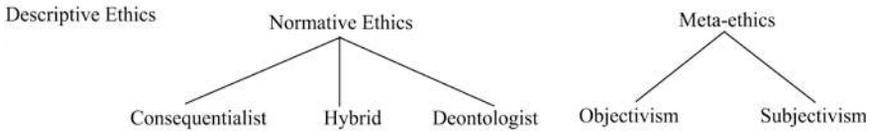


Diagram 10

Some Implications of These Meta-ethical Theories

I trust you can see how this sort of theoretical dispute affects the way one thinks about professional ethical issues. Consider, for instance, a central candidate aim for schooling: ‘producing good citizens’. Let’s say that what counts as a ‘good citizen’ gets clarified, in part, as: one who accepts the values of society. But say also that to achieve such an aim seems to involve some sort of (perhaps subtle) indoctrination in which students’ (positive) freedom of thought, at least concerning some moral values, is interfered with. Let’s assume that you have wrestled with the clash between the ‘freedom of thought’ value and the ‘shared societal values’ value and sorted out your priorities in favour of the latter. Now consider our two meta-ethical theories. A glance back at our two views, subjectivism and objectivism, should indicate that the view of the meta-ethical *objectivist* is the one that seems, on the face of it, to be the most favourable to providing some sort of additional legitimacy to the ‘value-indoctrination’ view. Why is this? Well, consider our astronomy examples of a while ago. You might feel on solid ground if you teach a child that the earth is round and not flat (and justified in your attempts to revise any tendency for him to believe in the latter) just because you take former proposition to be *true*. That is, as a defence against the criticism: ‘Why don’t you let students have the freedom to believe whatever view about the shape of the earth that they wish?’, you might feel it adequate to reply: ‘But that would mean letting some of them believe what is false’. So, in a parallel way, if you felt that the values on the basis of which you were intervening were *objectively true* values, you might feel that that was all the justification you need.

If, on the other hand, you felt that moral matters were essentially *subjective*, that goodness and badness and so forth were ultimately no more than a matter of taste, or feeling, or whatever, then you might feel much more uneasy about instilling a set of moral values in others, or acting on their basis to restrict another person’s freedom.

Consider our ‘undergraduate cheat’ example. Say that Joan, a friend, hears about it and challenges Jane (the cheat): ‘You shouldn’t have done that, cheating is wrong’. Jane responds that it isn’t wrong so long as no one in authority finds out. According to the meta-ethical *objectivist*, either Jane or Joan has a false belief, is, so to speak, a sort of ‘moral flat-earthier’. According to the *subjectivist*, there are no truths to be had in the moral domain and they have no more than different stances concerning the activity of cheating. If Jane (or Joan) see their moral views in an objectivist way, then it places a different interpretation upon how the nature

of their disagreement is construed than if a subjectivist theory of ethics were to be believed instead. Either way, an exploratory dialogue between them could occur but it would be conceived of differently depending upon which meta-ethical theory was accepted much as an exploratory dialogue between two people about the beauty of a painting is of a somewhat different kind to an exploratory dialogue between two people as to the shape of the earth.

So much (for now anyway) concerning why we should be interested in such a meta-ethical dispute. What about the meta-ethical dispute itself, which theory is correct – subjectivism or objectivism? As you will realize by now, I have no intention of trying to supply an answer. Meta-ethicists are by no means agreed on this matter. However, as is the case with so many debates within philosophy, to merely note that an issue is not fully resolved is not an adequate ground for avoiding thinking about it. The issues are usually too central and important to our ways of thinking about the world and our place in it to be ignored or left in the background of your mind, unnoticed. Such is the case here. Moral principles and value judgements don't just rule our professional lives, they intrude into every part of life. So the task is to think about things as best you are capable of, and have time for, and come to as intellectually satisfying a decision as possible. What I will be doing is presenting you with a range of arguments (concerning the subjectivism/objectivism controversy) for you to think about and begin to assess. In the interests of brevity, the following will be just the tip of an iceberg and, as noted before, there is a considerable philosophical literature on this and many introductory texts covering meta-ethics could be accessed to continue what I begin below.

Objectivism – More Detail

Naturalism Outlined

Let us start with objectivism. Before we can proceed much further you should realize that there are, broadly speaking, two sub-varieties of objectivism. Both are objectivist because both hold that goodness, badness etc. are matters of fact (and claims about them are thus to be properly understood as descriptive propositions, ones attempting to describe some aspect of reality). Where they differ is in their story as to *just what sort of fact* it is that the moral opposition is about.

The first variety, *naturalism*, holds that 'goodness' (and so on, I will limit the number of moral terms of I refer to from now on but a similar sort of analysis transfers across to each of them) is just a shorthand label for referring to other, quite ordinary, *natural* features of the world. That is, that goodness is, by definition, identical to some ordinary sort of property and is *not* some sort of funny special property of its own sort. (As we'll soon see, the other sub-variety of objectivism thinks that goodness *is* its own special sort of property and is sometimes called: 'non-naturalistic objectivism'.)

Even then, there is further disagreement among moral naturalists as to just *which* sort of ordinary property constitutes goodness. Let me illustrate.

Some naturalists would claim that, say, to judge some action as good is just to say of it that it is conducive to the survival of the human race. Because it doesn't matter what particular action, or even type of action, we have in mind here let's just use the all-purpose tag 'X'. So this variety of naturalism can be displayed as follows:

'X is good' *means* 'X is conducive to the survival of the human race'.

There are other varieties of naturalism; two common ones are:

'X is good' *means* 'X contributes to the greatest happiness of the greatest number of humans'.

and:

'X is good' *means* 'Most people in my society approve of X'.

