

A Practical Guide to Ethical Theory

Richard Field

© Richard Field 1999

Table of Contents

I. Ethics	2
II. Moral Reasoning: Principles and Problems	8
III. Theoretical Normative Ethics	22
IV. Teleological Ethical Theory	27
V. Deontological Ethical Theory	39
VI. Applying Ethical Theory: Some General Suggestions	53

I. Ethics

This course deals with the area of philosophical discussion and inquiry known as "ethics" or "moral philosophy." As these terms suggest, the primary focus of this area of inquiry is issues that arise in ethical or moral situations, situations which raise questions concerning what we ought or should do when the issue is not purely a matter of self-interest, but of right and wrong. An ethicist attempts to bring some clarity of thought to these issues: to define clearly the language that is used to discuss them, to reveal the forms of inference that underlie our reasoning about them, and to determine and justify principles that can provide guidance in resolving these issues by bringing into consistency our best thoughts and intuitions on these matters.

Ethics is therefore reflective and critical: it does not attempt simply to codify commonly accepted beliefs concerning moral questions, but attempts to find the grounds for reasonable moral beliefs. Because of this, ethics should be distinguished from what we commonly call "morality," which is simply a set of beliefs accepted by a given culture concerning what a moral agent ought or ought not to do in moral situations, whether these beliefs are a product of critical reflection or not. Although many people do at times consider the cogency of their moral beliefs in a critical manner, the philosophical ethicist attempts to do this in a manner that is broader in scope and more systematic in methodology than is typical in common moral reflection.

Why Study Ethics?

The aim of ethics, it should be noted, is not to give people good character. Research in developmental psychology suggests that factors such as family environment, parental disciplinary styles, peer group reinforcement and others are the basis for the development of what we commonly call "good moral character." The study of ethics can have no more influence on a person's character than the study of mathematics can alter a person's basic quantitative intelligence. What the study of ethics can do is to provide some understanding of basic ethical principles, and strategies of moral reasoning, that can be used in discussion and debate in support of positions on moral issues. It is not enough simply to have a viewpoint on a moral issue to enter into intelligent discussion on the moral issue, any more than it is enough simply to have a theory in order to do science. The more that is required in science is some evidence to back up the theory. In moral discussion, what is required is a reasonable analysis of the moral situation, and a considered application of moral values and principles that together suggest why the viewpoint should be accepted. Without this rational basis for offering the viewpoint, moral discussion offers nothing better than a dogmatic insistence that a given view is correct--a fruitless endeavor.

In this course, we will be considering some of the elements of good ethical reasoning, and the application of these to some contemporary issues of moral concern. We will not attempt to discover what the "truth" is on any of the particular moral questions we consider. Rather, the aim of the course is to provide an understanding of good moral reasoning that is sufficient to allow the student to enter into ethical discussions in an intelligent and reasonable manner. This endeavor presupposes a certain attitude on the part of the members of the course, an attitude that involves

two components: (1) a tolerance for other people's moral viewpoints (again, dogmatic insistence on an ethical position cuts off moral discussion rather than promoting it), and (2) an expectation on the part of one's own contributions to moral debate and discussion as well as those of others that viewpoints should be offered with justifying reasons for those viewpoints.

The Source of Moral Values

Viewpoints concerning the source of moral values have been and continue to be quite diverse. Religious ethics locates this source in an all-good god, who reveals through scriptural writings the rights and wrongs by which we should live, and ultimately judges our lives on the basis of the degree to which we conform to these standards in our conduct. A modern trend in ethical thought, which can be traced back to the writings of Thomas Hobbes and John Locke, understands moral values as based in the natural order in some fashion. Thus moral rights and duties can be discovered by considering the "natural manner" in which people interact with each other and their physical environment.

Who's right? Happily we do not have to decide the issue, notwithstanding its importance, since the critical evaluation of moral beliefs that is the aim of ethical inquiry does not require that we determine the source of moral values, but simply that we acknowledge that there are moral values and that a reasonable application of these values requires that they be applied in a consistent manner.

To understand what a consistent application of moral values is, consider an inconsistent application. The issue is slavery--a moral issue that has long been decided, but one that was the center of active debate in the early nineteenth century in this country. We commonly believe that people should be treated equally, that is, that people should be accorded the same respect in the manner in which we interact with them regardless of who they are, what their background is, and so forth. Slave holders in the early nineteenth century would typically accept this standard with respect to their family and the community around them, but not with respect to their slaves. This suggests an inconsistency in the application of the standard of equality of treatment unless there is some other relevant factor that requires that slaves should not be treated equally. One such factor commonly cited was intelligence: a person of African descent, it was thought, is not as intelligent as someone of European descent, and thus they could not be educated and could not live independent lives. The fact, however, that when given the opportunity slaves did learn as well as their slave holders' children, and that some of them demonstrated their capacity for great intellectual accomplishments, such as the famous author and abolitionist Frederick Douglass, showed that the claim of the slave's alleged intellectual inferiority was false. What we have, then, is an inconsistency of the application of a moral value, that of equality--an inconsistency that is easily discovered if the relevant moral value, equality, is articulated in clear language, and then applied to the manner in which slave holders treated their slaves.

Ethics, then, does not require that we determine the ultimate source or basis of moral values. It is enough that we admit that we do apply common moral values to our own conduct and the

conduct of others, that these moral values can be stated in clear and unambiguous language, and that they can then be applied consistently in different moral situations as guidelines to action, or, as above, can be applied to actual conduct to discern the inconsistencies of their application.

Ethical Relativism

But what if there are no common moral values? What if our moral beliefs are neither right nor wrong, but are simply "opinions" with no more objective validity than such matters of individual preference as whether someone likes butter or sour cream on their potato? The position that makes this claim is commonly known as "ethical relativism." If the ethical relativists are right, then there are no common standards of moral judgment, and the ethicist, who would try to define such standards, is working in vain.

There are a couple of different versions of ethical relativism. One version, individual relativism, claims that moral standards are relative to individual moral beliefs--that in effect what one believes is right is right with respect to one's own actions--and thus no sound moral argument condemning the actions of a moral agent can be offered if the conclusion of the argument is inconsistent with the moral beliefs of the agent. A typical way in which this position is expressed is the following: "everyone has different moral views, and will make different moral judgments under the same circumstances, so there's no point in trying to find a common moral standard because there is none." Another version of ethical relativism is so-called cultural relativism. According to this view, moral standards are relative to cultural moral beliefs. Thus although one can justifiably condemn the actions of a person that are inconsistent with the moral beliefs of that person's culture, no condemnation of the actions of a person of a certain culture can be justified if those actions are consistent with the moral beliefs of that person's culture. Typically this position has been offered to support the claim that we should respect the beliefs and practices of other cultures in order to avoid allowing our own cultural beliefs to be the basis of unfair assessments of these cultures. Thus the position might be expressed in this manner: "various cultures throughout the world have different moral standards, and it is simply a prejudice of our own cultural perspective to find fault with these views."

Although there are a number of good arguments that can be offered against ethical relativisms of either sort, one of the most straightforward ways of suggesting the implausibility of ethical relativism is to note the implications of the view. It is particularly difficult to be a consistent relativist when one is harmed or injured by others. I will offer here two hypothetical situations, one corresponding to each of the versions of ethical relativism, and state the implications of the respective versions in these cases. The extremity of the actions considered in these cases is meant simply to make the difficulty of accepting relativist conclusions clear; similar difficulties would arise in the consideration of less extreme moral offenses. Place yourself in these situations and consider whether you could accept the implications of ethical relativism.

Individual Relativism

The Situation: A man abducts you and your family for the purposes of his own sadistic enjoyment. He takes you and your family to a secluded farm house, and tortures your family members for five days before they finally die of the severity of their injuries. You miraculously escape just before he is about to do the same to you.

The Implications of Relativism: So long as the assailant in this case does not believe he has done anything wrong, he has indeed done nothing wrong. You have no basis to complain about his actions, even if he acted with full knowledge that you and your family objected to his actions, and that society views such actions as grossly immoral. Although you might wish to avenge his actions by harming him, if you do so your actions are no more praiseworthy (or blameworthy) than his. You may appeal to the law, but no court would be justified on a moral basis to convict or imprison the assailant. If this is done, the actions of the court are no more praiseworthy (or blameworthy) than those of the assailant. In fact, a legal system which would condemn the actions of the assailant would be no better on moral grounds than one that allowed such actions, or even one that would reward the assailant for his actions by paying him \$3 million from public funds.

Cultural Relativism

The Situation: There is a country, Xenophobia, that contains two distinct cultural groups: the majoritarians and the minoritarians. The majoritarians make up 90% of the population, the minoritarians the remaining 10%. There is a long history of suspicion and hatred in majoritarian culture against the minoritarians. According to majoritarian cultural history, the majoritarians are a culturally superior race that were the original inhabitants of Xenophobia, and they therefore see minoritarians as inferior and usurpers of their native lands. You are a minoritarian. The fragile truce between the majoritarians and minoritarians disintegrates over a murder case in which a minoritarian is charged with the assassination of a prominent majoritarian politician. There is a quick move in the legislature of Xenophobia to round up all minoritarians and place them in internment camps and to suspend judicial procedures in all cases where minoritarians are charged with crimes. The bill is quickly approved, and you and your family are interned. Your brother was charged with a robbery some years ago, and because the repressive law against minoritarians is retroactive, the commander of the camp takes your brother to the firing range where he is executed by firing squad.

The Implications of Relativism: So long as the actions taken by the majoritarians are consistent with their cultural beliefs, there is nothing wrong with their actions. You have no moral basis upon which to complain about their actions, even if the majoritarians understand that world opinion condemns their actions. You and other minoritarians might wish the world community (e.g., the United Nations) to place economic sanctions against the majoritarians, or even intervene militarily, but if it did, these actions against the majoritarians would be no more praiseworthy (or blameworthy) than those of the

majoritarians. Alternatively, you and other minoritarians might believe that you are justified in taking up arms against the majoritarians, if you can manage it. But an armed revolt against the majoritarians would not be just or right in any significant moral sense. Morally speaking, such an action would be no more meritorious than the repressive actions of the majoritarians to which you react.

If you find these implications objectionable then you cannot accept ethical relativism.

Although there are theoretical reasons why some philosophers (a small minority historically) have accepted moral relativism, most people who adopt relativism do so on the basis of reasons that involve unguarded judgment or simple conceptual confusions. As I noted above, some people in the past have adopted cultural relativism as a basis for insisting that people should respect the beliefs and traditions of other cultures, even if they are not consistent with those of one's own culture. But ironically cultural relativism undermines any moral basis upon which one might justify this normative claim. Upon a relativist view, if the people of a given culture believe that it is quite proper to annihilate another culture completely, there is no justification for objection.

Another common reason that people adopt ethical relativism is the erroneous view that if one rejects relativism, then the only alternative is to accept a position that they find unacceptable, commonly known as moral absolutism: the view that there are a certain set of moral rules that apply in all relevant cases without exception. According to moral absolutism, for example, lying to others is always wrong, even in cases where, for example, not lying would result in great harm to others. But rejecting ethical relativism does not require the acceptance of moral absolutism, but the acceptance of a broader position commonly called moral objectivism or moral realism. This view simply makes the claim that normative judgments can be justified on objective (nonrelativistic) grounds, but a moral objectivist is not necessarily committed to the claim that moral standards are absolute. Common moral standards can be understood to apply in different ways depending on the circumstances of particular moral situations. So a lie might be justified under certain circumstances.

Sometimes people embrace relativism as a reaction against traditional moral beliefs that they regard as oppressive and odious. It certainly is the case that people in the past, as many people of the present, have held certain moral beliefs that from a more enlightened point of view are clearly objectionable, such as the moral permissibility of slavery, the moral impermissibility of treating women as equals, and the like. But the rejection of ethical relativism does not imply that traditional moral beliefs must be accepted. In fact, the acceptance of traditional beliefs for no better reason than that they are traditional is a fallacy of moral reasoning, and quite antithetical to any reasoned moral position. Furthermore, the claim that traditional moral beliefs are wrong is a normative one, and thus can only be justified from a nonrelativist or objectivist standpoint. So those who embrace relativism as a reaction against moral conventions undercut any grounds upon which they can justify their reaction.

Another reason some people adopt relativism is from a pessimistic skepticism concerning the possibility of any definitive resolution of moral issues. They hear different views voiced on contemporary issues, all supported by what appear to be convincing arguments, and throw up

their hands in frustration, concluding that no truth can be found: "it's all a matter of how you look at it." But this reaction is premature: there is no reason to believe that because there are different and quite reasonable views on important moral issues, and the debates have yet to be resolved, that therefore they are irresolvable. Most important issues in any area of serious discussion--in science, mathematics, political thought--are resolved only after a time, often a lengthy time, of vigorous debate. There is no reason to believe that ethical issues should be any easier to resolve; in fact, the complexity of these issues suggests that they should be quite difficult to resolve in any easy manner. Moral skepticism also loses sight of the fact that historically many moral issues that were hotly debated in the past have been settled in ways that no reasonable person would object to today: that slavery is immoral, that women should have the same legal rights as men, that racism is immoral, that religious freedom should be protected by law, etc. Moral argument played a crucial role in the resolution of these issues. But from the point of view of moral relativism, we would have to conclude that the adoption of these beliefs in fact effected no advance in moral understanding--that if these beliefs had been rejected, our society would be none the worse for it.

The common reasons, then, offered for adopting ethical relativism are, on close examination, quite unconvincing, and in practice it is quite difficult for any reasonable person to adopt the view consistently. At any rate, in this course we will accept moral objectivism as the basis of our discussions: simply put, that moral positions can be justified on the basis of rationally supportable and objective criteria.

II. Moral Reasoning: Principles and Problems

We have seen in the previous chapter that ethicists attempt to move beyond the simple acceptance of ethical viewpoints--they attempt to offer reasons that recommend the general adoption of these viewpoints. The process whereby a given belief or judgment is supported by reasons that recommend its general adoption is known as "justification." Moral reasoning is simply one form of justification. In this chapter, we will consider first some basic terms and concepts that apply to all forms of justification, and then look more closely at special principles and common errors that pertain especially to moral justification.

Justification and Logic

Ethics, as any form of rational inquiry, is concerned with establishing beliefs or judgments that can be justified in an objective manner. The rational justification of a belief or judgment is expressed in the form of an argument. (Common synonyms for "argument" are "proof" or "demonstration.") In this context, the term "argument" does not refer, as it commonly does in ordinary discourse, to a debate or disagreement between people (although arguments can and commonly do enter into debates and disagreements), but rather refers to the process by which reasons are offered that recommend the acceptance of a given belief as true. Logic is the study of the fundamental principles that enter into the determination of the quality of an argument: the extent to which the reasons offered in support of a belief actually do provide support for the belief. We will consider here some basic terminology that applies in analyzing the structure of arguments, and determining their quality.

The Structure of Arguments

Any argument can be analyzed into two basic elements: (1) a conclusion, which is a statement that expresses the belief which is supported by the argument; and (2) one or more premises, statements (or propositions) which express the reason or reasons offered in support of the claim that the conclusion is true. For example, consider the following argument.

