



Postgraduate Dissertation coversheet & feedback form:

THIS SHEET MUST BE THE FIRST PAGE OF THE ASSIGNMENT YOU ARE SUBMITTING (DO NOT SEND AS A SEPARATE FILE)

<i>For completion by student:</i>		
Student Name Keith Dowling		
Student <u>University</u> Email:		
Module code: (HPPH7014)	S2student ID number: 2209177	Educational support needs: No
Dissertation title: <u>HOW OUGHT WE TO JUSTIFY OUR MORAL JUDGEMENTS?</u>		
Supervisor: Dr. Tristan Nash	Word count:14,978	
I confirm that this is my original work and that I have adhered to the University policy on plagiarism. Signed ...Keith Dowling Email submission: if you are submitting this form via email, please tick this box / (as a substitute for your signature) to confirm that you agree to the above declaration.		



Master's Degrees by Examination and Dissertation

Declaration Form

1. This work has not previously been accepted in substance for any degree and is not being concurrently submitted in candidature for any degree.

Name.. Keith Dowling

Date 16th April 2024

2. This dissertation is being submitted in partial fulfilment of the requirements for the degree of MA Philosophy.

Name...Keith Dowling

Date 16th April 2024

3. This dissertation is the result of my own independent work/investigation, except where otherwise stated.

Other sources are acknowledged by footnotes giving explicit references.

A bibliography is appended.

Name Keith Dowling.

Date: 16th April 2024

4. I hereby give consent for my dissertation, if accepted, to be available for photocopying, inter- library loan, and for deposit in the University's digital repository

Name Keith Dowling

Date 16th April 2024

Supervisor's Declaration.

I am satisfied that this work is the result of the student's own efforts.

Signed:

Date:

MA Dissertation

MA Philosophy

HOW OUGHT WE TO JUSTIFY OUR MORAL JUDGEMENTS?

Keith Dowling

Student No. 2209177

April 2024

Supervisor: Dr. Tristan Nash

Contents

Introduction		8
Chapter One	Moral judgements and Reasons	10
Chapter Two	Autonomy	17
Chapter Three	Universalisabilty	23
Chapter Four	Impartiality	27
Chapter Five	Practical Entailments	34
Chapter Six	Conclusion	40
	Bibliography	42

Abstract

What counts as a good defence of a moral judgement? In answer to this question, I begin with an account of the structure of the simplest moral judgements. I point out that for moral responsibility – praise or blame - to be appropriately ascribed to the person making it, such a judgement must be self-determined. However, we do not start from scratch in this. We can interrogate our received social values and also the obligations entailed by the social roles we fill. I then argue that for an adequate justification, our autonomous moral judgements must satisfy two formal conditions, universalisability and impartiality. More importantly, even if they are universalised and impartial, our judgement-making can be unsuccessful due to our lack of certain practical skills. These include the ability to properly identify the problem; the need to be sufficiently informed about the relevant non-moral facts in the case; the ability to identify the actual interests of those affected by our judgement; or to fully think through the alternatives open to us and be aware of the consequences of a preferred solution. Also our justification can be ineffective because of our failure to emotionally engage with the course of action prescribed, or to be insufficiently motivated to carry it out. It can fail argumentatively due to our inability to sympathetically explain our decision to others and, most importantly, due to our lack of experience in making decisions of this nature. Added to this, we must recognise that dissent is always possible and that we need to seriously consider and answer any objections. If these conditions are met, I maintain that this constitutes a good defence of a moral judgement. Moreover, I suggest that it is not then always an open question as to what our autonomous (universalised and impartial) moral judgements ought to be.

HOW OUGHT WE TO JUSTIFY OUR MORAL JUDGEMENTS?

No one seriously questions that I ought not knowingly infect other people with Covid, or that I should not steal from my co-workers in order to pay for my holiday, or that it is morally repugnant to drug a classmate in order to have sex with her. Often we express similar certainty about the moral wrongness of promise-breaking, lying, adultery, murder.¹ Of course, we need to be able to say clearly what makes such actions so obviously wrong. On the other hand, people have opposing views about questions concerning the morality of war, euthanasia, the treatment of animals, and numerous other 'grey area' issues. Often their views are mutually contradictory. Indeed, there are many who hold that there is not an adequate way of defending any of our grey area judgements.² I disagree. However I think that what counts as a good defence needs to be clearer.

So how should we argue in order to offer a good defence of a moral judgment? I will claim, firstly, that there are certain formal requirements that need to be met. However, I think we can improve on this. I will argue, secondly, that there are some practical considerations - many of which are entailed by the formal ones - that an adequate justification of a moral judgement needs to meet. These are the claims I will explore in this dissertation.

In **Chapter 1**, by way of clearing the ground, I will point out that to count as a moral judgement one or more factors in the situation evaluated will be picked out as a relevant moral reason. Speaking generally, these are action-guiding reasons that are usually thought to be other-regarding and overriding. However, they are not expressions of absolute principles but general principles that are always open to challenge and can be amended by subsequent thinking or experience.

In **Chapter 2**, I will argue that our moral judgements must be self-determined. To appropriately ascribe moral responsibility - praise or blame - the judgement-maker must be taken to be expressing their own informed, uncoerced choice. But more than this, albeit self-determined, our moral judgements must be (in some sense) 'objective'. The reasons we give in support of them must be reasons that another person could understand to be moral reasons, though they may not agree with them. With this in mind, I will argue that defensible moral judgments must meet certain minimal formal requirements. Which requirements?

¹ Or of moral judgements in support of honesty, kindness, integrity, Intellectual curiosity, etc

¹¹ See e.g., Hume (1751), Nietzsche (1886), Harman (1977), Foucault, (1984), etc.

In **Chapter 3**, I will point out that one condition usually given is universalisability. Our judgement-maker must accept that what she thinks she morally ought to do in a given situation is what anyone else ought to do in that situation if their cases are relevantly similar.

In **Chapter 4**, I will argue that a second (more controversial) formal requirement is impartiality. To be defensible, her moral evaluations must not prefer or benefit one person, or group, over another for “improper” reasons. I will look at both self-regarding and other-regarding justifications for this condition, but argue for impartiality in terms of the latter.

So our moral judgements have to meet certain formal conditions; they must be action-guiding, autonomous, objective, (universalisable, impartial) judgements about our own or another’s conduct mainly towards other people. But this well-trodden ground gives us only the formal conditions. The important question for us is: What counts as a good moral judgement-making performance?

In **Chapter 5**, I will argue that an effective use of the universalisability and impartiality conditions presupposes that the autonomous judgement-maker has mastered other elements of practical knowledge. The kind of things I have in mind include: (a) Our judgement-maker must properly identify the moral problem confronting her; which is not always as simple as it sounds. And when considering possible solutions to this problem, (b) to successfully universalise she needs to be adequately informed about the relevant facts. And to be evaluatively impartial, she needs (c) to consider the (usually) different interests and values of those affected by her judgement. And to satisfy both (b) and (c), she needs (d) to have thought through the alternatives that are open to her and (e) the most likely consequences of each possible solution. Also she needs to make sure (f) she has the appropriate feelings and (g) resolve to carry out her decision once it has been made. Added to this, there are other practical considerations to be taken into account. She needs (h) the ability to communicate her judgement to others in an appropriate and sensitive manner. And most importantly, she needs (i) plenty of experience of this kind of thinking. She will need also to recognise (j) that many of her grey area moral judgements will still be open to dissent and so (k) to consider any objections that other people have and to learn from them.

In **Chapter 6**, my conclusion, I will summarise the formal conditions we have considered and highlight the pragmatic considerations needed for an effective defensible moral judgement.

CHAPTER ONE

MORAL JUDGEMENTS AND REASONS

§1 Action-guiding judgements and reasons.

First of all, what is a judgement? Obviously, not all kinds of utterances are judgements. As K. Baier (1958:49) observes: ‘Judgements must be distinguished from guesses, conjectures, hunches, hypotheses, assumptions, suppositions and the like...’ We could add here that judgements are not merely the expression of desires, feelings, emotions, attitudes, or opinions.³ This narrows the field a bit. Karl Popper (1945:57) thinks that the important point about a judgement is that this is a way of stating one’s own position: ‘...a judgement indicates where the speaker stands’. And we might add it is usually a decision that the speaker wants her interlocutors to accept. But ‘a judgement’ involves more than this. Someone can get another person to accept ‘where they stand’ by an appeal to their authority (e.g. ‘Because I say so!’) or by coercion (‘Otherwise I’ll break your fingers’). When we talk of ‘a judgement’ we usually indicate that the speaker intends to rationally persuade her interlocutors that she thinks there are good reasons to show why others should accept it. This is to say, the conclusion of any judgement (let us call it X) needs at least one good reason (A) intended to justify X: the reason A is deemed to be an adequate justification of X. (What makes a reason “adequate” for a conclusion in moral contexts is the question I will attempt to answer in this essay.)

We can contrast the sense of judgement here with the notion of ‘opinion’. In the latter, usually a statement is put forward where no reason (implied or explicit) is given to support it.⁴ That this differs from a judgement becomes obvious when different opinions conflict with each other. In most cases, only one of them can be correct. But which one? The question is answered by each participant trying to give adequate reasons in support of their “opinion”; this is to say, they try to convert their opinion into a judgement by giving compelling reasons in support of it. And then (hopefully) they decide the matter by the quality of the arguments presented.