Some comments: First, note that each of these is a naturalistic theory, the sort of thing mentioned in the right-hand side of each definitional analysis is an ordinary, albeit complex, sort of fact (biological or psycho/sociological). On each view, in order to find out what particular sorts of action actually are right or wrong, good or bad, one might appropriately consult various scientific experts.

For example, social scientists might lend guidance as to whether X actually was approved of by most people in society or not. If, say, most of them did approve of tax avoidance, then, for the sort of naturalist who, in the above way, *meaning-equates* rightness with societal approval, tax avoidance is thereby automatically, as a matter of objective fact, right. Though not as straightforward, the situation is much like saying that, if someone is identified as an unmarried male adult human then, just in virtue of the *meaning-equation* of 'being a bachelor' and 'being an unmarried male adult human', that person is thus automatically, as a matter of objective fact, a bachelor. To learn that Bartholomew is a bachelor just *is* to learn that he is an unmarried male adult human. We have the *same fact* about him, just two meaning-equated labels for it. Similarly, for this variety of naturalist, to learn that tax-avoidance is approved of by society is just the same thing as learning that it is right. The latter is *not an extra* fact about tax avoidance, it is the same descriptive proposition presented again using a different label.

As mentioned earlier, one implication of this would be that two of our propositional types from Chapter 2, descriptive and moral, would collapse into one: descriptive. On this naturalistic view, there is no distinction to be made between those types. However, although each of the three rival analyses outlined is naturalistic, they are each defining moral terms differently when it comes to the

detail of the analysis (and there are more suggestions in the literature than just these three).

My *second* point then is that one should note that each is a claimed analysis of the *meaning* of (in this case) ‘good’; (similar things would be said about ‘right’, ... – there is a whole related family of moral terms and it doesn’t much matter which one we choose to focus on). So, each of these analyses itself constitutes what we were earlier calling ‘a *conceptual* proposition’ to the effect that such and such means so and so. Each hopes to be correctly reporting just what is *meant* by words like ‘good’ and it seems somewhat strange that they are so different – on the face of it, we seem to all have much the same meaning for these words even if, as a matter of normative ethical principle, we disagree about what sorts of things are right and what wrong. This is much like you and I sharing an understanding of ‘bachelor’ yet disagreeing as to whether Joshua is or is not one. As we will shortly see, this meaning focus gets these naturalistic theories into trouble.

Third, obviously enough, similar sorts of definition (but negatively put) could be given for ‘bad’, ‘wrong’ and so on.

Fourth, as you probably already realized, the last of the three analyses listed above (the one about society and its views) questions the distinction made earlier between what I called *descriptive* and *normative* ethics. It takes propositions about what *is* good and what society *asserts to be* good to be the same claims.

Fifth, despite superficial appearances, the second of the above analyses has to be carefully distinguished from the consequentialism (a *normative* ethical theory) of earlier in the chapter. I will return to this.

Diagrammatically, so far we have:

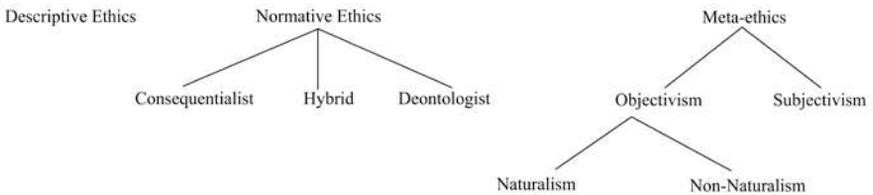


Diagram 11

This is the last diagram I will portray because, although we have more discussion of matters below, it gives the overall architecture of the spread of views covered in the chapter. Remaining to be discussed in more detail are *non-naturalistic* objectivism and *subjectivism*.

Standard Objections to Such Naturalistic Analyses

One or other of these three versions of naturalistic objectivism might seem to you to be appealing and roughly correct but keep in mind that these theories are supposed to be telling you what is *meant* by words like ‘good’. That is, they are supposed to

be presenting you with definitions reporting meanings in much the way that: ‘X is a bachelor’ *means* ‘X is an unmarried male adult human’ is presenting a definition reporting our conception of ‘bachelor’. As touched on above, you might think that the very fact that there are (so far listed) at least three rival definitions of radically different sorts is immediately rather suspicious; after all, you might think, there isn’t that much controversy about most definitions (as in ‘bachelor’ above). But some of the concepts in terms that we naively wield every day *are* obscure and ill-understood and controversy and disagreement may properly exist until things are pinned down and tightened up a bit. Moral language seems to be a classic case. (Certainly the issue won’t be settled by looking up a dictionary – see the sort of thing that we said in Chapter 8.) Still, the offerings are suspiciously *radically* different. Anyway, how can we decide on whether any, or none, of the above is a correct analysis of the nature of our moral concepts?

Basically, it’s a matter of seeing if any of them intuitively ‘fit’. That is, if, in response to any of them, you say: ‘Yes, that’s it, *that’s* what is meant by terms like “good”, “bad” and so on’, then you are accepting that conceptual analysis. However, we don’t have to leave things as unstructured as this. We can bring various features and implications of the suggested analyses into particular focus for our intuitions to bear upon. As an illustration, let’s discuss, say, the ‘human survival’ suggestion. The claim was, remember, that by ‘X is good’ is *meant* ‘X is conducive to human survival’. This implies that anyone who advances a substantive normative judgement to the effect that some act is good despite it involving the destruction of the human race is not just advancing an eccentric moral position, he’s talking literal *nonsense*. I mean this quite literally. That is, on this analysis, such a proposition would be held to be as literally incoherent as the proposition that Jones is a bachelor despite being married. But if we look at the proposition that some act is good despite it involving the destruction of the human race, it seems *not* to be incoherent. You might reject it but it doesn’t seem to be ‘married bachelor’-style nonsense. Yet, cashing out the meaning of ‘good’ according to the above analysis entails that it would be, because such a claim, when unpacked, would be saying that the act was conducive to human survival despite it involving the destruction of the human race! (Much as, when unpacked, ‘Boris is a married bachelor’ would be saying that he was a married, unmarried male adult human.) The proposition that wiping humanity out is a good thing might be an *odd* claim, and we would certainly like to hear how such a dramatically destructive act could be justified, but we *do understand* the proposition as coherent even if we end up deeming it to be in moral error. (Which, incidentally, you shouldn’t be too quick in saying. What if humans were waging an aggressive war against a peaceful and culturally and intellectually superior alien race and it was a matter of them or us being destroyed?)