It is clear that the world had no beginning, since if it did, there would have been a first moment before which there was no other moment, but every moment is preceded by some previous moment.

In this argument, the conclusion is "It is clear that the world had no beginning," and two premises are offered in support of this conclusion: (1) "if it did . . . other moment," and (2) "every moment . . . previous moment." The premises are indicated by the word "since." Other common premise

indicators are the words and phrases "because," "for," "for the reason that," and "on account of the fact that." The phrase "It is clear that" indicates the conclusion. Other words or phrases that indicate conclusions are the following: "therefore," "thus," "so," "it follows that," and "we can conclude that." At times when arguments are offered in ordinary discourse, premise or conclusion indicators are not used. When this happens, identifying the conclusion requires that we determine the statement in the argument that the author of the argument is trying to support.

It can be helpful to make the structure of an argument explicit when writing it out. When this is done, the convention in logic is to write out each premise as a separate statement, often numbered or lettered, followed by the conclusion. Sometimes a line is drawn between the premise(s) and conclusion so that their status in the argument is clear. If we do this with the argument above, we get the following:

- (1) If the world had a beginning, there would have been a first moment before which there was no other moment.
 - (2) Every moment is preceded by some previous moment.
- The world had no beginning.

Notice that when we put the argument in this form, the premise and conclusion indicators can be removed. For clarity's sake, we also interpreted the phrase "if it did" in the original expression of the argument in terms of its antecedent "the world had a beginning" when restating the first premise. Other conventions apply in some forms of logic that determine the order in which the premises are written, but we will not consider these conventions here. At times in this course we will express arguments in this conventional form to make the logical relationship between the premises and the conclusion clear.

The Evaluation of Arguments

There are two standards by which the logical quality of an argument (the degree to which the premises support the conclusion) can be determined: deductive and inductive.

Deduction

The deductive standards are called validity and soundness. An argument is valid if on the assumption that the premises are true, the conclusion must be true (cannot be false). An argument is invalid if on the assumption that the premises are true, the conclusion still could be false. For example, consider the following argument.

- (1) If the moon is made of green cheese, then the moon is edible.
- (2) The moon is made of green cheese.

The moon is edible.

This argument is deductively valid, which is to say that if premises (1) and (2) were true, the conclusion would have to be true. Notice that the fact that premise (2) is false does not effect the validity of the argument, because validity has nothing to do with the actual truth or falsity of the premises or the conclusion.

The fact that valid reasoning is not dependent on the actual truth of the premises of an argument allows for the possibility of hypothetical reasoning: by assuming that something is the case (that is, by making a hypothesis), we can discover what also must be the case. For example, we can assume that the moon is made of green cheese and discover that this implies that it is edible. Hypothetical reasoning is a crucial tool in ethics as well as the natural sciences and mathematics.

Of course, if we want our deductions to provide real understanding, the premises that we start with must be true. When a deductively valid argument has premises that are actually true, it is called a sound argument.

Induction

An argument may fail to meet the standard of deductive validity and yet still be a useful argument in that the premises provide some support for the conclusion. For example, no one could reasonably deny that the following argument provides some support for its conclusion.

(1) From the dawn of recorded history, it has been reported that the sun rises above the horizon once every 24 hours.

(2) The circumstances that might prevent the sun rising over the horizon once every 24 hours (the sun exploding, the Earth leaving its orbit), are highly improbable.

The sun will rise tomorrow.

This argument is deductively invalid, but by another standard of quality, the inductive standard, it is quite a good argument. The inductive standard of quality is known as strength. A strong argument shows that the conclusion is highly likely to be true given the truth of the premises. In other words, the premises provide strong support for the conclusion. A weak argument, on the other hand, would provide very little support for its conclusion. Notice that whereas the deductive standard defines only two values in assessing the quality of an argument--either an argument is valid or invalid--the inductive standard defines a continuum of values, since there is a whole spectrum of degrees to which the premises of an argument support an argument, from very strong to very weak arguments.

The Nature of Moral Reasoning

Like any form of reasoning, moral reasoning attempts to establish that some proposition or claim (a conclusion) is true on the basis of premises which support its truth. The essential difference between moral reasoning and other forms of reasoning is that the conclusion supported by the reasoning is a moral judgment. There are basically two forms of moral reasoning: (1) reasoning that attempts to demonstrate the truth of some general moral principle, and (2) reasoning that attempts to establish the truth of particular moral claims on the basis of general ethical principles. We will consider some examples of the first form of moral reasoning when we study theoretical normative ethics and consider some arguments offered by modern ethicists in support of a variety of ethical theories. But for the most part we will be considering moral reasoning of the second type, a form of reasoning sometimes called "moral deliberation."

Moral Arguments

Joan Callahan offers a good account of the basic structure of moral arguments.¹ One premise of such an argument will state a general moral principle (GMP) that provides the normative criterion that is used by the argument. A second premise states a factual claim or claims (FC), and the conclusion states a derivative moral judgment (DMJ) that is more specific in application than the principle stated in the first premise. This structure is deductive in nature, since if the GMP is true, and the FC correctly applies the moral criterion of the the GMP, then the conclusion must be accepted. The general form of the argument can be represented in this fashion:

(1) Any action having properties A, B, C . . . is right (or wrong).	GMP
(2) <u>This action has properties A, B, C.</u>	FC
This action is right (or wrong).	DMJ

The properties cited constitute the moral criterion defined by the principle. For example, the principle of honesty states that any statement of a falsehood that is offered with the intention to deceive is wrong. The properties of the action in this case are that the action is a statement, that the statement is false, and the action is intended by the agent to deceive someone.

Although this structure seems quite simple, moral reasoning can be quite complex, involving not one but several moral principles, a variety of factual claims, and at times a string of arguments that leads to the final conclusion. All of these elements can be placed in a variety of orders, making it difficult to sort out. Consider this argument, for example, from John Stuart Mill's *On Liberty*.

¹ Joan C. Callahan, ed., *Ethical Issues in Professional Life* (New York: Oxford University Press, 1988), p. 14.

II. Moral Reasoning

. . . Even opinions lose their immunity [to moral sanction] when the circumstances in which they are expressed are such as to constitute in their expression a positive instigation to some mischievous act. An opinion that corn-dealers are starvers of the poor, or that private property is robbery, ought to be unmolested when simply circulated through the press, but may justly incur punishment when delivered orally to an excited mob assembled before the house of a corn-dealer, or when handed about among the same mob in the form of a placard. Acts, of whatever kind, which without justifiable cause do harm to others, may be, and in the more important cases absolutely require to be, controlled by the unfavorable sentiments, and, when needed, by the active interference of mankind.²

In this argument, the conclusion is stated first: that opinions may be sanctioned by blame or punishment when their expression under particular circumstances constitutes "a positive instigation to some mischievous act," or, in other words, an incitement to others to cause harm. The general moral principle of the argument is stated last: that actions which without justification do harm to others may and in some cases should be controlled by "the unfavorable sentiments" (i.e., disapproval) or "active interference," such as punishment. The factual link between the general moral principle and the derived moral judgment is the claim that expressing opinions under certain circumstances can inflame peoples' passions, which can lead to harm. But notice that Mill doesn't simply state this claim, but offers an example that he believes makes the point. This example is stated, furthermore, in the context of a more specific moral argument applied to this example, an argument that uses the same general moral principle as the earlier argument. Thus Mill's statement "many justly incur punishment when delivered orally to an excited mob assembled before the house of a corn-dealer" can be analyzed in the context as the following moral argument (I've restated some of Mill's claims for the sake of clarity).

- (1) It is morally permissible, and in some cases morally obligatory, to sanction actions that incite others to do harm. (GMP)
 - (2) Expressing the opinion that corn-dealers are starvers of the poor will likely incite others to do harm when delivered to an excited mob assembled before the house of a corn-dealer. (FC)
-
- Those who do this justly incur punishment. (DMJ)

Thus Mill's argument is complex, interweaving two arguments--one main argument, and a more specific argument applying the same general moral principle to a particular case.

Analyzing and understanding moral arguments can, then, be a complicated affair. Some rules of thumb apply to the interpretation of moral arguments offered in discursive form.

1. Try to identify first the conclusion (DMJ) that a writer or speaker is attempting to support by his or her argument.

² John Stuart Mill, *On Liberty*, ed. Alburey Castell (Arlington Heights, IL: AHM Publishing, 1947), p. 55.

2. Next, answer two questions: (a) What general moral principle (GMP, a general moral claim other than the conclusion) does the writer or speaker offer as a basis of the argument? (b) What fact or facts does he offer that are relevant to the conclusion of the argument (FC)?
3. If either the general moral principle or relevant factual statements are missing, try to state these explicitly. Typically if a writer or speaker does not state these explicitly, it is because he or she believes that they are either so widely accepted that they need no argument, or that at least they are easily understood by the audience of the argument.

Formal Principles of Moral Reasoning

There are many criteria that can be applied to assess the validity of moral reasoning, but there are two formal principles of moral reasoning that are particularly important and fundamental. These principles are "formal" in the sense that they concern the form or structure of moral reasoning rather than its content.

The first of these principles is quite easily grasped from the analysis of the structure of moral arguments that we just considered. All moral arguments must include both a statement of a general moral principle (a normative claim) and a statement of relevant facts (descriptive claims) if they are to be valid. Put in another way, no derivative moral judgment follows from simply a description of facts. This principle of moral reasoning was first clearly stated by an eighteenth century British philosopher by the name of David Hume. Hume noted that there is a logical barrier to deriving claims concerning how we ought to act from descriptive claims concerning the way things are. This barrier has come to be known as the "is/ought gap" in ethical thought. One cannot validly claim that one, for example, ought to respect others merely on the basis of the factual claim that respecting others assures that they will treat us well. We must add some normative standard to the effect that being well-treated by others is a good that we ought to seek. The naturalistic fallacy considered below is one common way that this principle of moral reasoning is not observed.

A second basic principle of moral reasoning is that such reasoning must be consistent: we cannot hold inconsistent moral positions in different situations. This principle of moral reasoning is typically called "the principle of universalizability." This principle can be stated as follows:

If one judges that an action is right (or wrong), one is committed to judging that any other action that is like the first action in all morally relevant respects is also right (or also wrong).

This principle simply states that rational moral judgment must be consistent over all cases of actions that are similar to one another in respects that are relevant to moral judgment. If, for example, one judges that it is wrong to cause other human beings great physical pain because

pain is inherently bad, then one is committed to saying that it would be equally wrong to cause nonhuman animals great pain for the same reason. If one judges that this is wrong in the case of human beings but permissible in the case of animals, and one cannot spell out how human pain is in a relevant way different from animal pain, then one's moral position is inconsistent and thus fundamentally irrational.

The principle of universalizability is the basis of all moral reasoning. If this principle were not accepted, then moral reasoning would be impossible. Could we reject this principle? To do so would be to, in effect, deny that any moral judgment is more correct or true than any other, that there are no moral truths that can be accepted on rational grounds, and it seems quite impossible for us in a practical sense to believe this, as we saw in the discussion of ethical relativism in the previous chapter. When other people do harm to us, or treat us in ways that accord us little respect as human beings, it seems quite impossible for us to simply shrug our shoulders and say "I have no reason to complain about this." Even if we might question the principle of universalizability in theory, we all accept it in practice.

Fallacies of Moral Reasoning

A fallacy is an error in reasoning which prevents the reasons offered for accepting a certain conclusion from providing any real support for the claim that the conclusion is true. In other words, a fallacious argument seems to offer reasons for accepting a conclusion as true, but on closer examination it becomes clear that the argument provides no good reason for doing this. Logicians have identified many such fallacies--fallacies that arise quite commonly in debates and discussions of important issues. Below, I will concentrate on a number of fallacies that arise most commonly in these discussions.

The Naturalistic Fallacy

The naturalistic fallacy is a result of not recognizing the significance of the distinction between normative and descriptive claims, and the necessity of the former in any cogent ethical argument. The fallacy is committed when someone offers simply a description of "natural" facts as the basis for accepting a derivative moral judgment, without citing normative principles as well. Such an argument is never cogent, since descriptions of facts never imply the truth of some moral judgment. Consider, for example, this moral argument against disciplining children.

Punishing a child for some wrongdoing is never right, since punishment causes emotional pain.

The conclusion is a moral judgment, "punishing a child for some wrongdoing is never right," and in support of this conclusion a fact is cited, "punishment causes emotional pain." The fact itself,

II. Moral Reasoning

even if we admit that it is always true, does not show that the conclusion is true. The argument is based on an unarticulated normative claim, namely "Causing emotional pain is never right." But once this claim is made explicit, a number of questions are raised which the original argument doesn't address: "Is emotional pain always bad?" "Isn't it the case that in some circumstances linking a punishment with wrongdoing can inhibit wrongdoing in the future?" Since the original argument doesn't state the normative claim, it avoids raising some very important issues that must be addressed in order to evaluate the truth of the conclusion.

One recent moral debate where the naturalistic fallacy is often encountered is in the abortion debate. Consider these parallel arguments, one supporting pro-choice, the other supporting the immorality of abortions in the first trimester.

Pro-choice: Science has shown that the fetus in the first trimester lacks the ability to u, v, w, therefore it is morally permissible to abort a fetus in the first trimester.

Antiabortion: Science has shown that the fetus in the first trimester has the ability to x, y, z, therefore it is immoral to abort a fetus in the first trimester.

Neither of these arguments support their conclusions, since the facts that science uncovers concerning the abilities or inabilities of a first trimester fetus simply do not, by themselves, imply that abortion is right or wrong. What is needed, here, is some consideration of the normative question of what sorts of abilities a living thing must have for it to be regarded in a moral sense to be what in ethics is called a "person," someone who is a holder of the right to life. Thus, with respect to the first argument, a further premise is required that says "Everyone who has the ability to x, y, z, is in a moral sense a person," and the second argument requires a premise that says "Everyone who lacks the ability to u, v, w, is not in a moral sense a person."

Other examples of the naturalistic fallacy:

It is simply a fact that in our species, homo sapiens, sexual reproduction requires the mating of a male and a female, and it is quite clear that the diversity of two genders developed through evolution as a more advantageous way to perpetuate species. Therefore, it is clear that homosexuality, which fails to perpetuate the species, is morally wrong.

[Comment: Variants of this argument have often been offered as a basis for a moral objection to homosexuality, but the reasoning is quite obviously fallacious. It simply does not follow from the fact that gender developed in evolutionary history as a basis for reproduction that homosexuality is morally wrong. To make this argument, one would have to maintain that any employment of physiological traits of the human species that is inconsistent with the employment that favored their perpetuation in evolutionary history is morally wrong. On this basis, one would have to conclude that playing the violin is morally wrong, since clearly hands did not develop in primate evolutionary development to allow for the playing of musical instruments.]

Men need to feel dominant and to let off steam at times. So it is understandable that men get aggressive, even violent, at times. It's simply part of their nature.

[Comment: Even if it is true that men are more aggressive naturally than women, a claim often cited but as yet never proven, this would not mean that aggressive actions on the part of men are morally permissible.]