I need now to say something briefly about “action-guiding” judgements. In straightforward cases, with an action-guiding judgment our protagonist intends to guide her own or another’s conduct, or to evaluate their actions once they are done. The decision X is our

³ A.J Ayer (1956:107-108) seems to think this is all there is to it. In moral judgements we are reacting emotionally to a thing ‘in words’, in a way that is analogous to cheering or booing.

⁴ It is sometimes suggested, I think quite wrongly, that moral judgements are “just opinions” and it is implied that one opinion is as good as any other - and that we cannot usefully criticise any of them.

judgement-maker's considered choice of what should be done - her choice for performing some-or-other action X, from the various possible actions X, Y, Z open to her. And the reasons given A, B... are the considerations that count for her preferring X. For example, 'It's going to rain later today (A) so take an umbrella with you (X)'; 'You'd better practice now (A) if you want to play this weekend (X)'; 'You're going to a funeral (A). Don't wear florescent socks (X)'. The reason A is intended to suggest the reasonableness of the conclusion X. Whether or not the interlocutor accepts X is usually a matter of whether the reason(s) given in favour of X are thought to outweigh those that might be raised against it; e.g. 'I'm not going out today'; 'I haven't been picked to play this weekend'; 'No one will see my socks'. I intend the examples here to be outside the bounds of morality; they refer to things to which moral right or wrong are not obviously applicable. However, in the simplest cases, moral judgements are also action-guiding; *vis.* a decision and the accompanying reasons about how one morally ought to act, or what one should have done. So what distinguishes a moral action-guiding judgement from the rest?

§2 Moral reasons and judgements.

Although we might make them in descriptive terms, in a moral judgement one or more factors in the situation evaluated will be regarded as a decisive moral reason. Moral reasons are usually non-factual, (however see §10(b) ahead). They specify what a person morally ought to do, to believe, to choose, or what is considered to be a duty or a right - (and just as often, they refer to what someone ought *not* do, etc.). Thus, in support of her considered decision 'You ought to do X', the judgement-maker will offer a justifying reason A....because you promised her you would'; or '...because it would be kind of you'; '...it would be cowardly not to'.⁵ In other words, moral reasons are obligation-making considerations, expressing thoughts like the virtue of 'keeping promises', of 'doing kindly things', or of 'not being cowardly'; reasons intended to guide our choices or decisions (and subsequent actions) that are thought to be morally right or wrong. More importantly, when we give ourselves or our interlocutor a moral reason for a decision X, we reckon that X is then something that we are accountable for doing, or that we can be blamed for not doing or that, lacking an adequate excuse, failure to do X makes us culpable. There are two further conditions we need to note here, that are claimed to be distinctive of

⁵ Ideally, it is undertaken before any action actually happens, when the judgement-maker tries to determine what ought to happen.

moral judgements

Traditionally understood, moral judgements primarily involve actions that affect other people. They refer to, as Bernard Williams (2012:12) writes: ‘...the demands, needs, claims, desires, and generally the lives of *other people*...’ In short, they are other-regarding. The point here is that often our judgement-maker will have the option of choosing to do what best suits her own desires or interest – a self-regarding act - when this means sacrificing the interests of others. But as Phillipa Foot (1972: 309) notes, moral judgements ‘...necessarily give reasons for acting, irrespectively of the agent’s “desires or interests”’. This point underscores the idea that in contexts demanding a moral judgement, our actions can benefit or harm other people. For instance, most of us have family attachments or we play roles within a community – and within such relationships, more often than not, the decisions we make affect other people. And if I am faced with a conflict between pursuing my own self-interest or of pursuing the interests of other people, normally (from this traditional moral perspective) I ought to disregard or sacrifice my own to promote the interests of others. However, this issue is not as clear-cut as my discussion might suggest.

Some philosophers maintain that self-regarding considerations can be – and others say they “should be” - the primary motive for our moral reasons and judgements. For instance, John Rawls (1972:136) tells us that in determining what is morally right ‘...each person is allowed to do whatever, in his judgement, is most likely to further his own aims’. For Rawls, it is possible to decide how to judge morally based solely on what I take to be in my rational self-interest (more on this point shortly). Michel Foucault puts the latter point more bluntly. He (1983b: 263) thinks that moral judgements *should be* concerned with self-interest; they ought to concern ‘...how the individual is supposed to constitute himself as a moral subject of his own actions’. Foucault thinks that in moral judgements one has a primary duty to enhance one's own well-being, one morally should ‘... develop one's own talent, cultivate virtues like courage, temperance and prudence ...’ Moral judgements, in other words, ought to be instruments of rational self-interest.

This seems to overlook the point made earlier that many of our actions can harm other people, either intentionally or unintentionally. Knowing this, I can choose to act fully aware that my self-interested action will have bad consequences for them. And surely, when I know this – that my action is likely to harm their interests - then at least I should consider whether or not I morally ought to proceed with it. At least, for the moment I will assume that actions which indisputably affect others are basic within the moral arena; and that they have moral priority over the self-regarding alternatives.

To add to our difficulties, Dewi Phillips (1977:140-57) has argued that moral considerations should always override non-moral ones. He (1977:140) writes: 'Moral considerations, we want to say, must be taken into account..' On the other hand, as Susan Wolf (1982: 420) points out '...things that matter morally are not the only thing that matter to us...moral values are... just one set of values among others'. *Contra* Phillips, Wolf thinks that often non-moral reasons should be decisive. Who is correct? Should even the weakest moral reason trump a strong non-moral alternative? For instance, in my despondent state on the way to a funeral, morally ought I smile at a passer-by who could thereby be cheered by this act? We might, quite reasonably, think here that the implied self-regarding consideration overrides the other-regarding nicety. However, I will argue that once the formal conditions (to be discussed ahead) have been met, there is a strong conceptual connection between one's moral judgement and the view that this should override all other relevant (e.g. personal, political, religious, scientific) considerations (p.36). Let us turn to a more pressing difficulty.

§3. Absolute or general principles.

The moral reasons our judgment-maker gives will be of varying degrees of generality. For example, the reason that 'One ought not to tell lies' may be defended on the grounds that 'lying diminishes the trust between people'. The former could be regarded as less fundamental than the latter. And no doubt, in turn, we could find a more basic reason than "One ought not diminish the trust between people'. To some, it seems obvious that this process cannot go on forever. There must be a stopping point; one fundamental moral reason(s) which we all regard as the end of the matter; so that nothing further needs to be said. For instance, if 'murder is wrong' is such an absolute principle then murder is always wrong. It is not right under any circumstances, even if a certain culture glorifies homicide, or even if the murder in question is of a very powerful, grossly evil politician. Similarly if female genital mutilation is morally wrong, as an absolute principle it was always wrong - and the fact that it is the practice in some cultures today does not make it right. It means that this culture institutionalises a moral wrong. To put the point here differently: there is one – or a set of - absolute principle(s) that entails all other moral beliefs and against which other beliefs and conduct can be gauged. We have room to briefly consider four attempts to justify this absolutist approach.

One such justification maintains that we know there are such truths by intuition. In this vein, G.E. Moore (1903: vii) writes '...Just as you cannot, by any manner of means, explain to anyone who does not already know it, what yellow is, so you cannot explain what good is'.

'Good', like yellow, is 'out there' and, what is more, it can be apprehended by intuition. In a similar way, since it is grounded in the moral good, an absolute principle – e.g. One ought to enhance the trust between people' - is "out there" and can be known by intuition.

This proposal is unsatisfactory. If absolute principles are "out there" for us to intuit, how is it possible that well-intentioned people come up with different moral intuitions? Of course, it is possible that someone, with sufficient ingenuity, could come-up with an intuition with which we all agree e.g. 'Boiling children in oil for fun is always wrong'. But it seems unlikely that there could be more than one or two such examples. They would be insufficient as ultimate support for the wide-ranging, every day moral judgements we need to make. Furthermore, if we have only Moore's approach to help us, what is to be done when disagreement does occur? When one person intuits a (putative) absolute principle which another does not? We might say that one of them needs to "get into" their intuitions more fully. But how do we know which one? Each disputant can accuse the other of having a mistaken intuition. The trouble is, on the intuitionist account (or an 'instinctive' account, see §8 ahead), there is no rationally defensible way of deciding who is correct. To make matters worse, when the self-evidence of one or the other's intuition is denied, the key assumption of this approach is undermined.

A different attempt to support the 'absolute principle' approach is by an appeal to authority. An obvious example here is The Ten Commandments⁶: e.g. 'You shall not murder'; 'You shall not commit adultery'; etc. They are unquestionable because they are given to us by God. Along similar lines, it is claimed that each culture has a set of tried-and-tested moral norms, some of which are regarded as absolute by all members of that culture. We are told, for example, that 'You should tell the truth', 'You should wait your turn', 'You should be tolerant of the other person's views' are absolute principles for British people.⁷ And indirect philosophical support for this empirical claim is offered by Michel Foucault (1984a: 25) who writes: 'By 'morality' one means a set of values and rules of action that are recommended to individuals through the intermediary of various prescriptive agencies...'. The moral judgements anyone makes are nothing more than an expression of their own cultural norms. And, apparently, in British culture some norms are thought to be beyond question.

The irrelevance of these or any appeal to authority can be shown in the following way. Either there are good (or bad) reasons for each one of the so-called absolute principles, or there are not. Now if the

³ *The Bible* Exodus 20

⁴ 'What does the UK think is right and wrong?' <https://www.studying-in-uk.org>.

authority in question has no reason or a "bad" reason for a moral principle - it is, say, an arbitrary whim or a way of maintaining their power - then these are not rational grounds for you or me to accept them. The more likely alternative is that the authority has "good" reasons for them. But then if there is a good reason for A, surely, it is because of this that one should recognise principle A as absolute, not merely because an authority says so - albeit God, or most British people.⁸ So let us turn to two attempts to give good reasons for thinking that a certain principle(s) is absolute.