So, what’s happened here is that we’ve drawn out an entailment of the proposed analysis and shown that it is counter-intuitive. The analysis would have it that a certain sort of moral view is *incomprehensible* yet our linguistic intuitions tell us that it isn’t. So, we’ve tentatively decided that that analysis is intuitively unsound.

It's *supposed* to be telling us what we mean by words like 'good' but, if you accept the foregoing, it doesn't.

There is more to be said on everything covered in this chapter, but it seems that that particular naturalistic theory of what one is saying of some act in calling it good is to be rejected as unsound. Let's look briefly at the others we have listed so far. Recall that, as *naturalistic* analyses, each claims to identify 'goodness' as just *meaning* the same as, or being shorthand for, some expression picking out some fairly ordinary natural quality. And each has an analogous implication to the above. That is, they would, respectively, classify the propositions:

'X is good but does *not* contribute to the greatest good of the greatest number of humans'; and

'X is good but is *not* approved of by most people in my society'.

as *un-understandable literal nonsense*.

Indeed the same (and, in my view, for what it's worth, rather crippling) objection seems to apply to any such attempt to automatically meaning-equate goodness with some *other* (and natural) property. For *any* such property it seems possible for there to be *coherent* disagreement about whether or not that sort of thing is good. One can coherently debate whether or not human survival, human happiness and so on are good things in a way that one can't coherently debate whether or not bachelors are married. Or so it seems; perhaps there are flaws in the above argument, although it is an argument that is long-standing in moral philosophy and generally well thought of.

A Possible Confusion I want to now (briefly) return to the distinction between the 'greatest human happiness' naturalistic *meta-ethical* analysis considered here and the *normative* consequentialist theory we met earlier. One version of the normative consequentialist principle might be: 'Acting so as to contribute to the greatest happiness of the greatest number of humans is good'. At first glance the normative and meta-ethical theories look very similar in that each connects goodness and human happiness. Appearances deceive however, for the *types* of connection differ.

Say someone disputed the *normative* 'greatest human happiness' view that we saw earlier; this might be an animal liberationist who was concerned about scientific research using animals, perhaps. She might say: 'Look, I accept that human happiness is important but not to the exclusion of that of other species. Sometimes maximizing human happiness is *not* good because we should trade off some human happiness to gain increased happiness for members of other species'. A normative dispute might then occur, one of the type that you have teased out as a formal dialogue between 'author and critic'. But each participant in the dialogue *understands* the other; they *share* meanings. (Indeed, such mutual understanding is a necessary condition for disagreement – that was why we earlier spent time in working definition style clarification to get all concerned onto the

same wavelength. The words ‘good’, ‘should’ and so forth are *common* linguistic property and the dispute is *not* about what ‘good’, or ‘should’ *means*. The *meaning* is shared and the dispute is about what actions fall into, say, the category: good. It is exactly parallel to you and I *sharing* an understanding of the meaning of ‘bachelor’ but disputing whether Boris is, or is not, a bachelor (‘He had a secret marriage a decade ago’ – ‘No he didn’t’).

But were one to have a *meta-ethical* naturalistic version of the ‘greatest human happiness’ view, one as listed in this chapter, then things are different. In that case, if some animal liberationist raised the very *same argument* then a long normative ethical critical dialogue would *not* ensue. Our meta-ethical naturalist would presumably respond with something more like: ‘Look it just *doesn’t make sense* to say that it’s *good to lessen human happiness* (in order to increase that of other species); you don’t understand the language properly, “good” just *means* “maximizing *human* happiness”. So that’s like saying that human happiness will be maximized by lessening it – which is, of course, incomprehensible nonsense’.

Such a naturalistic retort falls flat. It has none of the force of saying (in response to someone who’s remarked that Boris seems the most happily married bachelor she’s met) ‘Look it just *doesn’t make sense* to say that Boris is a married bachelor; you don’t understand the language properly, “bachelor” just *means* “unmarried (male adult human)”’. It is *not nonsensical* to say that human happiness should be lessened in favour of increasing that of other species. Whatever you think of animal liberationist theses, they are at least coherent in a way that married bachelor ones are not.

The difference is not that ‘bachelor’ is a trivial example with clear meaning whereas ‘good’ is more abstract and difficult, something where our intuitions are unclear. It is indeed more abstract and difficult and, in some aspects, our intuitions might not be clear but on the following, at least, our intuitions *are* clear: whatever else is a possible meta-ethical theory concerning ‘good’ (and the rest of the family of moral concepts) simply *equating* goodness with the maximizing of human happiness as *meaning* the same thing is, I suggest, intuitively to be rejected.

I have found in the past that many students have considerable initial sympathy with naturalistic meta-ethical theories. I think that this is because they don’t fully understand just what the implications of such theories are. In particular, I think many students confuse having some strongly held *normative* moral principle (like the consequentialist one discussed above) with having any such pet moral view ‘automatically true’ in virtue of the very meaning of the words (as per a naturalistic objectivist account) – hence my spending some time on the distinction. I suggest that, if you are getting a bit lost here, you reread the section and raise the matter with your tutor/instructor; the point is quite an abstract one.