Tu Quoque (pron. too kwokwee; Latin for "you also")

This fallacy is a quite familiar one, often used by those who are attempting to absolve themselves of responsibility for moral wrongdoing. This fallacy is committed when someone suggests that a certain action is not morally objectionable because others have done the same or similar things. Typically the person who offers this argument is someone who has been accused of moral wrongdoing, and it is directed against that person's accusers, as if to say "you cannot accuse me of wrongdoing, since you are guilty of wrongdoing as well." The argument is clearly fallacious, since the question of whether others are guilty of wrongdoing is irrelevant to the question of the accused person's responsibility for moral wrongdoing: as the old saying goes "two wrongs do not make a right."

Examples of *tu quoque* arguments:

Member of Congress: I'm not going to try to tell you that I didn't take illegal campaign contributions. I freely admit this. But if every person who ran for public office who took illegal campaign contributions were kicked out of office because of it, the halls of Congress would be considerably more empty than they are now.

"[Soviet] Minister of Culture Yekaterina Furtseva publicly berated an American correspondent for 'poking his nose into our [Soviet] internal affairs' when he asked a question related to the case of disgraced novelist Alexander Solzhenitsyn [a Russian writer who was a critic of the Soviet government in the 1960s and 70s]. If you cannot punish the killers of your government leaders, you have no right to be interested in such questions,' the culture minister retorted.

[Source: Hartford Courant (December 20, 1970)]

Moral Conventionalism

Children often attempt to convince their parents that there is nothing wrong with something that they have done or wish to do because "all the kids are doing it." This is an instance of a very familiar fallacy known as "moral conventionalism." More generally, one commits this fallacy if one suggests that common or conventional practices in society provide the normative standard by which the moral status of actions should be judged, thereby suggesting that if people in one's

society commonly perform some action, the action is morally permissible. One problem with moral conventionalism is that often it is committed with exaggerated claims concerning the prevalence of the action that is adjudged to be morally permissible. When children, for example, say "all the kids are doing it," this often is quite untrue. But there is a deeper logical problem with moral conventionalism--it implies that "common practices" within a given social group cannot be morally impermissible, which on the face of it is clearly untrue. If this were the case, then the practice of slavery in the southern states of the United States prior to the Emancipation Proclamation would have been morally permissible simply because it was common practice--an absurd suggestion.

Examples of moral conventionalism:

Historically one can find very few examples of governments of modern nations that have never taken extreme measures at times to promote the security and welfare of their people. There is indeed good evidence that the CIA funded assassination plots against leaders of foreign nations who were regarded as threats to US interests, but looking at history, one cannot seriously raise objections against this.

"The Inquisition must have been justified and beneficial, if whole peoples invoked and defended it, if men of the loftiest souls founded and created it severally and impartially, and its very adversaries applied it on their own account, [funeral] pyre answering to pyre."
[Source: Benedetto Croce, *Philosophy of the Practical*]

Moral Legalism

A fallacy that is closely associated with moral conventionalism is moral legalism. Whereas moral conventionalism appeals to common or conventional practices as a normative standard, moral legalism appeals to laws of the state or, more generally, any codes of conduct that are accepted within some social group. Thus one commits this fallacy if one argues that some action is morally impermissible because the law of the land or some other generally accepted code of conduct forbids the action, or that an action is morally permissible because the law or some code of conduct does not forbid it.

There are two main reasons why this form of reasoning is fallacious. First, the law does not forbid all actions that can on reasonable grounds be regarded as ethically impermissible. There are good reasons for this: legislating morality is not consistent with the values of a free and open society. If in every case of moral wrongdoing the government stepped in to prosecute the wrongdoer, we would find that the most personal aspects of our lives and personal relationships would be subject to the government's oversight--a way of life that would be demeaning to personal freedom and dignity. Another good reason for not legislating morality is simply the impracticality of enforcing such broad legal proscriptions: there is simply too much immorality in society for any judicial system of reasonable size to handle. Accepted codes of conduct adopted

II. Moral Reasoning

by clubs, professional organizations, etc., do typically make stricter demands of individuals than does the law, but even these codes of conduct can fail to proscribe actions that are immoral. Thus an action that is allowed under the law or a particular code of conduct can still be immoral.

A second, and more significant, reason that moral legalism is fallacious is that it simply does not follow from the fact that a legislative body within a state or a social group decides that some action should be proscribed that therefore the action is morally impermissible; nor does it follow from the fact that such a body does not proscribe an action that the action is morally permissible. Historically it is not difficult to find examples of actions that can reasonably be considered to be morally permissible that were in fact illegal in their day. To take one quite obvious example from the history of our country, the founding of the United States as a separate nation from England was an illegal act. Thomas Jefferson, in the Declaration of Independence, offers a moral justification for the action, despite its illegality. A perusal of historical and present day legal systems will also yield abundant examples of the legal acceptance of actions that are on any reasonable standards clearly immoral, from slavery to mass extermination of cultural and religious groups.

None of this implies, of course, that we have moral license to do anything we want to do, legal or not. There are very good moral justifications that one can offer in support of abiding by the laws of the state or the ethical codes adopted by other social groups. But in moral discussion and debate, it is not enough to cite a law or an ethical code in support of a moral judgment--ethical principles and standards must be offered in support of adopting a law or ethical code as a reasonable normative standard.

Examples of moral legalism:

It is perfectly legal, in fact is a legal right, for landlords to evict tenants who have not paid their rent, any time, whatever the circumstances. There's nothing wrong with this.

[Comment: The legal right of landlords to evict their tenants does not mean that there can be no moral objection to this under certain circumstances: if, for example, a tenant has no place to go but the streets, and the landlord would suffer no extreme financial loss by allowing a tenant to stay where he/she is until suitable alternative housing can be found.]

There is a law still on the books of our state which outlaws acts of sodomy, even between consenting adults. Therefore it is clear that sodomy is morally wrong.

Moral Prudentialism

Prudence is a normative standard of action that requires rational agents to act in their own best interests. Although it is certainly wise to follow the demands of prudence, and we do this every day when we brush our teeth, take medicine for an illness, and perform other actions that we believe will be a benefit to us, a fallacy is nevertheless committed when it is argued that a certain action or practice is morally permissible or required because it will benefit (is prudent for) one

person or a certain group of people. Thus if someone were to suggest that each person has a moral responsibility never to help others when doing so detracts from their own interests, the fallacy of moral prudentialism is committed.

We will be considering in class an ethical theory known as "ethical egoism" which does not recognize moral prudentialism as a fallacy--in fact this theory says explicitly that the demands of prudence are in fact moral obligations, and thus we as moral agents are obligated to seek our own interests in all circumstances. Although this theory has been accepted by some ethicists, it is generally rejected on the grounds that it leads to absurd conclusions. We will consider these objections in class.

Other examples of moral prudentialism:

The practice of granting aid to foreign countries is immoral since the wealth expended in this manner could be used to benefit our own people.

There's nothing wrong with cheating on your taxes. After all, it saves you a few bucks in the long run.

Slippery Slope Argument

A familiar tactic in ethical debate is to suggest that some practice should not be allowed because it will lead to dire consequences for society. If sex education is allowed in the schools, the result will be rampant sexual promiscuity; if women are allowed equal status in the workplace to men, countless numbers of men will be jobless; if physician-assisted suicide is allowed, we will end up executing the physically disabled ("just as the Nazis did!"). The general strategy here is to suggest that even if a certain practice of itself is not morally objectionable, accepting it as morally permissible will lead us inevitably to consequences that are morally objectionable, and therefore the practice itself should be proscribed on moral grounds. The problem with this strategy of argument is that it fails to establish that the foreseen, undesirable consequences will be inevitable. Considering the examples used above, there is no reason to believe that the availability of sex education in schools leads to sexual promiscuity, that equal status for women in the work place will lead to widespread unemployment amongst men, or that allowing physician assisted suicide will inevitably lead to executing disabled people. Without strong empirical evidence (which is typically not offered by people who use slippery slope arguments) the conclusion simply has no support.

Other examples of slippery slope arguments:

II. Moral Reasoning

We dare not allow people in a persistent vegetative state to be taken off life support. If we do, it will be no time before we start taking people with terminal diseases off of life support, and the old people, and finally anyone who is not deemed to be useful to society.

There is an inherent danger in allowing any consideration of religion in public school classes, since although teachers may begin by teaching this material in a noncommittal fashion, it won't be long before they begin to indoctrinate students, and our schools will end up being evangelical brainwashing centers.

Argument to the People (also known in its Latin form, *argumentum ad populum*)

An argument to the people in support of a particular moral judgment involves an appeal to certain popular prejudices and biases shared by a large number of people. Unlike moral conventionalism, which appeals to common practices, an *ad populum* argument is designed to inflame the passions of a certain audience, passions that will lead that audience to accept a certain moral conclusion. For example, consider a hypothetical example of a conservative religious leader proclaiming the following to his equally conservative congregation:

No one who is in the faith, and true to that faith, can accept the scourge of abortion perpetrated by the godless heathens that are destroying our society.

The reference to "godless heathens" and their acts of "destroying our society" will no doubt encourage many people to consider accept the implied conclusion of this argument, that abortion is morally wrong, but even if those who advocate the right of abortion in our society were primarily atheists (which in fact is clearly wrong from recent surveys), and even if these atheists were doing things that may be thought to be destroying the fabric of society, these facts would not establish by valid moral reasoning that abortion is morally wrong. In fact, these facts are quite irrelevant to the question of the moral status of abortion. This example suggests generally why *ad populum* arguments are fallacious: the prejudices and biases appealed to in such arguments, and the emotions they incite, are simply irrelevant to the moral issues they address.

Other examples of arguments to the people (*ad populum* arguments):

Social programs designed to aid the poor are immoral. After all, they are poor because they are lazy, and we shouldn't reward laziness.

[Comment: Even if it were true that poor people are generally lazy, which is a quite simplistic and dubious explanation of the economic inequities in our society, this argument does not establish why laziness should be a relevant factor in evaluating the moral status of social programs designed to benefit the poor. Shouldn't lazy people be entitled to some minimum economic and personal well-being? If not, the reasons for this need to be explained. In fact this argument, however, is not designed to offer any such

II. Moral Reasoning

valid basis for its conclusion--it simply relies on the emotional impact of the popular abhorrence to laziness.]

Experimenting on animals is morally wrong! Would you experiment on your cat or dog?
[Comment: We typically become quite emotionally attached to our pets, and this argument relies on that widespread emotional appeal. But our emotional attachment to our pets does not imply that it is morally wrong to experiment on any animals, just as the fact that people become emotionally attached to their houses implies that it would be immoral to raze a house.]

Red Herrings (sometimes known as the fallacy of "Irrelevant Premise" or in its Latin form "*Non Sequitur*," which means "it does not follow")

Red herrings are issues or points raised which are irrelevant to the conclusion that is drawn from these points or issues. The strategy of this argument is to divert the attention of the person hearing the argument from real issues, ones that need to be addressed in order to evaluate the truth or falsity of a conclusion, to bogus issues, but in such a way that it is not clear to the hearer that the cited issues are bogus. Consider a moral argument offered against the view that nonhuman animals are objects of moral concern (that is, that some forms of treatment of nonhuman animals are morally impermissible):

It doesn't matter what we do to animals, people are what matter. Human beings have moral rights--you cannot morally do harm to any human being. Animals are not people--they have no rights.

All that this argument says in support of the claim that it doesn't matter, on moral grounds, what we do to animals is the undeniable fact that human beings have moral rights. But this fact is quite irrelevant to the question of whether nonhuman animals should be regarded as objects of moral concern. From the fact that human beings are objects of moral concern, it does not follow that no nonhumans are objects of moral concern. The issue of the moral status of human beings raised in the context of this argument is, therefore, a red herring.

Other examples of red herrings or *non sequitur* moral arguments:

Suicide is not morally objectionable, since it is a choice made by a person who is not satisfied with their life, and we all have choices to make in life.

[Comment: There is no doubt that in many cases people who kill themselves genuinely choose that course of action, but this is wholly irrelevant to the question of whether suicide is morally permissible or not. Things we choose to do can be morally permissible or impermissible: that we choose a course of action has no bearing on the action's moral status.]

Some people have objected to the past practice of dumping chemicals, saying that this pollutes the natural environment. Some have even suggested that this is immoral. But you

II. Moral Reasoning

know, our environment is already filled with chemicals. Water itself is a substance with a chemical composition. Plants have chlorophyll, and that's a chemical too. In fact, every substance in nature is a chemical. So there isn't anything wrong with dumping chemicals into the environment--they're already there!

[Comment: The objection to the dumping of industrial wastes has to do with their toxicity and their destructive effects on the natural environment. This argument simply misses the point.]

III. Theoretical Normative Ethics

Theoretical normative ethics attempts to define reasonable normative principles that can be used as the basis of moral arguments. It is important here to distinguish such principles from the formal principles of moral reasoning that we considered in the last chapter, such as the principle of universalizability. Formal principles determine only the general structure of valid moral reasoning without any consideration of particular moral values--such as honesty, justice, etc.--that enter into the content of moral reasoning. On the other hand, normative principles express those moral values that provide the moral criterion by which one can determine whether a given act is right or wrong.

Prima Facie Principles

A well-known twentieth century ethicist, W.D. Ross, argued that moral issues can be understood as conflicts between certain prima facie duties, duties that can be expressed in terms of a number of commonly accepted principles which we will here call simply "prima facie principles."³ ("Prima facie"--pron. prima fashee--is a Latin term meaning "on first view"). These principles define what we might call the initial moral presumptions concerning how we should act as moral agents. A moral presumption is simply a presumption that someone ought to act in a certain way under certain circumstances unless some special justification is offered that defeats the presumption, that is, a justification showing that it is morally acceptable not to act in accordance with the presumption. When such moral presumptions come into conflict, then a moral issue arises.

Consider six prima facie principles that the vast majority of ethicists and moral agents generally would accept.

1. Principle of Honesty: Generally, moral agents have the obligation to tell the truth.

There is little doubt that truth-telling is a prima facie obligation that we commonly try to follow in our own conduct and expect others to follow as well. There is good reason for this. The communication of accurate information is the basis for effective collective action amongst people. Just imagine working with a group towards a common goal where everyone in the group is lies to everyone else--very little could get done.

The most serious question in applying this principle is determining under what circumstances it is violated. Certainly when someone offers a false statement with the intent to deceive someone else (a lie), the act would commonly be understood as a violation of the principle. But there are less direct forms of deception. What if a business partner tells me initially that he will not go on

³ W.D. Ross, *The Right and the Good* (Oxford: Oxford University Press, 1932).

vacation during the summer months, but then changes his mind, and allows me to persist in my belief that he will not go on vacation simply to avoid any personal conflict. Is allowing me to persist in a false belief being dishonest in a way that is proscribed by the principle of honesty? Perhaps, although when caught in the deception, the business partner might claim that he did nothing wrong since he never knowingly lied. There can be, then, reason for dispute concerning whether certain cases of deception violate the principle or not. We will return to these issues when we discuss honesty in professional ethics.