One such account is J.S. Mill's (1861:6) version of the utilitarian principle: '...actions are right in proportion as they tend to promote the greatest happiness for the greatest number, wrong if they tend to produce the reverse of happiness'. Actions that conform with this principle are (always) morally right; actions that fail to conform are wrong. However, Mill's argument for this foundational principle is notoriously flawed. He (1861:32-33) tells us: 'No reason can be given why the general happiness is desirable except that ...each person's happiness is good to that person and the general happiness, therefore a good for the aggregate of all persons.' Each one of us thinks their own happiness is good therefore one ought to think that 'the greatest happiness of everyone' is good. I hope that it suffices to point out that Mill commits the Fallacy of Composition here. It is like reasoning that because each husband loves his own wife, any husband ought to love everyone else's wife!

The 'greatest happiness principle' is badly flawed for other reasons. For instance, as many critics have pointed out, this (absolute) principle can be used to justify *prima facie* morally objectionable actions; e.g. I seem to be able to justify deliberately taking the vital organs from an unwilling P in order to relieve the suffering of a greater number of other sufferers Q, R, S. Worse still, there are many examples where maximising the greatest happiness (e.g. most people living in Nazi Germany) involves the suffering of victimised minorities (Jewish people, the physically or mentally disabled, gypsies, etc.). Also, as I will argue shortly, the value of one's personal commitments cannot easily be traded-off against considerations of the general happiness. On grounds like those above, the ultimate status of the utilitarian principle can be called into question.

We have room to consider one last absolutist claim; Kant's (1785:85-86) contention that values like 'keeping one's promises', 'developing one's talents', 'helping others in need', and 'not committing

⁸ We might add that in the past often the moral "certainties" of most British people were mistaken; e.g. the acceptance of slavery, the treatment of women, etc.

suicide' can be shown to be absolute since they are entailed by the categorical imperative. Of the latter, Kant (1785: 68) writes: '...what I have to do is ... ask myself only 'Can you also will that your maxim should become a universal law?' Using this rule, he maintains, that if the maxim 'make a lying promise' were to be universally acted on, then the point of the particular lying-promise would be lost since we could not have an institution of promise-making which a deceitful promise presupposes. Promises would never be taken seriously. Thus 'keeping promises' (developing one's talents, etc.) are absolute moral principles.

Again, we are faced by many problems.⁹ For instance, quite often we are forced to choose between two so-called absolute principles. So what should we do when we cannot do both? Here is a (somewhat extreme) example of such a dilemma. Using the categorical imperative 'One should never tell lies' and also 'One should be loyal to one's friends'. Now suppose a madman with a knife is bent on murdering my innocent friend, who is hiding in my house. And the would-be murderer asks me if I know where my friend is. Should I lie to him? Or should I tell the truth and be disloyal to my friend? Which principle should be regarded as absolute? With only Kant's categorical imperative to hand, it seems impossible to know.

I hope I have said enough to show that an appeal to ultimate principles - at least the candidates considered above - is defective. We might disagree with someone's intuition that a principle is ultimate, or their appeal to an authority that simply asserts an absolute principle. In such cases, if we are to be persuaded rationally, a justification for the claim is required. But the attempt to justify the principle must not be fallacious, or (where there is more than one) result in mutually conflicting directives. More importantly, I think this whole approach of trying to identify absolute moral principles is misleading because it misrepresents what is involved in defensible moral judgement-making. Before we turn to this, let me summarise the main points in this Chapter.

I have argued that in our moral judgements some factors, in the situation evaluated, are regarded as moral reasons for the conclusion that we morally should/should-not act in a given way. Such reasons are usually thought to be other-regarding and overriding. But they should not be taken to be instances of so-called "absolute principles"; for they are always open to challenge. This latter claim, moreover, is supported by a further important consideration. When making a moral judgment, the judgement-maker is taken to be self-determined and to be giving their own informed decision.

⁹ As noted, Kant thinks we have a duty to 'develop our talents'. Are we really morally obligated to do so? If someone squanders her piano-playing talent has she done something morally wrong?

CHAPTER TWO

AUTONOMY

§4. Moral judgements must be autonomous.

Morality allows individuals to be, indeed insists on their being self-determined, in the sense not only of deciding what they are going to do, but in determining for themselves what they morally ought to do. In short, we must make up our own minds about moral issues; we need to be autonomous. As Robert Sparrow (2007:7) writes: ‘To say of an agent that they are autonomous is to say that their actions originate in them and reflect...ends that they have themselves, in some sense, chosen. Their ends result from the exercise of their capacity to reason on the basis of their own past experience...’

This somewhat daunting claim is supported by the following consideration. To praise or blame someone for an action she has done presupposes that *she* has - or could have - decided freely and rationally to act in this way. Moreover, it is because we think that she has formed her own judgment and then chosen to act in accordance with it, that we hold her responsible for her behaviour.¹⁰ I think this point is important and is worth stressing. The idea of a person being morally autonomous entails the idea of that person forming her own judgement - appreciating the force of some-or-other relevant moral reason about the matter in question and then making her own decision and acting accordingly. And this rules-out her merely conforming with a received code of behaviour, or an unexamined set of principles. But how is she supposed to make her own autonomous judgements? It would be very odd to suppose that she could start all of her moral thinking from scratch.

One answer is that the autonomous moral judgements she makes will very often include her endorsement of - or sometimes her disapproval of – the kind of thing we noted earlier, a prevailing cultural or religious norm. She can decide for herself whether or not, in the present circumstances, she morally ought to tell the truth, or wait her turn, or not commit adultery, etc. The point is, to be autonomous precludes her unthinkingly conforming with a preferred code – but not of ignoring it completely. She needs to interrogate her received values - question and probe them. She needs to ask: What good reasons are there for this

¹⁰ I have been asked: What should we make of the person who is unfailingly kind, honest or loyal, but has never given this a second thought? Of course, they will be praised – but the praise is only fully merited if before they act, they have thought the matter through and decided that this is what they morally ought to do. Similarly, we condemn someone who performs evil actions, even if it is obvious that ‘...he never realised what he was doing’ (see H. Arendt on Eichmann (1965:287-288)). The condemnation is far stronger if we realise that his evil actions have been fully thought-through.

commandment, cultural norm or practice? And is this a reason that ought to guide my thinking and behaviour towards other people? The point is, if she thinks she morally ought to conform with a generally accepted norm, this is because she sees the reasonableness of the norm in question - not simply because it is the general practice. Her autonomy is manifest in her ability to reflect and decide in this way.

There is another important source of received values which our self-determining judgement-maker needs to interrogate. These result from what Christine Korsgaard (1996:100-102) calls 'one's practical identities'. All of us have practical identities. Some arise, for instance, from our family relationships (I'm a husband, a father), from our jobs (I was a lecturer), from our friends (my fellow students in the MA graduate cohort), or our wider social or philosophical allegiances (I'm a long-suffering member of a political party; I am a vegetarian), and so on. As Korsgaard (1996:101) writes: '... (these are) descriptions under which you value yourself...under which you find your life worth living and your actions worth undertaking'. And they are roles that help to define your relationships with other people. They indicate how you should behave towards them; or, as Korsgaard writes: '...(they) give rise to unconditional obligations'. Husbands morally ought to be faithful to their wives; fathers should help to provide for their children, as a lecturer I should have done my best to guide my students' studies; as 'their friend' I should stand by my pals in their adversity, and so on. Each of the roles that make up my practical identity place certain obligations on me. And I can't violate these obligations (at least completely) without, in the process, risking this aspect of my identity. A husband who doesn't see that he morally ought to be faithful to his wife is a bad husband; a father who doesn't see that he ought to help provide for his children is a bad father. One further point. None of us creates the moral obligations entailed by our various practical identities. They have evolved over time. And they are one of the raw materials, so to speak, out of which our autonomous judgment-maker can base her reasons for acting.

So how exactly does autonomy fit into the account above? How is she autonomous if our protagonist is required to follow the moral demands of a role? I think this is a difficult question. One answer is that we are assumed to choose freely - without manipulation - and rationally to take on a given role. (I could have chosen not to have married, not to have had children, not to have been a lecturer, etc.) And if once in a while our judgement-maker considers one of her particular roles and the entailed obligations, and decides to continue (or not) in that role and to meet the entailed obligations, then she is behaving in an autonomous way. Similarly, our judgement-maker can be morally autonomous (in this weaker way) when she decides to join the army and thereby to obey

the orders of her commanding officer - or she joins the church and obeys the moral advice of her priest. At a given point in her life, she chooses freely and rationally to join the army, the church, *et al*, and thereby takes on a role with prescribed values that seem to restrict her autonomy. However if, from time to time, she considers her role and the entailed obligations and decides to continue (or not) in that role and to meet these obligations, then she is behaving as an autonomous agent.¹¹

What about the few complex philosophical or political ideas that one feels committed to in a profound way? Let us suppose that, following Aldo Leopold (1949), our environmentally concerned judgement-maker decides that there are good reasons for thinking that the interests of 'the global eco-system' itself should have the greatest moral standing - in contrast to the welfare of the individuals (human beings, animals, vegetation) dependent upon it. (A dead ecosystem means there are no living individuals within.) Once again, it is an autonomous human judgement that marks-off the ecosystem's place as sovereign in the moral order of things, as it is a considered human choice that marks off what counts as an ecosystem. So once again the basic point remains – in such important matters, one is or should be the author of one's own judgements and actions.