Super-naturalistic objectivism! First cousin to these naturalistic views, and open to the same sort of complaint, is the view that by ‘X is good’ is *meant* ‘X is commanded by God’ – a sort of *supernaturalistic* objectivism if you like! For a start, it would be a matter of some controversy as to which god (or gods) the

theory is to have in mind here; but whichever way we jump, this view has the same sort of problem as before. It implies that the moral propositions and judgements of atheists, or religious rebels, or those believing in some other sort of god to the preferred one are *literal nonsense*. For even the most religiously pious of you, this should be a counter-intuitive consequence. It is *conceivable* that there is some sort of supreme being and it might even be the case that he/she/it is perfectly good. If this were to be so, then one would expect such a creature's every command to be good. But this would just be the way things turned out, it is not what is *meant* by saying that something is good. An atheist can, in the fullest sense, employ our ordinary moral concepts to make moral judgements without being caught up in some sort of self-contradiction just because of her atheism.

Summary so far

Let me just sketch back over what we have been doing so far. We distinguished three sorts of enquiry about moral values and here focused on just one of these, the one asking about the *nature* of values and value judgements – that is, meta-ethics. Concerning the nature of moral values and judgements, we distinguished two main camps of meta-ethical theorist – the ethical objectivist and the ethical subjectivist. With its air of factuality, objectivism seemed a more comforting theory for those wishing to uphold, and perhaps impose, values. Looking more closely at objectivism, it emerges that, within that broad theoretical camp, there are sub-divisions. One of these varieties of objectivism is ethical naturalism. It itself has sub-varieties, depending on the particular natural property taken to be meant by 'good' (and some were illustrated). As it turned out, these sub-varieties seemed not to need detailed individual scrutiny because *all* versions of naturalism seemed to suffer from intuitively unacceptable implications (as outlined a short while ago).

So, the story to date is: one variety of objectivism fairly solidly clobbered, or so I suggest to you – what now?

Non-naturalism Outlined

Another variety of objectivism is, predictably enough, *non-naturalistic* objectivism (henceforth 'NNO' for short). There are other names in the literature for varieties of this view like: 'moral realism' and 'intuitionism' (the aptness of which will be clearer as we proceed) but I will stick to 'NNO' to emphasize the contrast with *naturalistic* objectivism. Like its naturalistic brother, NNO *is* an objectivist theory, that is, it *does* hold that goodness, badness and so on are objectively existing properties of the universe and are *not* just 'in the eye of the beholder'. The difference here is that it's *not* held that claiming that some action is good meaning-equates to making a claim involving fairly *ordinary* natural properties; rather, according to the NNO theorist, such a proposition asserts the presence in that action of the objective *and distinct* property of goodness. 'Goodness' is *not*

seen to be just a 'shorthand' way of referring to something otherwise familiar, some sort of natural property (like: 'what society approves of'). Rather, it's its own beast – a *different sort of property in its own right*. So, if you like, we can, on this view, divide the world's properties up into two sorts – 'ordinary' natural ones (of varying complexity) like size, number, mass, conduciveness to human survival and so on, and 'special' *moral* properties like goodness, rightness, badness and so on. (In contrast, the ethical naturalist had only one list – ordinary natural ones – some of which are *also* labelled 'goodness', 'badness' etc.) Generally speaking, NNO has been more kindly thought of in the meta-ethical literature than has moral naturalism. In particular, it is not open to the meaning-equation problems that naturalistic theories got into when they insisted that what was actually comprehensible was incomprehensible nonsense.

Now this is all very fine sounding but, as usual, we should look at what arguments can be advanced in favour of such a theory and what against it. The great merit of the theory, like any objectivist view, is that it holds out the hope that moral *truths* are available, that some views about what conduct is good and what bad, are just true and others are just false. It would be comforting to think that someone who holds the view, say, that having sexual relations with children was morally permissible was not merely morally unusual in her views but in moral *error* in some objective sense, a 'moral flat-earther' whose views were just, as a matter of (moral) fact, *false*. And, with respect to the full spread of our professional ethical controversies, it would be nice if there were true answers to be had. It would be especially comforting and motivating if you could be assured that *your* views were not *merely* your views (like your taste in music) but captured moral truth.

On this view, a moral theory, like a physical theory, is a hypothesis about the nature of reality and moral disputes would be disputes about what the universe is really like with some of the people in dispute, no matter how firm and sincere their conviction, being in error, *objectively* in error, their view being just false. And this would be in just the same sense that, *if* the earth goes around the sun and not vice versa, then the geo-centrist theorists just have a *false* view. And *if* Jesus BarJoseph was an obscure Jewish revolutionary with delusions of grandeur and not the son of God (as, say, there is no such entity), then theists just have a *false* view. And *if* the speed of light is not a constant, then modern physicists just have a *false* view. Similarly, *if* having sexual relations with children is right, not wrong, then most of us just have a wrong view – a *false* hypothesis as to the moral facts on this. All of these views on various topics would just be different ones about different aspects of reality and be true or false depending upon what the objective facts of the universe happened to be. (On this view, it might seem appropriate to rename ethics as 'moral science' – a term I understand to have been in vogue in the past.) Of course, all of this could be said of naturalistic objectivism as well. The above paragraph is really just reinforcing the message of objectivism of any sort. The difference is only that the moral principles are not held by the NNO theorist to be about ordinary physical reality of about some special sort of moral aspects of the universe.