2. Principle of Promise-Keeping: Generally, moral agents have the obligation to keep their promises.

Like the principle of honesty, the principle of promise-keeping raises some questions concerning when it applies. The issue concerns what constitutes a promise. Certainly a direct statement by someone to the effect that the promise to do something constitutes a promise, but then there are times when we presume that a promise is made tacitly, due to the nature of a given human relationship. Parents, for example, might be understood to tacitly offer their children a promise to foster their development and protect their interests, even though parents will rarely state such a promise.

A commonly accepted exception to this principle is recognized when the moral agent who offers a promise is later unable to fulfill it. We cannot be expected to fulfill moral demands that we simply cannot fulfill, and this applies to promise-keeping. Thus if I promise to coach the Little League Baseball Team on Saturday afternoon, and Saturday morning I am in a traffic accident that causes serious injuries, I am released from any moral responsibility to fulfill the said promise.

3. Principle of Nonmaleficence: Generally, moral agents have the obligation to do no harm to others.

This principle forbids the doing of harm, where harm might be understood as both physical and psychological harm. Thus a parent who either physically abuses their child or causes extreme emotional distress in their child violates this principle. The commonly accepted exception to this principle is cases where doing harm is the only way to prevent harm to come to oneself (self-defense) or to others.

4. Principle of Beneficence: Generally, moral agents have the obligation to do good for others.

"Good" in this context, like "harm" in the principle of nonmaleficence, must be interpreted broadly, to cover both physical goods (nourishment, health) and psychological goods (security, happiness). Beneficence and nonmaleficence are often confused in cases of preventing harm from coming to others. Preventing harm is required by beneficence, not nonmaleficence, since preventing someone from coming to harm is not a failure to do harm oneself, but doing something good for another.

5. Principle of Autonomy: Generally, people have the right to live their lives as they see fit so long as doing so does not interfere with the correlative rights of others.

This moral principle corresponds to the political value of freedom of action and speech as it is protected in the US Constitution, and in other open societies. Here, however, it applies not simply as a restriction against the government, but as an obligation for all moral agents: we have the obligation to allow people to make their own decisions, to live as they wish, so long as the exercise of this freedom does not impede our exercise of our own rights.

There are two contexts in which the moral demand of autonomy applies. First, in those cases where our actions are sufficient as a means to pursue our own interests, we have the right not to be interfered with by the action of others--a right of noninterference. However, in those circumstances where we require the aid of others to attain our interests, such as a cancer patient who needs a skilled surgeon to remove the cancer, we have a right to control the actions that others take for our benefit. The patient, for example, has a right to control which procedures the surgeon takes to bring the cancer under control--a right of control.

6. Principle of Equality (Justice): Generally, people should be treated in a manner that accords to each an equality of respect.

This principle is more fundamental than the other principles in the sense that it expresses as a normative principle an implication of the principle of universalizability: the standard by which we treat one person is the same standard that we should use to treat all people.

It should be noted that this principle does not imply that everyone should be treated in the same way. People have different needs, and enter into different types of relationships with others involving specialized demands and requirements. We recognize this in our daily lives, and treat people differently in accordance with the particular demands of the context. One doesn't treat one's child like one treats, for example, the checker at a supermarket. What the principle does demand is that each person should be treated in accordance with the same criteria. This underlies, for example, the fair administration of the law. Everyone, whatever their background, should be treated according to the same legal standards. This does not imply that everyone will be treated in

the same way by the police, the courts, etc. (a convicted felon will not be treated in the same manner as an ordinary citizen), but the same standards will be applied regardless of who they are.

The Use of Prima Facie Principles

One can use these principles in a simple and straightforward way as the basis of moral justification. That is, one can use them as the general moral principles within the premises of a moral argument, and derive moral judgments concerning actions that have been or might be taken in a moral situation.

W.D. Ross suggested another, and potentially more valuable, use of such prima facie principles. These principles can aid in the definition of a moral dilemma by offering an explicit manner by which we can express the dilemma in terms of conflicting principles. Consider this case, for example.

Jones is a minister who is concerned about one of his parishioners who has experienced a great deal of stress of late from a recent divorce. In confidence, the parishioner tells Jones that he is planning to kidnap his son from his estranged wife, who currently has custody. Jones does not wish to violate a confidence, and yet he is sure that the plan, if carried out, will cause more pain in the end for all concerned.

A proscription against breaking a confidence between a minister and a parishioner falls under the principle of promise-keeping. When parishioners tell clergy about their private lives, the clergy is bound by an implicit promise never to reveal that information. On the other hand, in this situation Jones believes that intervening and preventing the kidnapping, by revealing his parishioner's plan to his estranged wife and/or the police, will prevent more harm for everyone concerned--an action that is justified under the principle of beneficence. Jones' dilemma, then, can be understood as a conflict between the principles of promise-keeping and beneficence. The issue becomes clear once prima facie principles are applied to the case.

In other cases, a conflict may arise in the way that one prima facie principle applies to a given moral situation. Consider the following situation.

Dorothy is a police officer called to the scene of a terrible accident. She determines that there are two people involved in the accident that need immediate first aid or their injuries will be fatal. But unfortunately there is no other help coming for a number of minutes, and while Dorothy is treating one victim, the other is likely to die. She is torn concerning what to do.

The principle of beneficence in this case applies with respect to two objects of moral concern, and requires that Dorothy aid both accident victims, although she likely will only be able to save one. Beneficence, in this case, conflicts with itself.

Solving Moral Dilemmas

We have now seen how we can clearly define moral dilemmas as conflicts of prima facie principles, but how can we go about solving moral dilemmas? W.D. Ross suggested that we can do this by simply investigating all of the morally relevant facts that apply in a moral situation, and then determine which of the conflicting prima facie duties is more important in the situation. The more important prima facie duty then overrides the less important, and constitutes our "actual duty" in the situation. Certainly this can be done easily enough in many situations, and we often do this in a spontaneous and unreflective manner. Thus, if we see a frightened man pursued by his potential murderer, and the latter asks us where the man went, beneficence overrides honesty and we tell a lie.

But there are less tractable dilemmas--dilemmas where it is not at all obvious which of the conflicting prima facie duties is more significant in the situation. In these cases, moral intuition fails us. What are we to do? One other way to resolve the dilemma is to attempt to define some "higher" moral principle, that offers a simple moral criterion for judgment that can be used to arbitrate the conflict between the prima facie principles. This is the approach that the vast majority of ethicists have taken historically, and the fruit of their efforts comprise the body of theoretical normative ethics. It is to this material that we will now turn.

Normative Ethical Theories

Two approaches to normative ethical theory have predominated in the last three hundred years or so: the teleological and the deontological approaches. The basic distinction between these two approaches is to be found in the different ways that they understand what determines the moral status of an action, that is the status of a given action of being either morally right or wrong, permissible or impermissible. The teleological approach (also known as "consequentialism") determines the moral status of action by the causal consequences or outcomes of action. That is, an action is right or wrong according to whether it leads to good or bad consequences. Thus, judging whether a given action is right or wrong requires that one consider the actual or probable consequences of the action. Is stealing wrong? A teleologist will say that we should consider what comes of stealing: who is benefitted? who is harmed? We will be considering two normative theories that take this approach: ethical egoism and utilitarianism. On the other hand, the deontological approach denies what the teleological approach affirms: it denies that the moral status of an action is determined by the consequences of the action in favor of a view that human rights must be respected and moral obligations fulfilled even if an optimal outcome is not achieved. Deontology claims, then, that some criterion other than consequential values determines the moral status of actions and should be used in making moral judgments. There have been a great variety of suggestions as to the moral criterion that should be used. We will consider two of these suggestions stemming from the most prominent deontological ethical theories in modern

III. Theoretical Normative Ethics

times: Kantianism (named after its founder, the philosopher Immanuel Kant), and contractarianism.

IV. Teleological Ethical Theory

President Harry Truman had a difficult decision to make in 1945, near the close of World War II. The Japanese military refused to surrender even though it was quite clear that their defeat was inevitable. Truman's military advisors, however, were estimating that a land invasion of Japan might result in the deaths of as upwards of 1,000,000 American service men and prisoners of war, as well as many thousands of Japanese casualties, both civilian and military. There was another option: drop the atomic bomb on a Japanese city. Estimates of the deaths and injuries for such a drop were high, but not as high as the the estimates for land invasion. Truman, of course, made the choice to drop the bomb, first on Hiroshima, and then, when an offer of surrender did not materialize, a second bomb on Nagasaki. The gamble worked: the Japanese surrendered. But the cost was indeed horrific: two populous Japanese cities destroyed, with upwards of 200,000 Japanese civilians killed either immediately or by radiation poisoning.

It is not difficult to understand Truman's thinking. Dropping the bomb had the potential to save lives. Thus, although dropping the bomb would certainly cause widespread death and destruction, and although it was far from certain that a land invasion would cause as many deaths as estimated, there was still a high probability, in Truman's view, that lives would be saved. Truman's decision remains one of the most controversial political decisions of recent times. Could further diplomatic efforts have brought about an offer of unconditional surrender without further bloodshed? Could a drop of the atomic bomb off the coast of Japan, or on an isolated military target, have convinced the Japanese of the futility of continuing the war without the horrible loss of life caused by the drops over Hiroshima and Nagasaki? What about the longer range risks of the nuclear arms race, that did eventually put the entire world at risk of nuclear holocaust? The issues are complex. But despite this, the principle behind the complexities of this ethical debate is quite straightforward: which alternative would bring about the end of the war in the least costly manner?

Teleological (Consequentialist) Ethics

The case of Truman's difficult decision reveals something important about the manner in which we make moral judgments in many situations: our judgments often boil down to thinking through the consequences of our actions, and doing what in the end we believe will bring about the greater good. Ethicists commonly call this approach to moral judgment "teleological ethics" (from the Greek roots *telos* = end or aim + *logos* = reason), or "consequentialism." The basic intuition behind teleological ethics is that the purpose of moral judgment is to bring about what is good in the world, and avoid what is bad or evil. The task of teleological ethical theory is to define in explicit terms the principle behind consequentialist moral judgment, and resolve some fundamental issues concerning its application.

In this course, we will consider two consequentialist ethical theories: ethical egoism and utilitarianism. We will find that these two theories represent consequentialist moral judgment in

quite distinct ways. The common ground of these theories, however, is the view that in moral judgment, the consequences are what matter, and are all that matter.

Ethical Egoism

Only a small minority of ethicists have considered ethical egoism to be a viable normative theory, and very often ethicists who on first glance appear to be offering an ethical egoist position turn out to be offering a version of rule utilitarianism instead (a theory we will be discussing below). Nonetheless, at times people do make decisions that are best represented by ethical egoist reasoning, and so it is worthwhile to consider the theory, if only as a way of understanding a common error of ethical reasoning.

The ethical egoist makes this fundamental normative claim: all moral decisions should be made on the basis of a consideration of what serves the interests of the moral agent him or herself, that is, the person who is making the decision. Accordingly, the proper basis for moral judgment, for the ethical egoist, can be accurately expressed in the form of the following normative principle:

One ought to do whatever is in one's own best interests.

It is important to be clear about what the ethical egoism is not saying. First, the ethical egoist is not saying that we, as moral agents, ought to act selfishly, that is, that we should never promote the interests of others. Acting in the interests of others is quite consistent with ethical egoism, so long as acting in other people's interests is at the same time acting in one's own interests. Given the fact that very often what is good for others is also good for ourselves, it is a fair bet that an ethical egoist would do things in many circumstances that promote the interests of others.

Second, an ethical egoist is not saying that a moral agent should do what he or she wants to do. Our desires are often inconsistent with our best interests, as is clear from the fact, for example, that many people desire to smoke even though they understand that smoking is not in their long-term interests with regard to their health. Rather, the proposed criterion for moral judgment for the ethical egoist is what is sometimes called "enlightened self-interest": the most reasonable judgment, in light of all available evidence, concerning what in fact will promote one's interests, both physically and psychologically.

As we noted earlier, it is not enough simply to propose an ethical viewpoint; one must also offer an argument for why the ethical view should be accepted. So, why should we accept ethical egoism? The argument for ethical egoism is typically based on a factual theory of human motivation, sometimes called psychological egoism: the claim that in fact all voluntary action is motivated by the fundamental aim of achieving some good for oneself. According to psychological egoism, altruism is an illusion. We never truly act for the benefit of others. Those actions that appear to be altruistic always turn out, upon closer examination, to be motivated by self-interest. So the parent who suffers the economic burdens and the inevitable trials and

tribulations of parenthood, in what appears to be a self-sacrificing manner, is really looking for their own self-satisfaction--such as the satisfaction of pride in the accomplishments of one's child. On the basis of this theory of human motivation, the ethical egoist concludes that we should, as a matter of moral obligation, seek our own interests.

This argument for ethical egoism raises the issue of the naturalistic fallacy: isn't the ethical egoist fallaciously arguing for a normative claim on the basis of a descriptive claim? Clearly this is what the ethical egoist does. But against this charge, the ethical egoist might make this point: if it is true that we all in fact seek our own interests, it is unlikely that we could do anything else. Thus, even if one could reasonably suggest that we should seek some other aim as a basis for moral action, the suggestion would be pointless, since we simply could not follow it. Thus, the best we can do as moral agents is to seek to fulfill our own interests in the most reasonable and circumspect manner possible.

This point is well taken only if psychological egoism is true. But there are good reasons to question whether it is true. One reason is the fact that there are countless cases on record of self-sacrifice: parents who sacrifice their well-being, at times their very lives, for the sake of their children; soldiers who in battle sacrifice themselves for the sake of their comrades in arms. These cases clearly suggest the implausibility of psychological egoism. Typically the psychological egoist will explain such cases by claiming that the aim of an apparent self-sacrifice is a feeling of satisfaction, as indicated above in regard to parenting. But in cases where someone sacrifices their life for the sake of others, we must conclude, if the psychological egoist is right, that the satisfaction that the person receives briefly before their death outweighs the satisfactions of living out their natural lives--a rather far-fetched idea.

The error behind psychological egoism seems to be a confusion of wants and interests. It does appear to be true that we always do what we want to do in some sense of the word "want." When we claim that we do what we do not want to do, typically this refers to an unpleasant means that we utilize to achieve a wanted end. I might not want to go to the dentist, considering the experience in its own right. But I do want healthy teeth, and in the light of my understanding that the unpleasant experience of going to the dentist is necessary for dental health, I want to go to the dentist for the sake of the end. The same might be said of the self-sacrificing act. The self-sacrificing altruist does not want to die, but in light of the fact that one's own death is necessary to save others, the altruist might be said to want to sacrifice his or her own life for the sake of others. But now, psychological egoism fallaciously infers from the fact that we can understand all acts as proceeding from personal wants to the conclusion that all acts proceed from an assessment of personal interests. This inference is invalid deductively, and inductively there is little evidence to support it. People who sacrifice themselves for others may want, in the sense explained above, to sacrifice themselves, but this in no way implies that they have a personal interest in the sacrifice, or that they act from a perceived personal interest. The most straightforward way of interpreting such actions is that they are instances where people sacrifice their own interests for the sake of others. The psychological egoist, because of the confusion of wants and interests, searches for some manner of understanding the motives of such acts in terms of self-interest, and is led by faulty logic to the implausible explanations of self-satisfaction considered above.