This all seems a bit formidable. And the question arises: Given these limitations, wouldn't it be more rational for our judgement-maker to choose not to have any practical identities at all? But this overlooks the fact that people are almost always interconnected, their lives intertwined. As Korsgaard (1996:102) writes: '...it is not possible for a completely isolated individual to have anything resembling a human life.' For the latter, a social setting is required, that provides roles for us to fulfil – husbands, fathers, churchgoers, ecocentric holists, and the rest.

§.5 Objections to autonomy.

A big problem here is to accept the cogency of moral autonomy (in the way summarised above) in a world that is correctly described by scientific laws. If the causal principle is correct, it applies just as much to the desires, choices, commitments of all sentient creatures, including our judgement-maker. Surely, her moral thoughts, choices, commitments, have a chain of antecedent causes? So aren't her moral thoughts always subject to such determining forces? And if they are, what room is there for her to exercise her so-called autonomy?

One defender of the view that moral autonomy is compatible with a

¹¹. if she decides to continue in the role with its entailed obligations, to count as her (autonomous) moral decision she will need to satisfy the formal and practical requirements discussed ahead. Accordingly, "the tax benefits I will enjoy" would not be the right sort of reason

universe that is correctly described by causal laws is Harry Frankfurt (1971). He distinguishes between first-order and second-order desires and volitions. Roughly and briefly, a first-order desire is something I can be said to want, e.g. 'I want to offer you (the reader) a summary of Frankfurt's position', so I have a first-order desire. And this desire is also a volition – a desire I can act upon - because I am now trying to give such a summary. Now here comes the difficult bit. Frankfurt (1971:6) writes: 'Human beings are not alone in having desires and motives, or in making choices.... It seems to be a peculiar characteristic of humans, however, that they are able to form...second-order desires'. Having a second-order desire is a matter of having a particular reflexive relationship with certain of my first-order desires. Not only can I have a first-order desires, *vis.* I can "want" something – but, also (and very often) I can want *not* to want it. I can consider a particular desire and think that I would rather not have it (*vis.* I wish that desire would go away). For example, I really fancy a glass of beer (a first-order desire). But I am about to drive my car. Accordingly, I have a second-order desire not to want to drink alcohol. Moreover, my second-order desire can trump the first-order desire: (So "No beer for me, thank you"). Following Frankfurt, having second-order desires is what defines my autonomy. I can endorse and act upon certain desires and reject and not act upon others. So wanting a particular first-order desire not to issue into an action, is an example of my self-determination. Frankfurt (1971:17) writes: 'The enjoyment of a free will means the satisfaction of certain desires – desires of the second order'.

A dogged determinist, of course, could challenge Frankfurt's approach by insisting that our second-order desires are just as much causally determined as the first-order variety. However I think the Kantian alternative, that equates 'autonomy' with a set of peculiarly human competences, can be defended against determinism. This locates our ability 'to deliberate rationally and then to choose and decide' in the (small) set of capacities human beings have that cannot be satisfactorily explained as 'a caused reaction to events'. Human beings can think and respond autonomously with choices, and our judgements and actions are thereby open to the normal moral assessments.¹²

§6 Subjective autonomous judgements are rationally indefensible.

A more challenging difficulty is that since they need to be self-determined, moral judgements will be "subjective" and thereby - the implication seems to be - they are not rationally defensible.

¹² Unfortunately I do not have sufficient space to offer a more detailed evaluation of these accounts.

In this vein, the claim that a judgement is subjective is sometimes taken to mean that the judger's emotional states (will always) determine their judgment. To judge that 'X is good' is no more than to give one's emotional reaction to X. As Ayer (1956:104) maintains, it amounts to no more than : '...the feelings of approval (or disapproval) which a person...has towards the thing'. Alternatively "subjective" is sometimes used to imply that an autonomous judgement could be grounded in some trivial or quite irrelevant factor e.g. 'We morally ought to eat mussels on Mondays in March' yet still count as their moral reason.¹³ To put the point differently, it is implied that since they are subjective, our judgements could depend upon the vagaries of individual psychology (in the ways noted). And this means they are not rationally defensible. However I will argue that, albeit subjective, our autonomous moral judgements need not be clouded by such flaws. They can be made in a calm, dispassionate state-of-mind, with recognisably relevant good reasons based upon a clear understanding of the issues involved.

Another reason why it is said that autonomous moral judgements are not rationally defensible is the fact that they are expressed in everyday language, with its ambiguities and idiosyncratic interpretations, due to the imprecise meaning of everyday words.¹⁴ More importantly, it is pointed out that people have different background beliefs, assumptions and (even) methods of thinking.¹⁵ The good/right/ permissible acts referred to in autonomous moral judgments will inevitably be grounded in the mesh of these different background beliefs and thereby will not say anything uncontroversially objective about the thing evaluated. Thus they are not rationally defensible.

It is worth our noting, in passing, that there doesn't seem to be the same methodological difficulties when a judgement is made in Science or, say, in History. Ambiguity in the language used, difference in background beliefs, have no place in scientific judgement-making.¹⁶ Such judgements are defensible in that they are given in terms of verifiable facts (by observation, hypothesis, testing, etc); the reasons and procedures provided will be those incorporated in (to use a brief phrase) 'the scientific method'. Similarly in History, the kind of reasons that support a historical judgement refer to certain kinds of evidence; documents from the period in question, private diaries, and so on; historical judgement are rationally defensible insofar as they

¹³ see MacIntyre (1981:46).

¹⁴ Foucault [1971:378] shows how normative terms change their meaning more often than we suppose.

¹⁵ See Foucault (1966b: 281)

¹⁶ See Black (1946: 334-353)

meet these objective conditions.¹⁷ Unlike judgements in Science and to a lesser extent in History, an autonomous moral judgement, as we noted, is hampered by the lack of clarity of its key terms, due to their use in everyday language and in that they depend upon the particular background beliefs of the person making them.

Nonetheless, we have noted some of the less important conditions that need to be met in order for our autonomous moral decisions to be defensible (*vis.* normally they *are* action-guiding decisions, based upon other-regarding and overriding reasons). I will argue now that they become fully defensible when two other necessary conditions are met.

¹⁷ See Copi (1961:459)

CHAPTER THREE

UNIVERSALISABILITY

§7 A particular moral judgements must apply to all like cases.

It is often argued that a defensible moral judgement needs to be universalisable. As R.M. Hare (1963:15) writes ‘...if I call a thing a good X, I am committed to calling any X like it good’. If our autonomous judgement-maker thinks that something A is an adequate reason for judging X to be good or bad, right or wrong, in one particular case, it must be a reason for judging X to be good/bad, right/wrong in all like cases (unless there is a special reason for not doing so). In other words, to make a rationally defensible moral claim, another condition a judgement needs to satisfy is universalisability.

But a similar characteristic seems to be a property of all descriptive terms. As Hare (1963:13) writes ‘...we cannot without inconsistency apply a descriptive term to one thing and refuse to apply it to another similar thing’. When you call the small furry animal - with longish ears and a small, rounded tail that you see hopping out of its burrow - ‘a rabbit’, you are committed to calling any other thing that is exactly the same or similar in the relevant respects ‘a rabbit’. Regarded in this light, Hare thinks, universalisability is not a substantive moral principle (that is being smuggled into our account of moral judgements) but a logical feature of descriptive and moral terms when they are used in judgements. It is, he claims, one form of consistency.

So, according to Hare, to be consistent, our autonomous judgement-maker must be prepared to apply the moral judgment she makes in a particular case to any other relevantly similar case – unless she can show that the situations differ significantly in some way. In other words, she must treat like cases alike. When she is thinking about what she ought to do in the given situation, at the same time, she must consider the possible like-cases that her judgement commits her to and ask herself: Would I be prepared to treat these relevantly similar cases in the same way? Let me try to unpack this last point.¹⁸

Firstly, (at the pre-decision stage, where she is thinking hypothetically) it means that (i) her judgement about her own past or future behaviour must be consistent. If she thinks she had/has a good moral reason A to act in a certain way, she is thereby committed to recognising that, in relevantly similar circumstances, A is a good reason for *her* to act in the same way. She cannot think that her judgement ought to be universally acted upon and then, when the occasion arises,

¹⁸ The discussion here is along similar lines to my essay on the moral philosophy of Kant.

she ought not act in accordance with it. Such a person gives, as Sissela Bok (1978: 84) writes: ‘...the false appearance of virtue’. Less politely, anyone whose moral judgement X trades on the fact that they think X ought to be universally acted upon and then, when the occasion arises, behaves so inconsistently, we call a hypocrite.

Secondly, (when she is thinking hypothetically) she must include (ii) the spectator judgements that she makes concerning the behaviour of other people. If she thinks that A is a reason why ‘Smith morally ought to do F’ and Jones is in the same or a similar situation, she is rationally committed to judging that ‘Jones morally ought to F’. Considerations which justify or impugn one person must, in the same circumstances, justify or impugn another. We might think here of the way in which we say: ‘What’s sauce for the goose is sauce for the gander’.

Thirdly, it follows that (iii): the type of action she judges to be right for her to do in her own case commits her to thinking that this is right for anyone else to do in relevantly similar circumstances. Let me stress this point. If she decides that X is a right for *her* to do (for reasons A in circumstances P) then she must accept also that X is right for any other person to do in the same or a relevantly similar circumstances. Her moral judgement needs to go beyond “I” and “you” to the universal form of the judgement. We might think here of the way in which people regularly question the rightness of an action by asking: “What if everyone were to do that?”