Standard Objections to Such Non-naturalistic Analyses

Advance upon naturalism though it may seem to be, there are problems with NNO. One trouble lies with all of those ‘ifs’. Morality seems in rather worse shape than science (though perhaps no worse off than religion) in *establishing* or *justifying* its claims. So: even if there are some moral facts, how do we *know* what the moral facts are? My senses of sight and sound might tell me that, say, a child is being bullied before my eyes but my judgement that this is wrong goes beyond what those senses supply. So, what warrants my claim that it is wrong?

One suggestion in this context is that we have something like a ‘sixth sense’, an organ of moral intuition which just lets us ‘see’ moral properties of goodness and badness in much the way that the other senses are taken to apprise us of the presence of more ordinary properties.

There are several worries with this idea of a sixth moral sense. *First*, we don’t just make judgements about actions going on in front of us but evaluate things far distant from us in space and time (something generally not possible with our other senses, although astronomy is an interesting exception). For instance, I make the judgement that William the Bastard of Normandy should not have made Harold Godwinson swear an oath of fealty to him by threatening him with indefinite imprisonment if he didn’t. How can a sixth *sense* account of moral value judgements cope with cases like this? *Second*, we make moral judgements about whole classes of actions. That is, we issue general judgements like: ‘Stealing is wrong’. Yet we haven’t ‘seen’ all instances of stealing, so we can’t have ascertained the unfailing presence of the property of wrongness by applying our sixth moral sense to all cases. But perhaps I’m being too hard on the NNO theorist here, he might claim that the same goes for general propositions in science and claims by scientists about distant places and times; so mightn’t the moral objectivist defend his claims in much the same ways that a scientist would? Mightn’t we build up general moral views by induction from particular instances? – perhaps. Or have the hypothesis that all stealing is wrong as the best explanation of a spread of particular ‘observations’ that this, that and the other case of stealing is wrong? – perhaps. Clearly a whole meta-ethical enquiry looms here and I am only gesturing at some possible initial moves. A *third* worry with this sixth sense view is to point to the non-independence of the sense. Generally speaking, the other senses operate independently of each other; that is, I can smell things without seeing them and so on but this doesn’t seem to be so in the case of moral judgements. If I have *no* information from my other five senses, then I can make *no* moral judgement. Say someone is being murdered in front of my blindfolded eyes and stopped-up ears. Presumably (if she thought it wrong) the NNO theorist would claim that the air would be veritably reeking with the moral property of badness yet I would not be able to detect it *directly* with my claimed sixth moral sense. All very odd.

Of course, as with everything in this chapter, some way around the objections might be found. Nonetheless, some NNO theorists have responded to them by playing down the notion of a sixth sense. They claim that it is *not* as if we can

detect in any direct way *particular* instances of particular moral properties being present; rather, our intuitive powers allow us to recognize *general* moral truths about our world.

Contrast this with our previous NNO version. Rather than ‘seeing’ wrongness, say, in particular situations one is confronted with and then, perhaps, building up more general principles by induction, or whatever, from such particular items of moral ‘data’, one intuits the truth of the general principle *directly* and then applies it to make judgements in particular cases.

This gets us over the oddities of the ‘sixth sense’ notion but other worries remain with NNO. As to whether you consider these to be serious enough to disincumbent you to accept any form of NNO is for you to judge (however tentatively).

One family of worries concerns the ‘physics’, if you like, of what is going on here. If they exist at all, then rightness and wrongness seem to be very strange properties indeed. Quite what is their relationship with other, more ordinary, properties like mass and length? Should our account of the laws of nature include laws of ethics and, if not, why not? Of course, mere weirdness is not automatically a problem – consider the bizarre nature of the accounts current in micro-physics. However, the seeming absence of moral properties from our ordinary scientific accounts of the functioning of nature is problematic. Nor is it at all clear how inclusion could occur.

Another family of worries concerns our access to, or knowledge of, what situations, or types of situation, have goodness or badness present. We have seen the problems with the ‘sixth sense’ account of detecting goodness or badness. The intuition of general moral principles’ truth is also difficult to fathom. What is going on when a moral agent just ‘intuits’ that, say, all stealing is wrong (or that all stealing is wrong unless it is food and one is starving, or ...), or, at a more fundamental level perhaps, that acting for the greatest happiness of the greatest number of humans is always right?

One model that is sometimes appealed in answer to this question is that of mathematics. It is sometimes held that the basic axioms of various mathematical systems are some sort of necessary truths, that merely entertaining them leads us to see that they could not be other than true. Some philosophers have held that some normative ethical principles, particularly deontological ones, have that status. As usual, I enjoin you to hunt out some introductory texts on moral philosophy to pursue matters further and I will content myself with just one observation. Note that moral controversy rages on almost any matter that you care to name. This seems to sit uncomfortably with quasi-mathematical intuition of moral *certainties*.

In summary so far, there are two broad sorts of worry about NNO theories. One concerns the oddity of moral properties compared to ordinary ones and the other one, the challenge of the question: ‘Even if there are objective moral facts of the matter, how could we ever know them, how could we assure ourselves that we are not “morally colour-blind”?’.

It would be nice if we could satisfactorily resolve these challenges to NNO theories, because it might place our inclination to, for instance, intervene and run

a criminal's life for her (on the basis we are 'only doing what is good for her' or 'good for society') on a firm factual footing. In particular, we would not be open to the challenge: 'That's just your view' that seems always to be able to be mounted against someone who is over-riding someone else's wishes on the basis of moral principles construed in accordance with subjectivist meta-ethical theory (as we will see below).

I would like now to keep exploring this 'How could we ever *know* what the moral facts actually are – even if they exist?' objection to NNO.

The question is not an idle one. Our worry has been to establish some sort of objectively firm basis for the ethical values governing professional practice (or motivating professional reform). If it is indeed to be a *firm* basis we had better be fairly sure of our moral facts. If a *naturalistic* theory of some sort were to be correct then the problem would be no more difficult than those facing science but what if the right theory is some version of *NNO*?