Ethical egoism, then, is based on a theory of human motivations that is quite implausible, and thus the theory itself has little to recommend its adoption. We should also note that if ethical egoism is not a viable ethical theory, then collective ethical egoism is equally unacceptable. Such a collective form of egoism would say that one ought simply to serve the common interests of one's social group, whether that group is one's community, one's religious denomination, or one's nation, irrespective of how one's actions affect people outside one's social group. If psychological egoism is false, then we can act for the interests of people outside our social group, and there is no basis for the claim that we ought not do so when the effects of our acts on those interests are significant. This leads us to the normative theory of utilitarianism.

Utilitarianism

The utilitarian, like the ethical egoist or the collective egoist, is a teleologist in approach: the utilitarian will claim that the moral status of what we do is determined by the consequences of what we do. But unlike the ethical egoist or collective egoist, the utilitarian will insist that the principle of equality should be applied to the interests of objects of moral concern. In other words, the interests of all objects of moral concern must be considered on an equal basis in all ethical deliberation. One's own interests, or the interests of people within one's own social group, should not be given a favored status over the interests of others. This is not to say that these interests are irrelevant to moral judgment. It would be as illegitimate to disregard one's own interests in moral judgment as it is to disregard the interests of others. The claim of the utilitarian is simply that the weight given to the personal or group interests of the moral agent in moral deliberation ought to be a function of the degree to which those interests are affected by the action under evaluation, not on the basis of who has those interests. Thus, if I am attempting to determine the moral status of my actions, and my interests will be affected by my actions to the same degree that the interests of each of nine other people are affected, then I ought to give my interests no greater weight in moral deliberation than the interests of any one of these other nine people.

Why should we accept this view? Perhaps the best argument for utilitarianism is, simply, the failure of ethical egoism and collective egoism as viable ethical theories. If there is no basis, and as we saw there appears to be none, to claim within a teleological framework that we ought morally to favor our own interests or those of one's social group in moral judgment, then the correct moral viewpoint is that all interests should be treated with equal weight. Thus, the nonmoral value of the consequences of action for anyone who is affected by an action must be taken into account in moral judgment.

The next step in the development of utilitarian theory is to consider what sorts of nonmoral values define our interests as objects of moral concern. In other words, what sorts of consequences provide the criterion by which we can decide the amount of good or evil that is produced by our actions. Classical utilitarians such as Jeremy Bentham (1782-1832) and John Stuart Mill (1806-73) argued that the proper moral criterion was happiness, and that happiness can be understood as the presence of pleasure and the absence of pain. This view is sometimes

called "hedonic utilitarianism" (from the Greek word *hedone* = pleasure). We should note that Bentham's and Mill's understanding of pleasure was quite broad. Any experience that we would call "pleasant" was considered pleasurable: listening to music, reading good literature, achieving success and satisfaction in one's career. Similarly, "pain" was used to describe any disagreeable experience, whether it involved physical pain or psychological pain, such as emotional anguish and disappointment. According to this view, then, a moral agent ought to aim at the production of pleasures and the alleviation of pains of whatever sort for anyone affected by the moral agent's actions.

There are a number of problems with hedonic utilitarianism. One problem is that the view interprets pleasure and pain so broadly that the terms lose any meaningful reference to a specific aspect of experience that could be used to determine the value of our experience. The terms "pleasure" and "pain" are used meaningfully when they point to specific feelings within experience, such as physical pleasure and pain, that can be distinguished from other sorts of feelings. But when used as broadly as Bentham and Mill used them, the terms simply become synonyms for good and bad experience, and it becomes useless to employ them.

G.E. Moore (1873-1958) argued that any attempt to define what constitutes good and bad experience commits the naturalistic fallacy, and is therefore pointless. The best that we can say with respect to the moral criterion for judgment under a utilitarian perspective is that moral agents should aim at the production of good experience, and the avoidance of bad experiences. This more recent viewpoint is often called "ideal utilitarianism," and it is currently the accepted viewpoint among most working utilitarian ethicists.

We have considered an issue concerning the correct moral criterion for judgment within utilitarianism. There is another issue concerning how this criterion should be applied in moral judgment, an issue that has given rise to two distinct versions of utilitarian theory, commonly called act and rule utilitarianism. Act utilitarianism was the common approach among utilitarians up until the 1950s (although there is reason to believe that John Stuart Mill argued at times from a rule utilitarian perspective). Since then, rule utilitarianism has been adopted by some ethical theorists, so that today both theories are used in ethical discussions.

Act Utilitarianism

Act utilitarianism offers the most straightforward way of applying the utilitarian criterion of moral judgment. According to this view, the criterion should be applied to each individual action of a moral agent when determining the action's moral value (that is, whether the action is right or wrong, permissible or impermissible). This procedure of moral judgment can be expressed in what is commonly called the principle of utility, the basic normative principle of utilitarianism. This principle has been expressed in a number of different ways by utilitarian ethicists, but we might settle upon an expression of the principle that is consistent with ideal utilitarianism as follows.

IV. Teleological Theory

One ought to seek to produce the greatest possible balance of good over evil, or the least possible balance of evil over good, for all who will be affected by one's actions.

"Utility" is the technical term used by utilitarians to refer to the degree to which an action produces good and/or avoids evil. Thus if action A is productive of a greater good than action B, A is said to have a greater utility than B. Likewise, if A and B both produce evil, but A produces a lesser evil than B, then A is said again to have greater utility than B. Thus, returning to the example of Truman's dilemma, if it is the case that dropping the bomb cost fewer lives than a land invasion would have, dropping the bomb would be said to have had a greater utility than a land invasion.

According to act utilitarianism, then, the moral value of an action is determined by, or is a function of, the nonmoral value that will be produced by the action for all parties affected in comparison with the nonmoral value that is produced by all alternative actions that a moral agent might take in a given moral situation. It is important to stress, here, that determining the moral value of an action according to act utilitarianism requires a consideration of all available alternatives. An action is not right or obligatory simply because it produces good consequences. Nor is an action necessarily wrong if it produces bad consequences. The key to moral judgment is to weigh the relative utility of alternatives. Thus, an action that produces good consequences can, upon examination, be morally wrong or impermissible if there is some alternative that produces a greater good; and, as in the example of Truman's decision, an action that produces evil consequences can be morally right or obligatory if it produces the lesser evil of all alternatives available to a moral agent.

Applying Act Utilitarianism

To avoid confusion, it is important to sort out a couple of different issues that can be addressed from a utilitarian perspective. One issue, the issue that is our primary focus in this course, is the reasonable determination of the moral value of an action in a moral situation. This is the question of moral obligation that is pertinent in deliberation: What among all alternative actions is the right action to perform in a moral situation? The answer, from the perspective of act utilitarianism, is that the right, or obligatory, action is the one that will produce the best possible consequences, as we have seen.

Another somewhat different question that can be addressed from an act utilitarian perspective is the question of moral responsibility: Is a moral agent who has already performed an action morally responsible for the action, that is, is it reasonable to blame, or praise, the agent for their action. We will consider in some greater detail the basis for such judgments later on in the course. We should note here, however, that it is fallacious, from an act utilitarian perspective, to determine the moral responsibility of an agent based on information that was unavailable to the agent at the time they made their decision. The crucial point is summed up in the traditional proverb "Hindsight is 20/20." An action that a given moral agent reasonably believes will produce

the best possible consequences could very well turn out to have disastrous consequences. It would be unjust, however, to blame a moral agent for an action the consequences of which the agent could not possibly foresee. Consequently, in determining moral responsibility from an act utilitarian perspective it is important that a judgment be made in light of the information available at the time, and such a judgment might be quite different from the one that we would make from the privileged standpoint of hindsight of what action was actually the right one to take. For example, considering Harry Truman's dilemma, it would be quite consistent upon act utilitarianism to make the judgment on hindsight that it was wrong to drop the atomic bomb on Japan, but that Truman should not be blamed for taking the wrong action since he could not have foreseen the ultimate consequences of his action.

Returning, now, to the issue of moral deliberation, what should a moral agent do when faced with a moral dilemma? The answer of act utilitarianism is that the moral agent should do that which, in light of all available evidence and information, the moral agent determines, in their best judgment, is the morally right action--the action that will produce the best consequences for all concerned. Moral judgment, then, is a risky business. We can never be certain, given our less than perfect foresight, of what the consequences of our actions will be. Still, we do constantly predict the consequences of our actions, and we do believe that these predictions can be made reasonably in light of past experience. When I go to the store, I predict that there will be food available there to purchase for my dinner. When I decide to go to college, I predict based on available statistics that over the course of my working life the choice is the right one to make from the standpoint of future earnings. The act utilitarian will argue that in moral judgment as well, reasonable predictions can be made, and thus although such judgments always will be uncertain, they can still be reasonable.

Invariably, because of the relative uncertainty of prediction, moral judgment from an act utilitarian perspective involves considering the relative probabilities of the consequences of our actions. In some cases, when accurate statistics are available, the mathematics of probability can provide a precise way of dealing with these issues. Thus if a state legislature is considering whether to raise the speed limit on highways by five miles per hour, statistics showing death rates from automobile accidents as a function of the legal speed limit can be used to determine the probability that the death rate in accidents will rise by a certain amount if the speed limit is raised.

In most moral situations, accurate statistics will not be available, in which case a moral agent will need to rely on a less precise, intuitive sense of probabilities. Though intuitive assessments of probabilities are imprecise, they still can be reasonably made based on the available evidence of past experience. It is reasonable to judge, for example, that there is a greater probability that the grocery store I go to will have lettuce available for purchase than, say, imported Camembert cheese. Likewise, in moral judgment, it is far more likely that an accident victim will survive an accident if I call for help and render what first aid I can than if I simply ignore the victim and hope that a police officer notices him soon. Thus act utilitarianism would justify the guidance of the principle of beneficence in such an instance.

Whether or not precise calculations of probability are possible, act utilitarianism places upon any moral agent the derivative obligation to seek any and all available evidence or information that is relevant to determining the probable consequences of one's action, and consider in a serious and

conscientious manner, in light of this information, what the probable consequences of one's actions will be. To do anything less would constitute a forfeiture of one's obligations as a moral agent to make the best judgment possible. Thus, if Truman had failed to seek the best expert advice concerning the probable consequences of dropping the atomic bomb as opposed to ground invasion--if he decided on a whim to drop the bomb--he could be blamed for moral negligence based on act utilitarian principles, even if his action turned out to be the right one.

Rule Utilitarianism

What is wrong with act utilitarianism? A number of problems with the view have been cited by supporters of rule utilitarianism, but one of the most serious charges cited is that in certain situations act utilitarianism justifies actions that from an ordinary, commonsensical perspective are typically regarded as immoral. The reason for this is simply that such actions, under certain circumstances, can turn out to have the best consequences. Consider, for example, a person who enjoys the music of Beethoven and decides to shoplift a CD of Beethoven's music from a large and prosperous record store. He takes the CD home, plays it for hours and hours, which provides him pure enjoyment. He also plays it for the Beethoven Society, of which he is a member, which delights the members of the Society. The store, of course, has lost the ten dollars they would have received for the goods. The owners can recover their loss by raising prices, and if we distribute the lost revenue of the one stolen CD over the thousands of customers that buy CDs at the store, the rise in price required to recoup the loss from this one instance of shoplifting is minuscule. It's quite arguable, then, from an act utilitarian perspective not only to conclude that this act of shoplifting was permissible, but that it was morally obligatory in light of the great pleasure it provided the shoplifter and his friends.

This conclusion is clearly not in accord with common moral viewpoints. Most people regard shoplifting as not only illegal, but morally wrong. Does this mean that utilitarianism is fundamentally flawed? The rule utilitarian will say no. The point we should make in this case, according to this revised form of utilitarianism, is that even though the consequences of this individual act of shoplifting may be favorable, the overall consequences of the practice of shoplifting are not at all favorable. Returning to our example, the cost of recouping the loss of revenue of the one CD to each of the thousands of customers of the CD store might indeed be quite small. But what if everyone stole their CDs? Then the store would fold, the music companies who produce CDs would go out of business, and no recorded music would be available to anyone. Not a pleasant prospect for the music lover. Thus if shoplifting were common, the consequences would be quite unacceptable, which justifies the claim that any given act of shoplifting CDs is morally unacceptable.

Now, general practices within society are governed and controlled by the means of rules. Some such rules are codified into law. The rules of the road, for example, govern the general practice of driving. Other rules that govern general practices are informally adopted by people within society. There are no laws against rude behavior, but it is generally accepted that one

should not be rude to others. Still other rules are adopted formally within some specific social group to govern the practices of the members of the group. For example, many professional organizations have adopted codes of ethics that govern the general practice of people working within those professions.

Since general practices are governed by rules, and their consequences determine the moral value of actions, for a rule utilitarian, then the consequences of the general adoption and observance of certain rules must be considered to determine the moral status of actions. We can modify the principle of utility to express this procedure of moral deliberation as follows.

One ought to act according to those rules of action which, if generally adopted, will produce the greatest possible balance of good over evil, or the least balance of evil over good.

A couple of points should be made here to avoid misunderstanding. First, the rule utilitarian is not claiming that the existing rules adopted within society determine the moral status of actions. There is no guarantee that the rules that are currently in place will produce the best consequences for all, and considering the suffering that has been caused by tyrannical governments historically, it is quite clear that mistakes can be made in the social adoption of rules. Rather, what we need to consider in moral judgment is what rules or sets of rules would produce the best outcomes. If such rules are not in place, this is a justifiable basis for an ethical indictment of present society based on rule utilitarian principles. Secondly, the rule utilitarian is not claiming that the ethical status of actions are dependent upon the accepted practices within a society. This would be moral conventionalism, a fallacy that the rule utilitarian would avoid as any other ethicist. An action that violates accepted forms of practice is morally justified, under rule utilitarianism, if the rule governing that action would lead to better consequences, if adopted, than currently accepted practices.

Applying Rule Utilitarianism

Rule utilitarianism broadens the focus of moral deliberation and ethical discussion considerably when compared to act utilitarianism. The central concern here is not the narrow consequences of the particular acts of individual moral agents, but the more encompassing and long-range consequences of social practices observed by all moral agents within society. Thus the issue is not what will happen if I do such and such, but what will happen if everyone as a rule did such and such, as compared to other forms of accepted practice.