This aspect of universalizability is frequently transgressed. We apply one set of standards to ourselves - or to people we care about - and a different set of standards to other people even though their circumstances are relevantly similar. Consider e.g. the politician who, during the Covid pandemic, fulminates against people attending parties but who himself attends parties. Where the inconsistency is so obvious we point to the double-standards that such behaviour displays. Incidentally, condition (iii) seems to be a compelling reason why ‘I morally ought not steal from my co-workers to finance my holiday’ (see the start of this essay). I certainly wouldn’t accept it is right for my co-workers to do the same. Another not too distant application of (iii) is evident in our condemnation of “free riding”, where the aim is to enjoy a universalised benefit without incurring its cost oneself. For example, individuals that benefit from collective international action to reduce global warming but make no effort to control their own carbon emissions. Which brings us to (what I think is) the most important test our judgement-maker must satisfy, to see whether or not she is prepared to universalise her particular judgement.

She needs to consider (iv) cases in which her judgement X is a directive guiding the actions of other people towards her, or towards

someone she cares about. Thus if she thinks that it is morally right for her to do X to you then, to be universalisable, she must accept that it is morally right for you to do X to her in like circumstances. Entailment (iv) of the rule is often used to expose a common error, namely, where we apply a harsh judgement about a stranger(s), but we are unwilling to accept the same judgement when it is applied to ourselves. In everyday situations, we might ask: 'Would it still be all right if they were to do the same thing to you?'¹⁹ I should add that here we have a compelling reason for me **not** to knowingly infect other people with Covid - because I would think it morally outrageous if they were to knowingly do this to me.

We are considering the claim that anyone who makes a particular judgement is logically committed to universalising it in terms of (i)-(iv). Or are they?

§8. Objections to universalisability.

There are lots of objections to this condition. We will consider three.

(i) First of all, as we noted above, consistency is a requirement that applies to practical judgements generally, and not just to moral judgments. If I judge that I ought to wear a raincoat when it is raining, then I think that anyone ought to wear a raincoat when it is raining. However, to recognise this is not to deny that universalisability is necessary for moral judgements - what it seems to indicate is the triviality of this condition.²⁰

(ii) It has been argued that universalising is not even necessary for a perfectly good moral judgement. Sartre (1945:36) tells us how one of his students, during the Second World War, was faced with the following moral dilemma: Should he stay in France and look after his ageing mother, or should he go to Britain to join the Free French Forces? Sartre (1948: 36) asks: 'Who can give an answer to that *a priori*? No one'. Suppose, say, that the student decides to join the resistance forces. It doesn't make sense to add that only if he is willing to universalise it is this the morally right decision; or that he should consider anyone faced with the same dilemma who acted differently, morally blameworthy. Sartre points out that a rehearsal of the universalisability rule (which he calls 'the Kantian ethic') in no way would help the young man solve the dilemma. He adds that the correct answer for the student is to '...trust in our instincts'. But, as we noted earlier (§3), an 'instinctive response' is

¹⁹ This aspect of universalisability has been around a long time; see Matthew 7: 12.

²⁰ See e.g., Don Locke (1968: 25-44)

inadequate as a defensible moral reason in a judgement.²¹ Moreover I will argue shortly that children have special obligations towards their ageing parents and, in the context in question, this is a weighty, universalisable reason. Let us turn to another objection.

(iii) Even if we accept that universalising is a necessary condition, it is not a sufficient condition for a morally good judgement. There are plenty of *prima facie* immoral acts to which the rule could be (and has been) applied. For instance, people have universalised maxims in favour of sexual, religious and racial discrimination. As an example of the latter, consider the advocates of apartheid in South Africa, who thought that the colour of a person's skin is the benchmark by which they should differentiate morally between people. And let us not forget the Apartheid government regarded 'the colour of a person's skin' as the relevant moral reason for discriminatory policies concerning political participation, education, housing, employment, healthcare, and so on. Evidently, they were willing to universalise the judgements that gave rise to such policies. No doubt, some would even accept condition (iv) above, (*vis.* if their skin had been brown it would be morally right for them to be on the receiving end of these policies). With only the universalisability condition to hand, it doesn't seem as though we could challenge them morally.

(iv) This points to a further objection directed at the rule. As Peter Singer (1991:12) notes '...a wide range of ethical views including irreconcilable ones are compatible with this notion of universality.' Just as the racist can universalize their maxim, so can the committed anti-racist - the committed sexist and the feminist, the pro and antisemite, etc. all can universalise their maxims. The fact that such diametrically opposed and irreconcilable outlooks could satisfy the rule suggests that the weight of moral distinctions falls almost entirely upon the extra evaluative assumptions that are required in addition to it, in order to produce a morally defensible judgement. The significant moral content is to be found in these further assumptions. So our autonomous judgement-maker cannot expect this rule on its own to give her strong, defensible reasons about right and wrong conduct.

This prompts the question: Are there any other formal conditions that, if they are to be defensible, her judgements need to meet? Many philosophers argue that moral judgments must not only be consistent, also they must be impartial. As we will see, this is an even more controversial condition than universalisability.

²¹ It has been pointed out to me that whatever the student does he must do something wrong (the act he omits to do) and so he would have grounds for guilt. And this would apply to anyone in a similar situation. However, I think the student has a special obligation to his elderly mother (p. 34) and this gives anyone similarly situated the strongest (universalisable) reason to stay.

CHAPTER FOUR

IMPARTIALITY

§9. Moral judgements must be impartial.

It is often maintained that a defensible moral judgement needs to be impartial. As Roger Scruton (2003: 52) (optimistically) claims: 'As soon as we set our own interests aside and look on human relations with the eye of the impartial judge, we find ourselves agreeing over the rights and wrongs in any conflict'. Scruton tells us that an impartial judgment is one in which '...self-interest, prejudice, or bias... have no decisive influence'. He adds that impartiality (like universalisability) is not itself a moral rule, but a logical feature of moral judgements.

There are various ways in which an evaluation may be described as 'impartial'. Flipping a coin at the start of a football match constitutes one way. The chooser does not need to consider anything else in advance of the toss and they are blind to its outcome. In this way, any bias, prejudice or special pleading about the decision is ruled out. A different kind of decision procedure that is expected to be impartial is where the judge's choice is based on established criteria of suitability, merit or attainment. The decision-maker thinks through the rules and the facts that apply to all particular cases and, if they judge in accordance with these criteria, they are thought to have judged impartially. Those choosing between applicants for a job, lecturers marking student's essays, Judges in law courts (or of beauty competitions) are supposed to take this standpoint. Again, the important assumption is that they have ruled out bias, prejudice or special pleading. We are told that the same idea is a necessary condition of moral judgements.

Before asking if this is correct, let us consider four contrasting ways in which impartiality is interpreted in moral contexts. For some writers being impartial is a prudential, self-interested matter. Two contrasting approaches here are (i) the psychological egoists, who argues that we cannot help but judge every act from the viewpoint of our own interests. However sometimes it is prudent to appear to be impartial (but this is an unnatural psychological constraint on our behaviour). More interestingly, there is (ii) the ethical egoists' argument that in some contexts, it is in one's own long-term self-interest to genuinely judge and act impartially. Quite different from these self-interested approaches is the view (of most moral philosophers) that the needs and interests of other people are what matter morally. Here the employment of impartiality in moral judgements is seen in terms of either (iii) adopting

the viewpoint of an Ideal Observer, (*a la* Scruton above).²² Or it is argued that what counts as 'being impartial' is (iv) the attempt to take into account the viewpoints of all of those affected by one's impending judgement. Let us consider each of these interpretations in more detail.

§10 An assessment of the different approaches.

(i) *Impartiality for a psychological egoist*

The claim that (i) 'human beings are always biased towards their own interests' is commonplace.²³ As a matter of fact, everyone must always judge in her or his own self-interest; genuine impartiality falls outside of the motivational range of normal human beings. This is not an ethical claim. It tells us how psychologically we must act, not how we ought to act. We cannot help but judge every act from the viewpoint of our own best interests. Nonetheless '...a person may do things for which he has no sufficient motive'.²⁴ He can feign impartiality. And sometimes it is prudent to do so because otherwise other people '...will not be persuaded'. But the motive behind the judgement must be self-interest. A more cautious version of this view is that while people can be impartial when their own interests are not involved, the moment they might have to sacrifice their interests - or do something irksome - then they are bound to champion their own cause.

A number of philosophers provide collateral for the psychological egoist's view.²⁵ For instance, Henry Sidgwick (1874: 498) writes: 'I am concerned with the quality of my existence - as an individual in a sense, fundamentally important, in which I am not concerned with the quality of the existence of other individuals and this being so, I do not see how it can be proved that this distinction is not to be taken as fundamental in determining the ultimate end of rational action for an individual'. What are we to make of this?

David Hume rejects the idea completely. He (1751:210) points out, for instance, '(that)...we give praise or approbation to ...things that cannot affect our self-interest at all, because they are too remote from us

²² See also e.g., J.S. Mill (1861:318) who argued that we should be '...strictly impartial as a disinterested and benevolent spectator'.

²³ See e.g. Thrasymachus (Plato's *Republic.*); Hobbes (1651), and for modern support J. Griffin (1996:92) who writes '...our natural perhaps genetic partiality limits our will.