Look at the immense normative disagreement among humans on matters of morality. If some set of normative principles is objectively correct, then how could we know *which*? And what could you say to someone who disagreed with your ethics? That is, what on earth could you do to prove that you were correctly describing moral reality, not them? If there are moral facts out there waiting for us to, say, intuitively apprehend them, then there are a lot of people running around with faulty moral intuition. That is, people differ even on the most fundamental moral matters; they differ from place to place and time to time. How can we tell who are morally muddled and who not? The worry applies whether we have 'sixth sense'-style particular moral judgements in mind or 'mathematical principle'-style general moral principles as what is being intuited. A couple of the attempts to meet this concern are worth our attention.

One response is to say that we can tell which value propositions are true and which are false by seeing what *other people* think. That is to say, we might feel that, although some *individuals* can't intuit soundly, surely the *majority* can be relied on. This might be all very well except that it is unclear whose votes we should count – the majority *of what*? 'Those of our *society*' you might say – but meaning what and why? What counts as *your society*? – your town/region, state/province, nation ... ? However construed, a further worry is that the majority view *shifts* within a society over time and certainly from society to society and from group to group within any given society. Let me elaborate upon each of these concerns with this 'Trust the *majority* of society as moral truth see-ers' answer to our 'How would one *know* the moral truth?' challenge. I will deal with the former first, then the latter.

Say that you had opted for your *nation* as the most obvious unpacking of 'your society' for moral purposes. But why not a state or province within it (or, to go more fine-grained, a local community or, to go more coarse-grained, humanity as a whole)? The more compact the group that you look at, the more likely you are to find some intra-group consensus but, also, the more likely you are to find that the majority of that particular group is in ethical dispute with the majority view of

some *other* little group. The wider you cast your net, the more likely it is that there will be disagreement within the group.

There is also the issue of time. Societies change in various ways over time and one parameter of that change concerns the moral values endorsed. Of interest to many professionals would be the shifts in thinking about the relative importance of individual rights versus those of groups, and, particularly, of attitudes taken to individual freedom of thought and action and legitimate constraints upon them.

So, *when* do you fix the *time* for your survey? The views of any society that you will be surveying the ethical intuitions of will change over time. You might be a bit impatient with this concern and think that it is just obvious that *current* society is the correct group for the 'survey'. But why? Remember the context of the discussion. We were entertaining NNO as a meta-ethical theory and, in response to a 'moral blindness' worry, we advanced the view that, while individuals might have moral 'illusions' surely most of the group wouldn't. In this context then, one would have to be asserting of one's own current society that it is less prone to moral illusion than other societies in different times or places. How on earth could one justify that sort of proposition? Moreover, imagine others thinking in exactly that sort of parochial way at about their own society in some other time or place. What could one say to them? Sometimes people are tempted to respond: 'Well, what is right *for them* may not be what is right *for us*; it all depends upon the society that one is living in'. Again, though, keep in mind the context of this discussion, namely moral *objectivism* of an NNO type.

This view, if applied consistently to others in other societies, would have the extremely odd, if not quite incoherent, consequence that the moral *facts vary* from time to time and place to place (*unlike* the ordinary natural laws of, say, physics). It would also mean accepting that it was right for a Mongol hordesman to pillage, rape, burn, torture and so on; that it was right for Germans in the Thirties to persecute Jews, mediaeval Spanish Christians to persecute dissenters and so on. Also it is unclear how one should treat the not unknown case of society changing its mind on some issue as a result of a charismatic figure or propaganda campaign. If the majority view changes, does that mean that the facts change as well? – if so, they are oddly malleable and 'un-factual' sounding facts. Or should we say that either the earlier or later views of society were in error? That sounds better but leaves us with the original problem: how do we know *which* group speaks truth?

Given all the difficulties that we get into on that one, some would seek some other criterion beside majority opinion as a way of discriminating among rival and incompatible moral hypotheses. One such move is to appeal to some sort of god. On this view, the factually true moral hypotheses are those that the god tells us are true. (For theists, there is usually some sort of link held to obtain between the god believed in and principles about what is right and wrong and, as you will see, we have a version of such a link in each of our meta-ethical theoretical possibilities.) What can be said in response to this suggestion?

My first point here is to distinguish the view at hand here from our 'super-naturalistic' objectivism of a while ago. There, the view was that 'X is good'

meant 'X is commanded by God'. Unlike with that theory, the claimed link here is not one of *meaning*. Rather, it is the view that (*non-natural*) properties of goodness and badness are indeed an objective part of the fabric of the universe, are built into the 'physics' of reality. The only extra twist is the hypothesis that the universe was made by some entity and that his/her/its reports of the nature of the universe are thus our most reliable guide as to the moral (and other) aspects of reality. Note also that there is no question of the god's *commanding* being bound up with goodness (within our below discussion of subjectivism, we'll consider yet another theistic variant that does have such a link). Goodness and badness are just held to be *artefacts* of the creating god – woven into the objective fabric of the universe just as much as mass and electric charge and all of the other natural features of the universe are (although moral agents seem to have some creative role as well in that they make something, say, bad by doing it – all a tad weird).

As it turns out, there are considerable difficulties with this 'Ask some god or other what the moral facts are' way of resolving our difficulty concerning how to work out what moral hypotheses actually reflect reality, or capture the moral facts.