Consider an example. One current issue in biomedical ethics (the ethics of the medical sciences) that has arisen in recent decades concerns when it is morally permissible, if ever, to take a brain-dead patient off life-sustaining equipment. The issue gained notoriety in the early 1970s from the case of Karen Ann Quinlan, a woman who was in a persistent vegetative state as the result of an automobile accident, and whose family brought suit against her doctors to force them

to take her off life-sustaining equipment, an action known as "passive euthanasia" in ethical discussion. How would a rule utilitarian address this issue? First, we should note again that the issue under rule utilitarianism is not about the probable consequences of any single case, such as the Karen Ann Quinlan case. Rather, the pertinent issue for ethical discussion concerns the probable consequences of alternative forms of general practice that might be adopted in all cases of patients who are in a persistent vegetative state. We could, then, couch the issue in these terms: Would it be better for society in the long run if patients in a persistent vegetative state were passively euthanized at the request of their families as opposed to outlawing this practice? Considering the question in this manner, we must take into account the grief that is often experienced by families who have essentially lost a loved one to severe brain damage, but where the body of the loved one is kept alive by artificial means. We also have to consider the potential risks that we take if we allow such patients to die: Could some patients whose cases are not hopeless, who might recover, be mistakenly allowed to die if passive euthanasia is allowed in cases of persistent vegetative state?

Actually, the issues are not quite this simple. To fully address the question from a rule utilitarian perspective, we must consider not simply the alternatives of outlawing or allowing passive euthanasia in such cases, but all viable alternatives concerning different rule-governed procedures that might be utilized in handling such cases. We must consider, for example, the possibility of requiring that such cases be brought before a board of medical experts to assure that patients whose condition is not hopeless are not mistakenly euthanized. Perhaps the further safeguard of a hearing in a court of law should be required in such cases. Any and all viable alternatives must be considered if the ethical analysis of the issue is to be thorough based on a rule utilitarian approach.

This moral issue raises another complication that might require consideration when applying rule utilitarianism. Many people in society today believe, on the basis of religious convictions or traditional ethical viewpoints, that human life is sacred. For such people the taking of any human life, even in cases of brain death, is morally wrong. People who have these views might be greatly distressed or aggrieved if passive euthanasia were allowed in such cases. Should the rule utilitarian take such moral sentiments into account as one possible negative consequence of permitting passive euthanasia in such cases? What is curious about this issue, of course, is that the sentiments at issue are ones that are a result of a moral viewpoint on precisely the ethical issue under consideration.

A rule utilitarian cannot consistently rule out the consideration of any negative effects of the implementation of a certain rule or practice within society. Because of this, a rule utilitarian is bound by their position to consider even the moral sentiments of people who would be opposed to such implementation. The important point to consider, however, from a rule utilitarian perspective is that it is not only the immediate consequences of the implementation of a certain social practice that must be considered, but the long-range consequences as well. Moral sentiments do change over time, and often they change as a result of major changes in social policy and practice. Years ago, before the establishment of women's political and economic rights, many men, and indeed some women, believed that women should not be accorded equal rights in society as a matter of traditional practice and moral principle. No doubt such people were greatly disturbed by the

gradual advancement of women's rights in society. But such attitudes have diminished considerably over time. Few believe today that it is morally objectionable for women to play an equal role in society—quite to the contrary, most people regard the denial of equal rights for women as a social injustice. And few would deny that the economic and social opportunities now open to women have benefited the lives of women generally, and society as a whole. Thus history gives us good reason to believe that when reasonable changes in social policies and practices are instituted, any negative feelings arising out of conservative moral viewpoints are relatively short-lived compared to the long-term benefits of these changes. In light of this, concerns regarding negative moral sentiments diminish quite significantly in rule utilitarian ethical evaluation.

In closing our discussion of rule utilitarianism, we should note that many of the points we considered above pertaining to the application of act utilitarianism apply in a similar manner to the application of rule utilitarianism as well. Again, with rule utilitarianism, as with act, we inevitably must deal with the uncertainties of predictions, but in this case predictions concerning the effects of the social institution of rules. Again, as with act utilitarianism, we will need to deal with such uncertainties by weighing the relative probabilities of uncertain outcomes. The rule utilitarian again will presume that such assessments can be made reasonably--not in itself an unreasonable presumption, since lawmakers assess the probable outcomes of the institution of socially adopted rules, namely laws, all the time. In the application of rule utilitarianism, as with the act formulation of the theory, the aim of moral deliberation is to make the most informed and intelligent assessment of outcomes possible, and base moral decisions on this assessment.

Act or Rule Utilitarianism?

What is the better theory, act or rule utilitarianism? Teleological ethicists have debated this question, and no doubt this debate will continue. Resolving theoretical issues in ethics can be an extraordinarily complicated affair, as anyone who works within the field knows all too well. Theoretical discussion is indeed important (most of what I have said above is the product of long ethical discussion and debate within the field of ethics), but we need not resolve the theoretical issues once and for all to make use of these theories, since it is quite clear even from the rudimentary discussion of the theories offered above that both theories offer significant tools of ethical reasoning and analysis that enable us to articulate and justify moral positions on a great variety of ethical questions.

Rule utilitarianism makes very good sense in addressing moral issues, such as the issue of shoplifting we considered earlier, where isolated individual actions might provide overall benefits for objects of moral concern, but where such actions, when generalized to socially acceptable practices, can have quite detrimental effects. On the other hand, act utilitarianism makes much better sense as the guiding basis of moral evaluation in cases where an action brought under moral evaluation is so unique, with respect to the nature of the action itself or the circumstances in which it would be taken, that it is difficult or impossible to articulate a rule that could apply over a

IV. Teleological Theory

sufficient number of cases so that the consequences of its implementation can be assessed. It might be the case, for example, that the historical circumstances under which Truman had to decide whether to drop the atomic bomb or not were so unprecedented and unique in character that there is little chance that a set of circumstances that are similar in a morally relevant sense will ever be encountered again. Historians have often argued that, although we can take many lessons from history, some historical situations are so unique that they cannot be considered or explained except on a case by case basis. If this is true of Truman's dilemma, then act utilitarianism is clearly the preferable theory for moral guidance.

There are also those isolated cases where we believe that exceptions to generally acceptable and warranted rules are justified. In such cases, an act utilitarian approach provides a sensible direction to take in developing a moral argument for the exception. On the other hand, rule utilitarianism offers a potential basis for criticism of any justification of such an exception. The pertinent issue is whether such an exception, if allowed, would itself undermine the authority of a valuable social rule or policy. Recently, for example, some ethicists and legal scholars have argued that because of the solid and expanding evidence suggesting the existence of the so-called "battered wife syndrome" (where women who have been battered repeatedly by their husbands lose the ability to think sensibly about their relationship with their batterers and ways to end the relationship), leniency should be shown in cases where women have murdered their abusive husbands. If it is the case, as has been argued, that such women pose little if any danger to the public, and paroling them or offering them leniency initially upon conviction will enable these women to repair their emotional wounds and get on with life, then act utilitarianism can provide a basis for leniency offered on a case by case basis. But then, from a rule utilitarian perspective, one issue that should be considered is whether granting such leniency undermines or weakens the authority of the legal and moral condemnation of murder underwritten by the rule of law. If it can be reasonably argued that this might be the case, then this, at least to some degree, weakens the strength of the act utilitarian endorsement of leniency.

No hard and fast rules can be offered, unfortunately, to determine when the use of one theory over another is warranted as the basis of moral deliberation. Moral issues can be complex, and novel moral situations arise that can require new and creative approaches to ethical deliberation, even to such basic questions as what normative ethical theory makes sense in a given moral situation. Thus the suggestions offered here as to when act vs. rule utilitarianism make sense can be only suggestions or at best rules of thumb. As with any complex set of issues, although rules of thumb are useful, and might hold true in most cases, we must always be open to the possibility that they may fail, in which case we have no choice but to remarshal the best of our creative energies to deal with the ever surprising, and at times troubling, contingencies and perplexities of human life.

V. Deontological Ethical Theory

Eduard Pernkopf's *Topographical Anatomy of the Human Being*, first published in the late 1930s, is a classic of human anatomical literature. Originally published in seven volumes, the work is filled with some of the most exquisitely detailed illustrations of human anatomy produced to date. The work is still widely used in medical education and practice as a valuable resource of information on human anatomy, and its illustrations have been reproduced in medical texts throughout the world.

None of these facts are disputed, but recently a question has been raised concerning whether the work should be regarded with the admiration that it has received in the past, or even whether it should be used at all. The issue is Eduard Pernkopf himself, and suspicions concerning how he produced his monumental anatomy text. Pernkopf was a German nationalist, an avowed anti-Semite, and a member of the Nazi Party. In 1938 he rose to the rank of Dean of the Medical School at the University of Vienna, and subsequently removed all "nonAryans" (a code word for Jews) from the school's faculty. More troubling than this, a recent investigation revealed that during Pernkopf's tenure as Dean of the school, some 1,370 victims of politically motivated execution at the directive of the Nazi courts of Austria were used as cadavers for medical research at the University of Vienna. Some specimens from these cadavers were preserved at the university until as recently as 1998. Since vital records were lost, it is not known with certainty whether any of these cadavers were used as models for Pernkopf's anatomy text, but there is strong suspicion among critics of the text that this is the case. There are also suspicions that some victims of the notorious Nazi "death camps," which systematically executed Jews, might have likewise been used.

The two physicians who originally called for the investigation, Howard Israel of Columbia University and William Seidelman of the University of Toronto, asked that a commemoration of the victims of Nazi terror be included in subsequent editions of Pernkopf's *Anatomy*. Other critics have asked whether the work should be published at all. Howard Spiro of Yale University, for example, has argued "when you say some good can come out of this data, it seems to me we tell our children and grandchildren the end justifies the means. I am very much against using any such data any way at all."⁴

The factual issues in this case are complex: Does the *Anatomy* have any real worth in current medical work as opposed to other, similar anatomy texts? Would the use of Pernkopf's *Anatomy* encourage the relaxation of ethical standards in the use of experimental subjects, and a more callous and inhumane attitude towards people generally? All such issues are relevant to ethical analysis under a utilitarian framework.

But now, let us suppose for the sake of argument that the *Anatomy* does have medical value. For example, let us suppose that it will lead to the discovery of a more effective way to treat indigestion, so that sufferers of this condition will experience quicker relief of their condition. Furthermore, let us suppose that the relief provided by this treatment will outweigh, in terms of

⁴ "Doctors Debate Using Anatomy Atlas Created by Nazi," *Kansas City Star*, August 3, 1997.

the general quality of life experience, whatever anxiety, concern, grief, etc., is caused for some people (e.g., the descendants of victims of the Holocaust) by the continued use of the *Anatomy*. Would this marginal gain in the nonmoral value of outcomes offer sufficient and definitive justification for the continued use of Pernkopf's *Anatomy*?

It is likely that many people would deny this. Two arguments might be offered in support of this denial, neither of which are concerned with the outcomes of the use of Pernkopf's *Anatomy*. First, it might be argued that the involuntary use of the victims' remains constituted a gross affront to their inherent worth and dignity as human beings--an unjustifiable moral wrong whatever the short- or long-term consequences of this use. The continued use of the *Anatomy* will simply perpetuate the original wrong done, and thus is unjustifiable. Second, one might emphasize the significance of the fact that the use of these victims' remains was involuntary (that is, they never consented to this use). The involuntary use of someone to further the interests of others is always wrong, and again the continued use of the *Anatomy* would simply perpetuate the original wrong done.

Are these arguments simply irrational, emotional responses to the moral issue? A deontologist would say no. (The term "deontology" has its origins in the Greek roots *deon* = duty + *logos* = reason.) These arguments, for the deontologist, are quite legitimate and reasonable justifications of the normative claim that it would be immoral to continue to use the *Anatomy*. In fact, from a strict deontological point of view, moral decisions should never be made on the basis of the likely outcomes of actions, since the worth of these outcomes is completely irrelevant to the question of moral obligation. What is relevant is the nature of the action itself. Certain types of action--such as the involuntary use of a person's remains--are wrong regardless of their outcomes, according to the deontologist. The aim of deontological theory is to define the basic normative principles that provide sound justification for this ethical position.

Kantianism

One of the most prominent deontological ethical theorists of modern times is the German philosopher Immanuel Kant (1724-1804). Kant's moral theory has had a tremendous influence on modern ethical discussion, and the broad framework of his approach is still widely used in ethical debate. Kant argued that the moral status of an action is not determined by its consequences. We are not morally obligated to seek the best overall outcome by our actions, but rather to perform those actions that accord with our moral duty--the fundamental demand that we should treat others, and ourselves, in a manner that is consistent with human dignity and worth.

The aim of Kantian thought is to make this meaningful but nonetheless vague criterion for moral judgment more precise and explicit, so that it may be used in a clear way as the basis of moral judgment. The central issue in this endeavor is this: How can we determine what actions are consistent with a moral respect for ourselves and others? Kant believed that our capacity for rational thought, thought that is logical and consistent, is sufficient to determine this. The basic idea behind this approach is that whenever a moral agent acts in an intentional manner, the action

implicitly warrants or recommends (or "wills" in Kant's language) the same action for others, and if a moral agent acts in a manner that accords with moral duty (that is, whenever the action is morally permissible), then the action is one that the moral agent could rationally (that is, consistently) recommend (or will) for all other moral agents.

Take, for example, the act of brushing one's teeth. Teeth brushing is conducive to good dental health, which in turn enables one to eat a healthy diet and also avoid the dental pain associated with tooth decay. It is therefore good for any person with teeth to brush their teeth regularly, and the fact that I brush my teeth shows that I believe this and thus implicitly recommend the practice to others. Furthermore, there is nothing irrational or inconsistent about this recommendation. I in no way contradict any other fundamental desire or wish (that is, any desire or wish in life that I could not give up) by recommending that others brush their teeth, nor is there anything absurd about wanting everyone to brush their teeth. Put in another way, in the language Kant himself used, I can as a rational agent rationally (consistently) will that everyone brush their teeth, which means for Kant that my own action of brushing my teeth is consistent with my moral duties (that is, I do nothing morally wrong by brushing my teeth--the action is permissible).

But now consider another example. Let us assume, as Kant did, that every rational agent has a fundamental desire for self-preservation, a desire to live, that they cannot give up. (This assumption might be questioned in the case of the suicide, but even in this case it might be thought that the suicide on a basic level wants to live, but that he or she simply doesn't want to live with the emotional trauma or whatever else is making life at the moment unbearable.) Now, consider in the light of this assumption whether I as a rational agent could recommend the act of intentionally killing another human being for any other rational agent who desires to kill someone. Clearly not, Kant would say. For if I recommend that people kill others, then I am recommending that others (if they wish) kill me, and that is inconsistent with my desire for self-preservation. Thus it would be inconsistent for me to recommend others to murder (Kant calls this a "contradiction of will"), and furthermore I could not consistently, as a rational agent, intentionally kill someone else because, as we saw, this act itself implicitly entails the recommendation of this act for others. Thus I know that it would also be immoral for me to kill others, and that it is my moral duty not to kill others.

It is necessary to consider only one more step in Kant's argument before we consider his fundamental moral principle. Kant believed that rational agents, when they perform an action intentionally (rather than, say, out of blind impulse), act according to rules, or what Kant called maxims, that they adopt as guides for their behavior. When I intentionally brush my teeth, I am following the maxim "One should brush one's teeth," or if I were to kill someone for personal gain, I would be following the maxim "One should kill others whenever it is expedient to do so." Thus, the implicit warrant or recommendation that I offer to other rational agents whenever I, as a rational agent, act intentionally is in effect a recommendation that others adopt the very same maxim that I am adopting for my action. Now a rule or maxim that everyone must follow would be a law, a rule of action that is universal (applies to everyone). Above we noted Kant's view that an intentional action is rational and thus morally permissible only if one could rationally recommend it to others. We now can understand that this is the same as saying that my action is rational and morally permissible only if I could will that the maxim of my action become a

universal law. This, then, is the basic moral principle for human action, what Kant called the Categorical Imperative. Kant's imperative can be translated from the original German as follows,

I should never act except in such a way that I can also will that the maxim of my action become a universal law.