²⁴ Hobbes (1651:93).

²⁵ In his justification of utilitarianism Bentham (1789:16) starts by offering a similar view of human motivation: 'Nature has placed mankind under the governance of two sovereign masters pain and pleasure. It is for them alone to point out what we ought to do, as well as to determine what we shall do'. See also e.g. Adam Smith (1776), R Firth (1952) for support of this view.

in time or place'. No doubt, you and I would praise the courage of Thomas Moore who in the Sixteenth Century refused to abandon his Catholic principles, even though this led to his execution. Another reason for rejecting the psychological egoist's claim that Hume (*ibid*) gives is the fact that we morally praise some things that adversely affect our self-interest; for instance '...a brave or noble deed performed by an adversary commands our approbation, while ...it may be acknowledged prejudicial to our particular interest'. In both examples, it is not obvious how our moral approval can arise from self-interest. I have room for one more (quite different) objection - that Hume does not mention. What evidence does the psychological egoist have to support the claim that human beings must always be motivated by self-interest? Since this is a psychological claim, presumably the support will be in terms of observable evidence. However, as we have noted, it is the evidence to the contrary that is striking. We seem to have any number of motives - other than self-interest - that inform our different judgements, like integrity, loyalty, courage, sympathy, fairness, generosity, and so on. In short, to persuade us, the psychological egoist needs evidence for their empirical claim. Let us turn to the interpretation (ii) of impartiality; vis. sometimes it is prudent to be genuinely impartial.

(ii) *Impartiality in terms of enlightened self-interest.*

John Rawls (1972) thinks we can be genuinely impartially and that sometimes it is in one's own enlightened self-interest to be so. For instance, suppose that I am asked to decide between alternatives A or B before I know whether my own position will be that of A or B. If I do not know where my interests lie, prudentially I ought to be impartial so that if I find myself in the disadvantaged position of A, I should do my best to minimise this disadvantage in my initial decision. In other words, genuine impartiality is needed in those situations where we need to "play it safe" - to try to ensure that we are not seriously disadvantaged in the future. Famously, Rawls offers us a method by which this kind of impartiality can be achieved. He suggests that in a 'reflective equilibrium' one can adopt an imaginary 'veil of ignorance'²⁶ from behind which one's actual interests and allegiances are "forgotten". In this imagined situation (1972:137) '...no one knows his place in society, his class position or social status, nor does anyone know his fortune in the distribution of natural assets and abilities, his intelligence strength and the like'. So the best thing for the prudential judgement-maker to do, to safeguard their future unknown interests, is to be impartial.

I will consider some problems for the account above shortly. Before this, let us consider a related way in which Rawls tells us how our

²⁶ For Rawls (1971:17-22) this idea is the basis for the principles and structures of his ideal political system.

judgement-maker can be impartial.

(iii) Impartiality from the perspective of an Ideal Observer.

Later in his book, Rawls (1972:184) gives a slightly different account of impartiality. He suggests: 'Something is right ...when an ideally rational and impartial spectator would approve of it from a general point of view... (when they) ...possess all the relevant knowledge of the circumstances.' This is the perspective of an Ideal Observer, someone who, together with all of the relevant knowledge about the action or situation to be evaluated, has the ability to calmly weigh all sides of the argument and only then make a judgement.²⁷ Following Rawls, the impartiality of an Ideal Observer viewpoint is best for determining what is right or wrong. Is it?

One worry here is that the impartial observer (like the chooser from behind the veil of ignorance) would not be Ideal if they lacked knowledge of some-or-other crucial fact. This is a tall order for an ordinary human being to meet. Even more of a problem is the requirement that when they judge, they must have an unemotional frame of mind together with the ability to weigh evenly all sides of the issue. The judge would need to be unmoved emotionally, for instance, when faced by the various trials of family life, or when they know that children in a different part of the world are being massacred or, presumably, at the murder of a loved one of their own. Certainly, they must not emotionally respond in the way that most normal people respond where, as Hume (1740: 295) observed '...reason is, and ought only to be the slave to the passions'. It might be asked (quite reasonably): Is (iii) really the right frame of mind in which to make a judgement in morally testing circumstances - simply to dispassionately observe what is happening and not feel anything? Which brings us to a more telling concern I have with the idea of an Ideal Observer (and also with Rawls' judge from behind a veil of ignorance).

This is the implication that there could be such an individual amongst the ranks of ordinary human beings; *vis* a fully informed observer who dispassionately watches what happens (or knows what will happen) in morally testing circumstances. Only God or an Artificial Intelligence, surely, could meet such weighty demands? Further, *pace* Hume, since in challenging moral contexts such a viewpoint would not seem to mirror anything like an ordinary person's experience, it is not very clear that the judgments of an Ideal Observer would be too helpful anyway. Let us turn to interpretation (iv).

(iv) Impartiality in terms of the viewpoints of those affected.

To be impartial, Peter Singer (1972:19) maintains: '...we (must) give equal weight in our (moral) deliberations to the like interests of all those

²⁷ See also e.g., Baier (1958:201); Scruton (2003:52)

affected by our actions'. Our judgment-maker must endeavour to see the issue she is considering from the viewpoints of all of those affected by her impending judgement - and not simply consider her own interests, or those she favours. She must try to imagine herself in the shoes of those affected and think what her decision might mean to them. I think this point is important and worth stressing.

An impartial (iv) point of view requires that no one individual's interest is intrinsically more important than anyone else's - unless there is a morally relevant reason for giving them greater or different attention. In the political realm, to give 'equal weight to the like interests of all those affected' requires that the structure of society, its institutions and policies should not favour one, or some groups over others by giving them benefits and opportunities that are not open to all. (I will return to this point shortly.) In the personal realm, as we noted, (iv) directs an agent morally to take equal account of the wishes and interests of those affected by their judgement. However this is not to say that there cannot be exceptions to rule (iv) (see 'special obligations' ahead), but the reason for the exception must be well-grounded. Incidentally, we have here an obvious answer to the question: Why is it morally repugnant to drug a classmate in order to have sex with her? (see the start of this essay). Such an act shows a shameful disregard for the wishes and interests of the victim. So we have in (iv) a way I think we might interpret impartiality. I need to ask now if we are justified in thinking that (iv) is a necessary condition for any moral judgement we make.

One way we might try to offer such a justification is in terms of the universalisability condition.²⁸ Like me, other people have interests (wants, wishes, purposes); they too can suffer pain or experience pleasure; they too have complex autonomous lives as moral agents; they too can universalise their judgements, and so on. We are *prima facie* like cases in these and many other respects. And the point is that I think that my interests ought to be taken into account when the judgements others make directly affect me. And since we are *prima facie* like cases in these important respects, to satisfy the universalisability condition, I should take their interests into account when they are on the receiving end of my judgment.

§ 11 Objections to impartiality condition (iv).

Some philosophers have questioned the appropriateness of any form of impartiality.²⁹ For instance, in an interview, Foucault (1984: 287) states: "Care for others should not be put before care for oneself. The care for

²⁸ But impartiality (iv) goes beyond universalisability; I can universalise, as we noted, my racist, sexist, jingoistic reasons.

²⁹ See e.g. Griffin 1996, Wolf (1992:145).

the self is ethically prior in that relationship....'. In short, an impartial (iv) moral view is seen to be undesirable; rather, we should judge and act in a way that serves our own best interest. Of course, when your and my interests happen to coincide, there is not a problem. But when our interests conflict, we each ought to give priority to our own interests.

A major problem for this view is that it cannot be consistently advocated. For when your and my interests conflict, my advising you to look after your own interest will not be in my self-interest! It is (logically) possible of course that our interests will never conflict. However, in the actual world, conflicts of this sort happen all of the time. And in such a conflict, I cannot think that you should pursue your own best interests – for this could undermine my own. And this means that Foucault's approach cannot be consistently advocated.

A more serious problem for impartiality (iv) is that it seems to leave no room for the 'special obligations' that we all have; for instance, the moral obligations that parents have towards their own children (and vice versa) or that friends have towards each other, etc. As Bernard Williams (1981:15) argues '... prohibiting an agent from seeing their own personal commitments as no more significant than those of any other, violates the integrity of agents and undermines any reason they might have for being moral'. Williams seems to think that the requirement (iv) is fundamentally at odds with the partial nature of our personal relationships. However I think this misrepresents impartiality(iv). It requires that in our moral deliberations we give equal weight to the interests of all those affected by our judgement - but not where there is a morally relevant difference. We are not free, of course, to call any difference we like as "relevant". So what counts as a morally relevant difference?

Earlier we noted (p.19) Korsgaard's attractive idea that the obligations issuing from our practical identities differ from the moral obligations that, generally, we have to all other people. Wouldn't we say with Korsgaard, for instance, that there is a *prima facie* relevant difference that justifies a mother treating her own children in a partial way (or her partner, parents, friends, etc.)? For her obligations to her own children are based on different considerations to her obligations to the countless other children in the world. She wittingly took on the role of 'being their mother' and, due to this role, she is morally obligated to adequately feed them, to clothe them, to talk and to play with them, to listen to their "jokes", etc. A mother who shows this special concern for her children is being 'a good mother', a mother who ignores it (completely) is a bad one. The idea here is that 'they are her own

children' counts as a relevant difference.³⁰ So that when it comes to, say, a choice of either feeding her own hungry children or feeding other unknown hungry children, a mother should give preference to her own. Along similar lines, we might argue that faced with the choice between serving with the French army or staying in occupied France to care for his ailing elderly mother (see pp. 26-27), Sartre's student has a *prima facie* special obligation to his ailing mother. As Williams (1981:15) implies '...because she is my mother...' is justification enough for showing partiality in this case. In ways like the above, we might argue that the special obligations entailed by our practical identities count as warranted exceptions when applying (iv).