An obvious one is that it is hugely controversial, to say the least, whether any sort of god exists at all and, if so, what he/she/it/them is/are like. Even assuming that there is some such entity and that the god settled on indeed made the universe, how are we to have access to that god's advice as to the contents of the universe? Interpretation of religious tracts is a notoriously dicey and controversial business – that's how sects crop up. In any event, as might occur to you after the complexities of earlier chapters, what is available from such tracts is at best a series of simple maxims that might be crude starting points for one's moral principles but leave all of those complexities unexplored. (As the saying has it, 'The devil is in the detail' – not literally speaking of course!) So, how else might one get guidance from the god as to the sophisticated moral 'fine-grain' of the universe? Perhaps by individual revelation but, as the reports of individuals claiming such experiences conflict wildly, one would need some way of sorting out the simply psychiatrically disturbed from those who do indeed have access to the god's advice (if any).

In any event, let's assume that you feel that not only does NNO give the correct meta-ethical theory of wrongness and rightness, goodness and badness and so forth, but that, moreover, with some god or other's assistance or without it, you can *tell* objectively wrong situations from objectively right ones. All of this sounds very promising but there is still one last difficulty with the NNO view that I would like put to you (indeed it is a puzzle for naturalistic objectivism as well). Moral judgements are not *just* a matter of noting which acts are right and which wrong, of noting various aspects of the universe; they are *action guiding*. That is, to hold some act to be right is to hold it to be the right thing *to do*, one is obliged to do it. It's hard to see how any mere factual information about the universe (even about queer properties of the universe) could of itself *oblige* us to act in some particular way. Perhaps this point is just another way of wondering just what sort of property goodness is being suggested to be.

Subjectivism – More Detail

Given the difficulties in finding any sort of defensible version of moral objectivism, some theorists abandon it altogether and advocate the position of the ethical subjectivist. Crudely put, recall that moral subjectivism rejects any suggestion that there's any possibility of the truth or falsity of moral propositions. On this view, ethical principles and judgements are ultimately no more than a matter of taste, or preference. If I say that stealing is right and you say that stealing is wrong, then it is not as if one of us has uttered a falsehood. As neither proposition is to be thought of as attempting to describe reality, neither proposition can be thought of as failing to do that. So, you and I are held to simply have different *attitudes* towards stealing.

Standard Objections to Subjectivist Analyses

And this brings us to the major worry with this view. In contrast to aesthetic attitudes, moral attitudes (whether conceived of as subjectivism or as either version of objectivism would suggest) are standardly claimed to have *two distinctive features*, ones logically bound up with the concept of something being a *moral judgement*. First, they are said to be *universalizable*; second, they are said to be *prescriptive*.

What is meant by these terms? Well, if I say that it was wrong for Jones to steal then, to say that this judgement is *universalizable* is to say that I am committed, on pain of inconsistency, to saying that it would *also* be wrong for anyone else to steal. Or, rather, to keep it a bit more accurate, it would be wrong for anyone else *relevantly similar in relevantly similar circumstances* to steal (a necessary, but troublesome, qualification; but just what counts as *relevantly similar*? – much of the fine detail of equity disputes concerns the unpacking of this).

As for the *prescriptiveness* point, this notes that one isn't *just* expressing an attitude when one is issuing a moral judgement, there seems to be some sort of prescriptive element to it. In a moral judgement one tells people (including oneself) how to behave. So, if I say that people should tell the truth, then I am issuing an imperative, or prescription, instructing people, myself included, to tell the truth.

These two commonly accepted features of ethics cause trouble for the subjectivist. All very well, one might think, for me to have a certain set of preferences and govern and judge *my* actions in their terms. But if that is *all* they are, mere personal preferences, then why on earth should Jones or Smith or you be judged in their terms? And *how dare* I issue prescriptions as to how other people are to behave! This sort of objection has immediate and obvious relevance to the value clashes which led us into examining the nature of values in the first place. How can we be justified in imposing *our* moral judgements on other people if they are no more than mere preferences? Call this: 'the *prescriptive impertinence* objection'. There are four main suggestions as to how one might respond to it –

appeal to the authority of some group, some god, or to reason; or: 'tough out' the objection.

So, one response to this is to immediately insist that it is *not just your* view that you are imposing, it is the group's, say, society as a whole, with you (some professional or other in this case) merely acting as its agent and complying with its preferences. A couple of queries. What if you are out of moral step with the majority? Should you then act according to aims which reflect those majority value principles, even though they are ones that *you reject*? After all, on the subjectivist view under discussion, it is not as if their view is true and yours false (there's no question of truth or falsity); there are just more of them than there are of you. Anyway, if you are going to appeal to the majority, *which* majority? That is, a majority of which set of people. We tracked through a parallel worry earlier with the NNO theory. Various candidate groups will differ in their moral preferences so which group is to be deemed to have moral authority (and why)? Would the answer vary according to the sort of moral issue under discussion? That is, might different groups have different domains of proper moral authority? And what is going to be the status of dissenting moral minority groups (say of a cultural or religious sort) that are a minority within the designated group? And so on.

Also raised in this context of appeal to the group is the issue of universalizability. If, say, society's majority preferences are to be imposed, then *on whom*? (This is a manifestation of our earlier worry about unpacking 'relevantly similar'.) Are they to be imposed just on members of *that* society? Does that then mean that each society is to be pronouncing upon what is right for its *own* members but not on what is right for those in *other* societies (because relevantly *dissimilar*)? So, so long as some society is agreed that, say, Jews should be eliminated, or women not allowed to vote, or parents allowed to instil their religious beliefs in their children, or whatever (write in your own moral horror story and some society, some when and where, has favoured it) then that is the appropriate set of moral preferences to govern the actions of *that* society's members. Enter ethical relativism with a bang – a self-destructive one.