I can rationally will that the maxim "One should brush one's teeth" become a universal law, and therefore I know that it is morally permissible to act upon this maxim; I cannot rationally will, however, that the maxim "One should kill others whenever it is expedient to do so" become universal law, and therefore I know that it is morally impermissible (morally wrong) to act upon this maxim, and morally obligatory that I adopt and act on the contradictory maxim "One should not kill."

Superficially, the categorical imperative in this formulation (which Kant called the universal law formulation; there are others that he offered, one of which we will consider below) might appear to be equivalent to the so-called Golden Rule: Do unto others as one would have others do unto you. Kant denied this equivalence for two quite good reasons. First of all, the Golden Rule provides moral guidance only if one presumes a previous moral judgment: the judgment of how others should treat oneself. Thus as a principle it does not provide a basic justification for moral judgments, but rather provides a means of translating self-regarding moral judgments (judgments of how the moral agent should be treated by others) into other-regarding judgments (judgments concerning how the moral agent should treat other objects of moral concern). The problem can be illustrated if we note that a masochist (someone who wishes to be harmed) could not derive the principle of non-maleficence from the Golden Rule. In fact the masochist would derive a principle of maleficence: one ought to harm others. Clearly this is not what is intended as an application of the Golden Rule, and thus the application of the Golden Rule as a means of deriving non-maleficence requires the moral condemnation of the masochist's self-regarding desire that others harm to him or her. Kant argued that the universal law formulation of the categorical imperative requires no presumed moral judgments. This claim has been challenged by some of Kant's critics, but it appears at least that he is right in his contention that it does not require self-regarding moral judgments as the basis of other-regarding moral judgments, as the Golden Rule clearly does.

Another reason for Kant's denial of the equivalence of the universal law formulation to the Golden Rule is simply that the latter cannot be used to derive judgments concerning one's moral duties to oneself. Kant believed that we had moral duties to ourselves, duties where the moral agent and the relevant object of moral concern are identical. One such duty, Kant argued, was that of not committing suicide. This duty can be derived from the categorical imperative in much the same way that we derived the duty not to murder others above. An act of suicide, Kant argued, contradicts the fundamental desire of self-preservation just as the act of intentionally killing another.

Kant's universal law formulation of the Categorical Imperative might also appear similar to rule utilitarianism, since both involve a consideration of universal rules. The crucial difference, however, is that the rule utilitarian considers the probable consequences of rules adopted

universally, whereas Kant considers whether the universalization of the agent's maxim contradicts the desires or interests of the agent himself. Thus, even if at a given time in history we were to judge, for example, that the universal adoption of the rule "One should kill others when one wishes" would lead to better overall consequences (say, because it will reduce a problem of overpopulation that reduces the quality of life for everyone), Kant would still insist that the maxim of the kill when universalized contradicts the fundamental desire of self-preservation, and for this reason murder would still be morally impermissible.

Applying the Universal Law Formulation

First we should note one very clear distinction between Kant's theory and utilitarian theory: Kant claimed that moral judgments can be made with greater certainty than any utilitarian could consistently claim. The reason for this is that utilitarian moral judgment, as we saw, depends on factual judgments concerning the future consequences of present actions, and such judgments can never be made with certainty. Kant, on the other hand, rejected the view that the consequences of actions are morally relevant, and thus moral judgment does not rely, according to his theory, on the uncertainties of prediction. In fact, Kant argued that moral judgment was an instance of a priori judgment, that is, judgment that is justifiable "prior to" or independently of empirical evidence, just as we commonly believe that one doesn't require any empirical evidence to know that "2+2=4" is true. Thus, one does not have to know anything empirically concerning when or where a given act will be performed, or what the consequences of the act will be, to offer a moral evaluation of the action.

But even though moral judgment does not depend on empirically known facts for Kant, there is in Kantian theory one central factual issue that must be resolved in order to determine the moral value of an action, and that is simply what sort or type of action the action is. This is the all-important, and at times, complicated issue of formulating the maxim of the moral agent's action, since it is the maxim that determines the type of action that a moral agent does or will do. Kant offered no specific guidelines to how one should describe an action in the form of a maxim as the basis of moral judgment, but we can make some reasonable guidelines explicit.

First, since the maxim is a rule of action adopted by the moral agent him or herself, it must be formulated normatively, not descriptively, as a guide for action. Thus any maxim can be formulated in the general form "I (or 'one') should (or 'ought to') do such and such." (An acceptable, and simpler, alternative to this form is the imperative, "Do such and such.") This might seem peculiar, since we have just said that the maxim determines descriptively what the moral agent does or will do. In fact this is not at all peculiar--we commonly understand what people do factually on the basis of an understanding of the normative rule they are intending to follow. Thus, if we see an amateur gardener picking out bean plants from her garden, we will either make the assessment that she is weeding or in the process of replanting depending on whether we believe she is intending to follow the rule "I should remove weeds from my garden," and mistakenly believes the plants she is picking are weeds, or the rule "I should remove old

plants when planting new plants," fully understanding that she is removing bean plants. Another way of making this same point is that in defining a maxim, what we are after is a descriptive understanding of the moral agent's intentions. Unlike utilitarianism, it is not an action defined in terms of actual or probable consequences that is brought under moral evaluation in deliberation, but the action as understood in terms of what the agent intends to do. Thus a second rule to follow in formulating the maxim is to include in the maxim only what the moral agent him or herself understands is involved in his or her action. If we discover, for example, that our amateur gardener in fact believes that bean plants are weeds, we could not formulate her maxim "I should pick out the bean plants," since because of her ignorance she could not have been intending this.

A third rule that should be followed in formulating the maxim is not to make the formulation of the maxim too broad nor too narrow. The problem that arises when we formulate the maxim too broadly is that we fail to capture in the maxim precisely those aspects of the moral agent's intentions that are relevant to the moral evaluation of their action. Thus, in evaluating the action of killing another human being, it would not do to formulate the maxim of the killer as "I should kill." This maxim would cover any instance where one causes the death of something, such as our amateur gardener pulling up bean plants, and clearly it is not simply causing the cessation of life that is of moral concern in the case of murder. It is the more specific intentions of killing a person who has their own desires, interests, goals, etc., that is relevant in describing the intentions of the murderer.

A formulation of the maxim that is too narrow, on the other hand, will not work because it will undercut the possibility of applying the crucial moral test to the maxim, the test of universalization. The problem is that such a maxim will include such specific references to the particular moral situation that universalization will not be possible for morally irrelevant reasons. Thus "I should stand on the observation deck of the Empire State Building on July 20, 1999" is not universalizable (not reasonable to will as a law), but not because of any morally relevant concern, but simply because it would be unreasonable to want everyone in the world to be in a confined space at once. The rule here, then, is that in formulating the maxim, no specific references to the particular moral situation of the action under evaluation should be included.

Unfortunately there are no more precise rules to follow to determine precisely how narrow or broad a formulation is too narrow or broad. One helpful way of approaching the task, however, is to begin with a short description of the action, even a one word description if possible, and then proceed on the basis of the question: What rule does one intend to follow in all and only those cases where one performs this action? Thus, we could begin our attempt to characterize the amateur gardener's maxim by describing her action simply as "replanting the garden" and ask "What rule does one intend to follow in all and only those cases where one replants a garden?" If the act of suicide is under moral evaluation, one can begin with the question "What rule does one intend to follow in all and only those cases of suicide?"

One last rule to remember when formulating a maxim is not to presume any moral judgments in the formulation itself. The problem here is that we are likely in this case to prejudge the moral value of the action rather than allowing the test of the Categorical Imperative to determine this for us. Thus in evaluating the morality of killing others, we should not formulate the maxim of

someone who kills another person as "I ought to murder others." Since a murder is, by definition, the immoral killing of another, we are likely to conclude that the action cannot be universalized simply because of our moral conviction that murder is wrong, rather than on the basis of the contradiction of will that it entails. For the same reason, the "should" or "ought" of the maxim should not be understood in the moral sense of these terms. In fact, Kant himself believed that in most cases, moral judgment is required when ethical demands contradict personal desires. In such cases, the "should" or "ought" of the maxim will be a prudential term, as one might say "I should go to the movies to night" because one has a desire for a pleasant diversion. Thus, in the example of murder that we considered above, the maxim "I ought to kill other people" would use "ought" in this sense, that is, "I ought to kill others when doing so achieves something I desire."

When the maxim is universalized, the result will be what Kant called a "contradiction of will" if it is morally impermissible to act upon the maxim. There are two ways in which this contradiction can occur. One way is that the universalization conflicts with some fundamental desire that we, as human beings, cannot renounce. The previous examples of the immorality of murder and suicide are cases of this form of contradiction of will, since in both cases, in Kant's view, the universalization of the relevant maxim contradicts our fundamental desire for self-preservation.

A second way in which contradiction occurs is when the universalization of the maxim contradicts a desire that we would be acting upon if we were to act upon the maxim itself. Consider, for example, the act of cheating. First, what is involved in all and only those cases where someone cheats on an examination? What is involved, quite clearly, is an intent of violating rules that prohibit the use of certain resources (e.g., another person's test, or a crib-sheet), rules that apply to all students in the examination--with the intent to better one's score. Thus we might formulate the maxim of the cheater in this way "I should when taking an exam break the rules prohibiting the use of certain resources when it is to my advantage to do so." Note here that the desire that the cheater acts upon is to get a better score on the exam than he or she would otherwise obtain. Now consider whether the cheater could universalize this maxim, that is, whether the cheater could consistently want everyone taking an exam to break the rules in this fashion. Clearly not, since if everyone cheated, there would be no advantage in cheating, and the hope for a better score relative to other students' scores would be thwarted. We might also note that if everyone cheated, grades would not serve to indicate a student's personal accomplishment, and they would lose all value in the eyes of potential employers, thus again thwarting the desires of the cheater. Thus it is impermissible to cheat, and morally obligatory to follow the maxim--which is itself an objective moral imperative--"Never cheat." Note here that the universalization of the maxim is self-defeating in that such universalization thwarts the very desire the cheater acts upon. This is the crucial result of this second case of contradiction of will.

The End-in-Itself Formulation of the Categorical Imperative

Kant offered several other versions of the Categorical Imperative, all of which he believed were equivalent to the universal law formulation. Kant's demonstration of the equivalence of these

various formulations is long and quite complicated, and we need not consider it for the purposes of this class. But one of the alternative formulations has been quite influential in ethics quite apart from the universal law formulation, and is well worth our attention: the so-called end-in-itself formulation. Kant's statement of this formulation is as follows:

So act as to treat humanity, whether in your own person or that of any other, in every case as an end, never as a means only.

The importance of this formulation stems from its value in defining in a precise way the distinction between a "person," where we use this term in a moral sense, and a "thing." A person for Kant is an object of moral concern, and the only type of object of moral concern he recognized (Kant did not recognize animals as objects of moral concern).

In a general way, moral personhood in ethics has been understood to be a status held by any rational agent: a rational agent is the holder of certain moral rights, such as the right to life and autonomy, that cannot ethically be abridged by any moral agent. This understanding of moral personhood is vague, however, since it provides little understanding of why rational agents have this status, and how, stated in the form of a general principle, moral persons should be treated. Kant, in offering the end-in-itself formulation, answered both of these questions. First, rational agents, Kant believed, are self-directing or in his terms "self-legislating": they give themselves the rules by which they act. An insentient thing, a rock for example, acts by certain rules--the laws of nature--but there is no sense in which we could meaningfully say that the rock gives itself these rules. However, as we just saw, the rational agent does do this by adopting maxims. Thus, he or she in a sense tells him or herself what to do. And again as we saw earlier, if an agent is rational the rules will be one's that the agent can rationally recommend to all other rational agents.

What is the moral significance of this factual distinction between a thing and a person? Kant's complete answer to this is rather complicated, but in short he argues that since the will of the rational moral agent is inherently good (a person of good ethical character is to be valued simply because they are good), and it is the rationality of this will that is the basis of this inherent goodness, then it is immoral to thwart the autonomy of the rational will by using the moral person as a means. To say it another way, we may permissibly use a thing, like a rock, in any manner we wish since there is nothing inherently good in the way a rock behaves, but we may not permissibly use a person in any way we wish, since the self-directive autonomy of the rational will is an inherent good.

Applying the End-in-Itself Formulation

As we have seen, every moral principle defines certain factual issues as morally relevant for determining the moral value of an action, and it is useful to begin when applying a principle with a clear definition of these factual issues. The crucial factual question when applying the end-in-itself formulation is this: Does such and such an action constitute a use of a moral person as means

only? It is important to note the qualification expressed here in the word "only." In a cooperative society, people make use of other people constantly. The patient makes use of the doctor to get well, the audience makes use of performers on stage for entertainment, etc. There is nothing wrong, for Kant, with such usage of other people since this use is cooperative--it is not the use of a person "as a means only."

So then, what does it mean to use a person as a means only. It is not sufficient to say simply "we use a person as a means only when we force them to do what they wish not to do," since many cases of the use of individuals as means only in Kant's sense do not involve the use of force. To return to a previous example, Kant would argue that the cheater uses others as means only, but clearly the cheater does not force others to do something.

One modern ethicist, Onora O'Neill, has offered a fairly straightforward way of interpreting Kant's formulation on this point. To quote O'Neill directly, "we use others as mere means if what we do reflects some maxim to which they could not in principle consent."⁵ The idea here is that a moral agent uses a person as a mere means when what the moral agent does to that person or what the person him or herself does serves an interest of the moral agent himself in a fashion that prevents the person from consciously choosing whether or not to serve that interest. This might be done by the use of force, as in the case of murder, where the murderer serves his or her own interests by forcibly ending the life of another, which precludes the consent of the victim simply because it is a forcible act. But in other cases, the consent of the person or persons used is not precluded by force, but by subterfuge.

Consider again the case of cheating. The cheater acts by the maxim "I should when taking an exam break the rules prohibiting the use of certain resources when it is to my advantage to do so." As we saw earlier, the cheater could not consistently wish that this maxim be universalized, since it is the honesty of other students that the cheater depends upon to establish the meaningfulness and therefore the value of good grades. Thus the cheater uses the honest work of the majority of students to meet his own ends. But the fact that the cheater does this surreptitiously prevents any honest student from consenting to the cheater's dishonest use of his or her work. In this sense, the cheater uses others in a way to which they cannot in principle consent, or in Kant's language, as a means only.

The crucial question to answer, then, when applying Kant's end-in-itself formulation is whether the actions of a moral agent exploit in some fashion another person in a way that precludes that person's ability to judge whether they would consent to such exploitation.