Let me summarise my account thus far. I have argued that there are certain formal requirements that our action-guiding, autonomous moral judgements need to meet; they must be, universalisable and impartial (iv). And I pointed out how some judgements are fairly easy to justify - as being moral or immoral – by using these conditions. In many cases someone's action is ruled out as immoral by the universalisability condition (iii): What if everyone were to do that? e.g. my stealing from co-workers in order to finance my holiday. Similarly if my action clearly and substantially affects another person and I completely ignore their wishes, this is ruled out by impartiality (iv); e.g. I cannot knowingly infect other people with Covid, or drug a classmate in order to have sex with her. Needless to say moral issues where disagreement and conflict occur are usually far more difficult to figure out and moral decisions much more difficult to justify. However, even here I think there is more we can do.

I will argue now that there are further practical considerations, many of which are entailed by the formal ones, that an adequate justification of a moral judgement needs to meet. And, in some cases, paying attention to these practical factors can resolve difficult conflicts

³⁰ Unfortunately, there are no obvious practical identities comparable to the parent-child relation that obligates us to meet the needs of children in the Third World. Although like Singer (1979:229-242) I think that when the choice is between some trivial luxury item for one's own child and helping meet the basic needs of impoverished children, preference should be given to the latter.

CHAPTER FIVE

PRACTICAL REQUIREMENTS

§10 Practical entailments of formal conditions

(a) *The need to accurately identify the problem.*

Often difficulties arise because we misidentify what is morally at stake. Sometimes this is because we fail to consider all of the relevant details in the situation; or we fail to see which factors are morally relevant, or the weightiest. And sometimes the problem we identify initially may not be as straightforward as it at first appears, especially where there are conflicting values, beliefs, or options for action. Perhaps we even fail to see a moral dilemma - that we must choose between two (or more) competing possible actions and that in making a decision we must compromise the other possible choice(s). So before she makes a moral judgement, our judgment-maker needs a clear understanding of the moral issue (problem, challenge) that needs to be addressed. Let us consider briefly what this can involve.

(b) *The need to be adequately informed.*

Firstly, to accurately identify a problem, she needs to gather as much relevant information – non-moral facts - as possible about the troubling situation. Who are the people involved? What are their individual concerns? What assumptions are they making about this particular moral issue? Sometimes she will need to establish if the act in question was avowedly intentional, accidental, or due to negligence. Often the evident ill-will (or good will) of the person responsible for the act is relevant. Occasionally the relevant facts concern the perpetrator's expected state of knowledge; or their immaturity. Sometimes non-moral facts are needed to fix the time and place where the right/wrong act happened; sometimes they are needed to show that the act in question has been carried out on previous occasions. Our judgement-maker has to take into account also any of the moral commitments that might be relevant to the problem; e.g. an obligation entailed by a practical identity, or a promise already made, etc. Then she needs to reflect on any previous moral decisions that apply to cases of this kind. Why are they satisfactory/unsatisfactory? If the latter, what should she do differently?

She needs as much relevant information as possible *also*, to comply with universalisability (i)-(iv). She must know all of the relevant non-moral facts - or at least enough to make a well-educated hypothesis about the issue she is considering - just in case she

overlooks something that results in a universalised judgement that psychologically she is unwilling to accept. For instance, does her proposed solution break the law, or any of the usual social norms, or other generally held beliefs about acceptable and unacceptable behaviour? (In a multi-cultural society this is quite a difficult demand to meet.)

(c) *The need to be aware of the likely consequences.*

When our judgement-maker chooses to act she knows her actions will have consequences, not only for herself but also for others. So before making her choice (and then acting), she needs to have thought through the most likely consequences of her proposed solution. For instance, (to comply with the impartiality (iv)), she needs to consider how her proposed action is likely to affect the different interests of the people involved. Who is likely to be helped or harmed if it is implemented? And if someone will be seriously disadvantaged, does this matter? The point here needs a bit more unpacking.

(d) *The need for knowledge of those affected by one's judgement.*

Included in her appraisal of the most likely consequences, to comply with impartiality (iv), she needs to consider how her proposed course of conduct will appear to those affected by her action; the impact her decision will have on them. Like her, other people have (perhaps different) interests, beliefs, preferences, that she needs to consider - though she does not have to agree with them. The other person might well have expectations, perhaps stemming from the relationship they have with our judgement-maker e.g. the expectations of an elderly mother, or a promise she has made to her partner, etc. If she is going to see her proposed judgement from their point of view as well as her own, she needs to imagine herself in their shoes.

This is more complicated than it might seem. More than merely imagining herself in their shoes but still with her own beliefs, attitudes, feelings, to see her proposed action from *their* point of view, she needs to understand their perhaps different beliefs, emotional attitudes and feelings, as well as her own. For the same event can make different people angry, jealous, hurt, worried, pleased, excited, and so forth. Further, recognising what another person feels emotionally involves quite a range of skills. It involves her noticing - and understanding the likely significance of - things like facial expressions or bodily posture, observing in which circumstances a certain type of emotional state typically occurs, and so on. Does she still accept her proposed solution as being morally right when she takes all of the above into account?

(e) *The need to consider all of the possible solutions.*

With this background knowledge to hand, before making her decision, our judgement-maker needs to think through the range of possible solutions open to her - and in her mind trace the likely consequences of each of these alternatives. What are the projected outcomes for each potential solution? What reasons are there for prioritizing one alternative over the rest? Would a decision based on this or that possibility be impartial – has its predicted outcome any explicit or implicit bias? Has she considered each proposal in terms of her willingness to universalise it - could she consistently carry out a decision based upon it?

Even then, she is still not quite ready to decide on the best course of action. Simply complying with the formal conditions is not enough. She must be able to implement the judgement in question. This means she needs the appropriate emotional feelings, motive and resolution to carry it out.

(f) *The need for feelings compatible with the preferred solution.*

A problem with my account so far is that I am claiming that the use of certain formal conditions (universalisability, impartiality) are needed if we are to adequately defend our moral judgements. It could be argued that this emphasis on our reflective capacities leaves no room for the kinds of feelings or emotional reactions that are (always) constitutive of a moral response. Earlier we noted Hume's (1740:295): dictum: 'Reason is and ought only to be the slave of the passions and can never pretend to any other office...' When making a moral decision, we need at least to care about the thing or state-of-affairs we are judging. And 'to care about' and to willingly embrace an option is different from coolly determining it to be 'the best'. Thus a satisfactory account of moral judgement-making (though I think that Hume overstates the matter) must include the person making it having the appropriate feelings; the thing in question must matter to them. In other words, we need not only the ability to weigh the options impartially and to universalize them, we need also to engage with them in the right spirit - we need the appropriate feelings to do what we believe to be morally right. Thus the obligations of a parent, for example, requires not only a reasoned judgment but the care and concern they should feel when coming-to and implementing their decision

Incidentally, this seems to provide an answer to the difference between Phillips and Wolf we noted earlier (p. 14). I have argued that to be a defensible moral judgement, the judgement-maker must be willing to universalise her judgement X for all like cases. And now we have added the implication that this shows she cares about, or values,

or has dominant feelings in favour of X. She strongly believes that X is the right thing for anyone to do in this situation. Moreover, she thinks X should prevail over other competing moral or non-moral considerations; (otherwise, why universalise it?). Doesn't this imply (in support of Phillips) that she thinks that her moral judgement should override all other non-moral (personal, political, religious, scientific, etc) considerations that might also apply in the situation?

I have argued above that when she universalises her autonomous, impartial judgment, our judgement-maker needs the appropriate feelings, value or care about the matter concerned. . But, once again, there is more to the issue than this.

(g) The need for motive and resolve to carry out her preferred solution. When she is deciding upon X, she needs to be sure she can act on X, (Kant's formula 'ought implies can' is an obvious further condition that applies to the plausibility of her moral judgement.) More importantly, she needs to be sure that she is willing to act on X in a resolute way. It would be odd to call her conduct rational if, when she could, she did not purposefully do what she thought to be morally right. Otherwise it (the lack of resolve) looks like another case of someone claiming to have a moral belief that they do not really have (see p.25). And even when she has a positive motive and resolve with regard to it, there is still work to be done if she is going to persuade other people of the reasonableness of her judgement. For instance, perhaps she is bad at explaining herself.

§11 Pragmatic implications of the formal conditions

(h) The need to communicate effectively.

Another practical ability she will need, if she is going to use the formal rules effectively, is the ability to communicate her thoughts and feelings to others in an appropriate and sensitive manner. She could spoil things due to her ineptitude – by behaving in a gauche manner or by an uncouth action, or by lack of "know how". If our judgement-maker has difficulty in communicating with other people, her effectiveness - as someone who intends to do the right thing for the right reasons - is diminished. Or her argument could falter due to her lack of experience.

(i) The need for plenty of practice and experience.