In contrast to that 'majority preference' attempt, another response to this 'prescriptive impertinence' worry is advanced by many of those of some sort of theistic persuasion (especially those subscribing to some sect of major monotheistic Muslim/Judaic/Christian religions). It appeals to the views of some god or other. Distinguish this sort of theistic moral *subjectivism* from the two theistic theories discussed to date. Theistic *naturalism* had a *meaning* tie up between goodness and what the god wanted (but note that it isn't a meaning link up here, so this view is not vulnerable to our earlier worries about non-nonsense being deemed nonsense). Theistic *NNO* had objective moral properties, just ones that were *artefacts* of the god. On this view, however, appeal is being made to the *preferences* of the god (however we might ascertain them) as the preferences to which appeal is to be made as those having prescriptive force when various agents' preferences clash. So, on this variety, there is still rejection of the objectivists' idea of moral *facts* and morality is held to be subjective, a mere matter of *preference*,

but it is to be the god's preferences, not those of any grouping of humans, that count.

I will leave it to you to paint in some objections to theistic moral subjectivism as they will closely parallel ones already made against theistic NNO. Apart from these concerns, there is a more basic challenge. Remember that we came to this view as one response to the prescriptive impertinence objection. That is, why should I be judged by, or take any particular notice of, *your* preferences? That was the challenge. The theistic answer here is: 'Well of course no human's, or group of humans' preferences have any authority over you but God's preferences are different and do have that authority'. But, even if one allows that the god in question exists, the simple challenge here is to ask: 'Why so, why consider some non-human, supernatural creature to have prescriptive force over us?'. This proves to be an extraordinarily awkward question to answer in any sort of satisfactory manner. I'll leave you to play around with the arguments if you are of some sort of theistic persuasion and I will offer just one further comment. Most (Judaeo-Christian) theists' first go at answering the objection is: 'Because he created us'. Don't be too quickly satisfied with this response; it is very readily criticizable.

As a third response to the prescriptive impertinence objection, there is the suggestion that rational consideration of ethical issues will somehow solve this concern about clash of preferences. Some preferences, it is said, are forced logically upon us and the source of their authority is that they are necessitated by reason. This would be fine if it worked and there have been famous attempts at doing it (most notably by the philosophers Immanuel Kant and Thomas Hobbes). I don't wish to do more than gesture at considering this sort of view as it becomes rather complex fast. Suffice it to say that general opinion in moral philosophy would be that no existing attempt has escaped crippling objection. We considered a related point of view when talking about direct quasi-mathematical intuition of moral principles when talking about NNO. As usual, my advice is to flesh out this chapter's sketch with the standard literature in moral philosophy. View my comments as mere 'tasters'.

Finally, instead of trying to find some person or group whose preferences are seen to have some special moral authority, one can respond to the 'prescriptive impertinence' objection about judging others by one's own moral values (given that those values are admitted to be no more than deep preferences that one has about how the world is to be) by, as I put it earlier, 'toughing out' the objection. If some moral principles are indeed your *deepest* preferences, then what could possibly have more salience as a guide for living your life by (including your, perhaps interventionist, interactions with others)? What else can one sensibly do?

Another, related, objection to subjectivism concerns its plausibility as an account of what we seem to construe our ethical views to be. It seems too slight a status for them to be construed as mere preferences, albeit ones we might be willing to subscribe to, and indeed enforce upon, others. Surely, it is said, there is more to what is right and what is wrong than just subjective preferences, no matter whose. Surely people can just be bad, be in moral error – Hitler, say. Such intuitive

concerns would lead one away from subjectivism to some sort of objectivist view (though, as seen, they have their own problems).

So, in summary on subjectivism, it can be seen as a response to the severe difficulties of any variety of objectivism but itself seems open to the objection that, if moral views are *just* subjective preferences (albeit strongly held ones), how can it be anyone's duty to act in accordance with someone else's preferences?

All very difficult to sort out.

One last view that I wish to address, and which obviously has abandoned any hope of finding a solid basis for imposing values on others, is that of what I shall call 'the *moral anarchist*'. This view is most easily seen as denying the universalizing and prescriptive elements we spoke of above. It accepts that moral values are 'in the eyes of the beholder' and so is a version of subjectivism; but contends that one's own moral preferences should *only* apply to oneself and should neither be used to judge nor to prescribe the conduct of others, far less as the basis for controlling it.

Comments: first, one wonders whether this is really a theory of the nature of moral concepts at all, for plausibly the universalizability and prescriptive elements abandoned here are indeed inbuilt logical features of morality, part of the meaning of the moral concepts. Rather, anarchism is perhaps better construed as a rejection of morality, as the substitution of another, non-moral, way of governing one's life. Second, one might wonder what is going on here; has a second-order prescription been issued by the anarchist to the effect that other people *should* abandon other meta-ethical theories and embrace moral anarchism (no matter what they think)? If so, then there is a whiff of self-contradiction here! Another difficulty concerns conflict. Sooner or later people living their lives by their own individual anarchistic values will be in clash – man, as the aphorism has it, is not an island. What is the anarchist's account of what happens next – might is 'right', or what?

Summary

As you will have realized, normative and meta-ethics are complex areas and we have done no more than skate on the surface of things. That said, it is worth your while to at least be aware of a range of possible theories, some of which you might be sympathetic to (despite the objections). Professional life does involve intervention in the lives of others on the basis of the interveners' values and that does raise a concern about the status and nature of that basis for intervention.

Although we're not going to do so, should you wish to pursue normative and meta-ethical issues further, then the tools of clear, sustained rational dialogue that you have been introduced to in earlier chapters more or less transfer across to that task. Your argumentation within meta-ethics, at least, will not, however, contain much by way of moral propositions because you will be arguing *about*, not *within*, ethics. (And conceptual propositions will be of great importance and of some subtlety.) If you want to pursue such issues further, then, as I have said

before, I'd recommend that you seek enrolment in a moral philosophy unit within an arts degree or, less satisfactorily, buy from your institution's bookshop some introductory philosophy text that covers normative and/or meta-ethics and read it carefully.