Regularly thinking about everyday moral issues (from a young age) is a good way to improve our ability to identify the morally relevant features in a complex situation. We need plenty of practice to ensure, for instance, that our self-interest is not controlling the way we think, or that our moral decisions are really our own and we are not merely parroting

the views of others. Further, we need practice in thinking about the rationale underlying the various day-to-day ethical practices and recommendations which confront us, in order to discern what we take to be the best reasons, to confidently follow the formal guidelines when making an autonomous judgement, etc. In other words, we need plenty of practice and experience of the kind of thinking advocated in this essay. Just as importantly, we need to think about what I called grey area moral issues (e.g. war, euthanasia, the treatment of animals, etc.) and practice making our arguments better informed, clearer, more coherent, etc. And we need to do this frequently, albeit abstractly. After all, we expect practice and experience in almost all other areas where improvement or achievement is possible.³¹

Even then, there is no guarantee that our reasoning will persuade others; even though we meet the conditions outlined, many of our moral decisions will be met with dissent.

(j) *Moral decisions are always vulnerable.*

Other people often have different background assumptions - not stated or even known - which can lead them to deny X (with varying degrees of confidence). For instance, sometimes a (rational) interlocutor will accept the reasons A, B but not the conclusion X. Maybe they think the argument for X involves an unstated implicit reason(s) C, which the listener does not accept. Or perhaps they accept A and B and in normal circumstances would conclude X but, in the present case, they support an additional belief C, which they think prevails over A and B. They (the interlocutor) might say: "Well, normally X would follow, but in this case it does not because of C". Or, in the case of moral dilemmas (and other grey areas) the reasons considered seem to imply different incompatible possible judgements X and Y - and closer examination doesn't show that either X or Y has priority (see e.g. p.16). Often in such cases not making a judgement is not an option and my interlocutor and I choose differently. In which case, I should never be so confident in my judgement that I do not acknowledge the possibility that I could be wrong.

(k) *The need to learn from our mistakes.*

It follows from (j) that because many of her moral decisions, especially in grey areas, will be met with dissent, our judgement-maker should always be ready to consider and benefit from the responses of other people. In other words, she needs to be open-ended about difficult moral

³¹ For example, to successfully take a penalty in football, or to bat against an Indian spin bowler, or serve an ace at tennis, we expect the player to have practiced "in the nets" as much as possible.

questions and to learn from her mistakes.

The list (a)-(k) could no doubt be extended. There may be many other practical considerations that need to be taken into account if we are to effectively use the formal rules. But I hope that (a)-(k) will be sufficient for our purposes. When we are trying to judge morally, these are some of the skills and abilities entailed by the rules of autonomy, universalisability and impartiality.

CHAPTER SIX

CONCLUSION

In this dissertation, I started by claiming that what counts as a good defence of a moral judgement needs to be made clearer and that such a clarification could influence the way we resolve many of our moral differences. In Chapter 1, by way of clearing the ground, I gave a brief account of the structure of simple moral judgements; the kind of action-guiding judgements we make that affect other people, where praise or blame is appropriate. I pointed out that one obstacle to my approach are theories which base morality on absolute principles; where the best reason for a moral judgement is already determined. The main problem with so-called absolute values is that they are always open to challenge. In Chapter 2, I argued that for moral responsibility – praise, or blame - to be appropriately ascribed to a judgement-maker, her judgements and actions must be taken to be self-determined. The assumption here is that any rational person is able to reflect upon and decide for themselves about which action they should morally do. However, I argued they do not start from scratch in this. They can make autonomous decisions by interrogating the received values that are customarily used to decide conduct. Also in their relationships with others, the social roles a person chooses to fill (husband, father, *et al*) entail certain moral obligations. I went on to point out that whatever they choose to do (or continue to do), to qualify as a defensible moral judgement, this must conform to at least two (formal) conditions. In Chapter 3, I argued that one condition is universalisability; can the judgment in question be carried out consistently in all like cases. And I argued against those philosophers who question this condition. In Chapter 4, I added that to qualify as a defensible moral judgement one's decision should be impartial; and that my preferred account of impartiality recognises that when making it, the relevant interests of everyone affected by the judgement should be considered. And I argued against those philosophers who question this requirement.

In Chapter 5, I suggested some of the practical skills our moral reasoner needs to successfully implement their autonomous, universalised, impartial judgments. I pointed out that often their judgement-making is skewed due (a) to their failure to properly identify the problem; or (b) their failure to be sufficiently informed, or (c) to weigh the consequences of their preferred solutions, or (d) to identify the actual interests of those affected by their judgement; or to not having (e) thought through all of the alternatives that are open to them in the problematic context. I noted also that sometimes the formal procedure

breaks-down in practice due (f) to the judgement-maker's inability to emotionally engage with their judgement in the right spirit, or (g) to their not being sufficiently motivated to carry out the decision once made. Then I noted some of the pragmatic implications for the effective use of our formal conditions. They include (h) making sure that we have the "know-how" to explain our decision appropriately and agreeably to others; and, most importantly, (i) the need for plenty of experience in making decisions of this nature. Added to all of this, we need (j) to recognise that dissent is likely for almost all of our grey area moral decisions; and how we need (k) to honestly consider any objections that other people make and thereby to learn from our mistakes.

Lastly, it is worth emphasising that we should see universalisability and impartiality as procedures to be used when making a moral judgement rather than substantive moral positions. They do not tell us what our moral judgements should be but what it means to make a defensible moral judgment. Nonetheless if I have argued correctly (and well-enough) once we are in the habit of making our autonomous decisions in accordance with the formal conditions, it will not always be an open question as to what our moral judgements ought to be. We will be constrained to settle most moral questions in a manner which is an extended application of this kind of thinking.

BIBLIOGRAPHY

- Arendt, H. (1965) 'A Report on the Banality of Evil'. New York: Penguin
- Ayer, A.J. [1956] *Language, Truth, and Logic*. London: Gollancz Ltd
- Baier, K [1958] *The Moral Point of View: a Rational Basis for Ethics* Ithaca. New York.
- Bentham, J [1789] *Introduction to the Principles of Morals and Legislation*. Classic reprint.
- Black, M [1946] *Critical Thinking: An Introduction to Logic and Scientific Method*. Prentice-Hall Inc. New York.
- Copi, I [1953] *Introduction to Logic*. Macmillan. London
- Firth, R [1952] 'Ethical Absolutism and the Ideal Observer' in *Philosophy And Phenomenological Research*. Vol 12 [3] pp 317- 345
- Foot, P [1958] 'Moral Arguments' in *Mind* 67 pp 502-513
- Foucault, M [1984a] 'The Ethics of the Concern of the Self as a Practice of Freedom'. In Rabinow, P ed [1994] *Essential works of Foucault 1954-1984 vol 1 Ethics*. Penguin. London. P281- 301
- [1984b] 'The Return of Morality' in *Michel Foucault Politics, Philosophy, Culture: Interviews and Other writings 1977-1984* ed Kritzman, L , trans Sheridan, A Routledge. New York.
- Frankfurt, H [1971] 'Freedom of the Will and the Concept of a Person' *The Journal of Philosophy* vol LXVIII no 1 pp 5.20
- Hare, R (1952] *The Language of Morals*. Clarendon Press. Oxford.
- Harmen, G [1977] *The Nature of Morality: An Introduction to Ethics*. Oxford University Press. New York.
- Hobbes, T. [1651] *Leviathan* Everyman's Library 1949.
- Hume, D [1740] *A Treatise of Human Nature*. Dover Philosophical Classics
- [1751] *Enquiry Concerning The Principles of Morals*. Ed H Aiken Hafner Publishing. Connecticut. 1970
- Kant, I [1785] *The Groundwork of the Metaphysics of Morals*. [1785] Trans H J Paton as *The Moral Law*. Hutchison. London [1948]
- Korner, S [1955] *Kant*. Penguin Books. London.

- Korsgaard, C [1996] *The Sources of Normativity*. Cambridge University Press New York.
- Leopold, A [1949] 'The Land Ethic' in M. A. Zimmerman *Environmental Philosophy: From Animal Rights to the Radical Ecology*. Prentice Hall Inc. New York . 1998
- Locke, D [1968] 'The Trivializability of Universalizability.' In *The Philosophical Review* (77) pp.25-44.
- Macintyre, A [1981] *After Virtue*. Duckworth and Co. London.
- Mill, J.S [1861] 'Utilitarianism' in *Utilitarianism, On Liberty and Considerations on Representative Government*. Ed H B Acton. Dent and Sons. London. [1972]
- Moore, G.E [1903] *Principia Ethica*. Prometheus Books. Great Books in Philosophy Series. London.
- Nietzsche, F [1888] *Beyond Good and Evil*, trans R. J Hollingdale. Penguin Classics. London
- Phillips, D.Z (1977) 'In Search of the Moral Must. Mrs Foot's Fugitive Thought'. *The Philosophical Quarterly* vol 27. pp 140-157
- Popper, K [1945] *The Open Society and its Enemies vol 1*. Plato Routledge. London
- Rawls, J. [1971] *A Theory Of Justice*. Belknap Press Cambridge Mass.
- Sartre, J. P [1948] *Existentialism and Humanism*. Methuen. London
- Scruton, R [2003] *Animal Rights and Wrongs*. Continuum. London
- Sidgwick, H. [1874] *Methods and Ethics*. Hackett Classic 1981.
- Singer, P [1972] 'Famine, Affluence and Morality'. In *Philosophy and Public Affairs* vol 6 pp 229-243.
- Sparrow, R [2007] 'Killer Robots' in *Society for Applied Philosophy*. Vol 24, pp. 62-77
- Williams, B [1981] 'Persons, Character and Morality' in *Moral Luck* Cambridge University Press. pp 1-18
[1985] *Ethics and the Limits of Philosophy*. Cambridge University Press..
- Wolf, S (1982) 'Moral Saints' *Journal of Philosophy* vol 79 pp.419-439