Even with Onora O'Neill's useful interpretation of Kant's end-in-itself formulation of the Categorical Imperative, it can still be difficult in some instances to apply. Consider again, for example, the moral issue with which we began this chapter: the possible use of the remains of the victims of politically motivated executions in Nazi-ruled Austria for the purposes of medical research. Let us suppose that in fact Pernkopf used these remains as models for his famous *Anatomy*. Was this the use of moral persons in a fashion to which they could not in principle

⁵ "Ending World Hunger" in *Matters of Life and Death*, 3rd edition, ed. Tom Regan (New York: McGraw-Hill, 1993), p. 260

consent? Clearly the victims did not offer any consent and were not given opportunity of consent to the use of their remains. However, it might be argued that the use of a person's remains is never the use of the person him or herself because, simply, the person's corpse is not the person him or herself. The corpse is thus merely a "thing" in Kant's language, and no moral issues as to its use can be raised. Does this close the moral discussion? Perhaps not, for it might be retorted that although the corpse is indeed not the person, a person's corpse is a person's property, the disposition of which after the person's death should be controlled by the person him or herself. If this point of view makes sense, then just as we use a person when we make use of their property in a manner to which they do not consent, so Pernkopf used the victims by using their corpses in a way to which they could not, and probably would not have, consented.

We can encounter, then, some difficult issues when applying Kant's end-in-itself formulation, but nevertheless this formulation has provided perhaps the clearest account of the status of moral personhood that has been offered in modern ethical thought.

Contractarianism (Contractualism)

The ethical theory of contractarianism derives from the social contract political theory of Thomas Hobbes (1588-1679). Hobbes believed that the fundamental warrant of the legal obligations that a state imposes upon its citizens stems from a tacit agreement that those citizens have with the state to the effect that the citizens will obey the laws of the state in exchange for the security, comfort, and prosperity that the state affords its citizens. In Hobbes' view, we should obey the laws of our government for the same reason that parties to a contract should follow the terms of their contract: a voluntary agreement binds individuals to recognize and abide by, in both cases, certain obligations. The one crucial difference is that in the case of the citizens' contract with the state, the agreement is tacit rather than explicit. We do not as citizens sign a contract with the state, or pledge an oath to follow the laws of the state. Rather, our voluntary choice to live within the boundaries of the state and reap the benefits of the civilized life that it provides constitutes a tacit acceptance of the terms of the agreement--the social contract.

There have been many criticisms of social contract theory, criticisms that show it to be all but untenable as a political theory in the way that Hobbes presented it. But the general notion that our relationships with others within society do entail, or at least may be fruitfully understood to entail, implicit agreements or contracts between individuals is still cogent in the view of many moral philosophers. When we believe that we have been treated wrongly by a friend who has betrayed us, we might provide a warrant for this belief by suggesting that a friendship involves certain reasonable expectations concerning how friends will treat each other, that no reasonable person would enter into a friendship with another person if they believed that these expectations would not be met, that our former friend understands the content of these expectations, and that therefore these expectations have the force of an agreement that might be, although it hardly ever is in any particular case, stated explicitly as a set of rules: one should help a friend in need, not violate confidences, etc. Similar expectations pertain to other social relationships. We expect that

parents will care for their children, that children will care for elderly parents to the extent of their means, that doctors will use every viable means available and appropriate to cure their patients, etc. These expectations define our social relationships and account for the ways in which our social connections with others aid us in living a more productive and satisfying life. In each case, if these expectations are not met, we might not only feel aggrieved, but believe that the offending person has broken a trust with us, and thus has acted in violation of their moral duty to us. This is a moral judgment. We are making the basic judgment that it is not only explicitly stated agreements between individuals, such as legal contracts, that bind by mutual obligation, but also tacit agreements constituted by mutual, reasonable expectations that we all understand when entering into social relationships with one another.

This is the heart of the ethical theory known as contractarianism or contractualism. Moral obligations stem from those basic rules that underlie and sustain our social lives, rules that reasonable persons would agree are necessary to define the minimal standards of behavior within given social relationships, and therefore rules that no reasonable person would consent to be violated by others in those relationships. We can state, then, the basic moral principle of the theory in this manner:

One ought to abide by those rules of conduct that rational moral agents would agree to observe when entering into a social relationship for their mutual benefit.

Applying Contractarianism

It is important to notice the hypothetical meaning of the phrase "rules of conduct that rational moral agents would agree to observe" in the basic principle of contractarianism. Actual agreements between individuals, according to the contractarian, have no normative force since such agreements may not be "rational" or "reasonable." I might agree with someone that I should give them half of my salary, when there is no good reason for my doing this. This hardly means, however, that I have a moral obligation to give the person half of my salary--quite to the contrary, a contractarian will argue that this agreement is quite wrong in a moral sense, that it constitutes an exploitation of me on the part of the other person, because the agreement is not in accord with what a rational moral agent would agree to. Thus actual agreements may or may not be in accord with reasonable moral demands. For a contractarian, the issue of moral obligation must be based on the consideration of hypothetical agreements where it is presumed that the parties of the agreements base their decisions on reasonable assessments of mutual self-interest. Another advantage of considering hypothetical rather than actual agreements is that it allows us to consider from a contractarian standpoint moral obligations between people where one party in the relationship is not capable of making any reasoned judgment concerning their interests. Thus, for example, we can consider from a contractarian standpoint the obligations of a physician to a comatose patient by asking what the patient would agree to if they understood their medical condition and were entering into a therapeutic relationship with a physician.

Like ethical egoism, the basic relevant factual question for moral judgment for a contractarian concerns the personal interests of moral agents, assessed from the perspective of enlightened self-interest. When considering the obligations of moral agents to one another in a given social relationship, we must first consider how the relationship impacts upon the interests of the parties involved. Unlike ethical egoism, however, the contractarian will consider the question of self-interests from the perspective of reasonable negotiation. The central issue, then, is what rules would be accepted by all parties in the given social relationship as the result of a mutual negotiation where the parties attempt to maximize the fulfillment of personal interests while making the assessment that their interests will be better served within the relationship than without it. This might be likened to the process of negotiating the price of a car. The potential buyer and seller make the judgment that their mutual interests can be best served by the seller selling the car to the buyer. In the process of negotiating the price, the seller attempts to maximize his or her interests--getting the most money for a car--and the buyer attempts to do the same with respect to his or her own interests--getting the lowest price possible. Negotiations continue as long as the buyer and seller continue to believe that the relationship, once the terms of the agreement are negotiated, will serve their interests.

The central aspect of the assessment here will be reasonable expectations concerning the nonmoral value that each of the parties of the relationship are likely to acquire from the relationship. For example, considering on a contractarian basis the moral obligations of friends to one another will involve an assessment of the potential value of friendship: what personal interests are served in a friendship? On the basis of this, a judgment can be made of the reasonable expectations that friends will have of one another. For instance, one value that we realize within friendships is companionship--the joy of doing things with another person. On this basis, we can adjudge that potential friends will have the expectation of each other that the other person will want to spend time with them. A friend who is never around is in effect no friend at all.

It is on the basis, then, of an assessment of rational expectations concerning the personal value of a given social relationship that contractarian moral judgment rests. This point has an important implication for the application of contractarianism. In contractarian moral judgment, we must presume that the rational moral agents entering into the agreement are ignorant of the actual outcomes of the relationship, since, unlike act utilitarianism, it is not the nonmoral value that arises from action that determines the moral value of the action, but the terms of the agreement between moral agents. John Rawls, a contractarian ethicist, called this presumption "the veil of ignorance."

Consider again the obligations of friendship. Let's presume that a friend of mine is in danger--for no fault of his own someone is out to murder him, so that I take a certain personal risk by being around him. From a teleological perspective, I might make the judgment that it is better for one person to take this risk than two, and on the basis of this a might judge that it is permissible for me to abandon the relationship. But from a contractarian point of view, these actual consequences are not the issue--friend's expect loyalty from one another, and this loyalty binds me in obligation to my friend irrespective of the fact that in the end the friendship places me at personal risk. This provides a good account of why we generally condemn so-called "fair-weather friends."

Two other points concerning the application of contractarianism are in order. First, contractarianism can be used successfully to develop a justification for universal moral obligations, such as honesty, promise-keeping, etc. In these cases the reasoning should proceed with a consideration of the general expectations we have of all moral agents on the basis of the fulfillment of personal interests. Thus, in developing a contractarian justification of honesty, we need to consider how honesty is the basis of the fulfillment of mutual interests in general social interaction. Contractarianism, however, has the special virtue of being useful for considering moral obligations that pertain in more circumscribed and specialized social relationships. Thus, as we have seen, we can consider the moral obligations of the special relationship of friendship from this perspective. We could do the same for relationships between spouses, parents and children, professionals and clients, etc.

A second point is that contractarianism provides a basis for considering both symmetrical and asymmetrical obligations, depending on the nature of the relationship that is being considered. Symmetrical obligations pertain when all parties in a relationship have the same obligations to one another, such as the obligation of honesty that all moral agents have in common with respect to one another. Asymmetrical obligations pertain when there is differentiation of the obligations that the parties of a given relationship have to one another. For example, the physician has the obligation to treat his or her patient, an obligation that clearly the patient does not have to the physician. In relationships where there is a significant difference in relevant abilities of one moral agent to serve the interests of another, contractarian analysis can often lead to the determination of asymmetrical obligations.

Finally, consider an application of contractarian theory to an issue we considered earlier--cheating on examinations. One might initially consider the expectations that underlie the student-teacher relationship as a basis for contractarian analysis. The teacher expects students to be honest on examinations, and students understand this expectation when they enter the teacher's classroom. This understanding has the force of an implied contract, it might be argued, between student and teacher to the effect that in exchange for the opportunity to learn and receive a grade provided by the teacher, the student agrees to be honest on examinations in order to provide the teacher an accurate basis for evaluating the student's ability.

It is reasonable to suggest that there is such an implied understanding between teacher and student, but this analysis does not go quite far enough, since it does not show how the teacher as a rational and self-interested agent would reasonably demand honesty on the part of the student as a component of the implied agreement. The problem is that it is not at all clear why it would be in the best interests of the teacher for students not to cheat. Perhaps a teacher who would allow cheating in the classroom would be eventually fired by an administration that has an interest in keeping student's honest, but this simply raises the question again: why should student's be honest on examinations, and why should the administration demand such honesty?

A broader approach to the question is needed here. First, it is important to note that what goes on in any given classroom is a part of a broad social system or institution of education. This system comprises schools and universities on all levels, and involves the participation of students, teachers, and administrators of these schools as well as ancillary organizations such as the various accrediting agencies, the U.S. Department of Education, etc. Now consider: Who are the

beneficiaries of this broad social institution? That is, whose interests are served by this institution? First, students are beneficiaries to the extent that they gain the knowledge required to pursue life goals, including, although not necessarily confined to, vocational goals. Second, employers are beneficiaries insofar as they are able on the basis of this system to get employees who have the requisite skills for work they wish to be done. Third, we might argue that society more generally is a beneficiary since society benefits in multiple ways from a well-educated public. Finally, within the education system itself, schools and universities on higher levels of learning (for example, graduate and professional schools) are beneficiaries from the educational success of schools on lower levels, since such success provides the basis for effective training on these higher levels.

This analysis of the beneficiaries of the system of education provides a number of avenues to develop a contractarian argument. First, external beneficiaries such as employers would not reasonably accept the academic credentials of graduates from educational institutions which did not actively proscribe cheating for the simple reason that there would then be no way to assess the academic accomplishments of these graduates. Thus educational institutions can be understood to have an implied contract with such beneficiaries that cheating will not be allowed. Second, society in general, through government support of education, invests social wealth and resources into education. The implied agreement, it can be argued, that is the basis of this support is that students will in fact be educated, from which society can reap benefits. This implicit agreement is violated if there is no way to assess the effectiveness of this education, accurate grading being the primary means of this assessment. Third, a similar type of contract is implied within the educational system itself, between schools on different levels. High schools implicitly agree to enforce standards accurate grading for colleges offering baccalaureate programs, and such colleges have a similar agreement with graduate and professional schools. In all such agreements, the administrative officers and faculty of educational institutions uphold the contractual demand of assuring high academic standards, including those governing cheating.

Finally, we can consider the contractual implications of the students' interests in the educational system. The general honesty of grading in the educational system is, as we saw earlier in the Kantian analysis of cheating, an essential condition for the value of the grades earned within the system. If students typically obtained grades dishonestly, good grades would have little worth in the job market. On the basis of this, we can argue that there is an implicit agreement among all students within the system with each other and the educational institutions of which they are a part to the effect that each student agrees to be honest to assure the honesty of grading which is the condition for obtaining academic credentials that have professional value to the students themselves.

It is apparent from this example that the application of contractarian principles does require some broad and reflective consideration of systematic social relationships embodied in the various institutional and organizational frameworks of society. The crucial issue to begin with is to consider how people are benefited, or might reasonably expect to be benefited, by such social relationships. On the basis of this, we can consider from a contractarian perspective the implied contractual obligations of these relationships.

VI. Applying Ethical Theory: Some General Suggestions

After considering four very different ethical theories, it is natural to ask when considering moral dilemmas: Which one should I apply? There is no pat answer that can be offered to this question, but some suggestions can be made that might be helpful when dealing with moral issues from theoretical perspectives.

In recent decades, many ethicists have relinquished the goal of developing a single ethical theory as the basis of all reasonable moral judgment in favor of a more flexible, multitheoretic approach guided by the concept of reflective equilibrium. The central idea of this approach is that reasonable moral judgment attempts to establish an equilibrium or balance between three factors that come into play when considering moral issues: (1) normative principles that provide criteria for moral judgment applicable in a range of moral situations, (2) moral intuitions, or basic moral sentiments and viewpoints that come into play when we reflect upon particular moral situations, and (3) factual beliefs pertaining to morally relevant factors in moral situations. To say that moral reasoning aims at a balance of these three factors is to say that no one of these factors has paramount importance in moral justification. We cannot approach moral problems by assuming that one given normative principle will always be valid when developing a justification for a moral viewpoint, nor can we always rely on our moral intuitions. What is required is a consistency between these various factors--a consistency that has its basis in the principle of universalizability that we considered in an earlier chapter. The implication of this approach for the application of ethical theory in moral deliberation is that our choice of what theory to apply must be based at least in part on how well a given theory "fits" into a consistent moral viewpoint that takes into account basic moral intuitions and relevant factual beliefs. The question then is how do we determine what constitutes the right fit.

Recent studies in the moral development of children have revealed the such development proceeds along two parallel lines: (1) a development of empathy for the feelings and concerns of others that constitutes a morality of care, and (2) a sensitivity to issues of fair treatment that constitutes a morality of justice.⁶ These developmental paths are independent of one another in the sense that they at times direct moral sentiments in distinct directions. The morality of care raises concerns about how our actions affect the general physical and psychological well-being of others. The morality of justice disregards to a large extent such considerations in favor of the question of whether a given action treats others in a manner that is consistent with the equal treatment of all. From the former perspective we might feel a certain sympathy for someone who has, for example, treated us unkindly for reasons of personal trauma or distress, even though from the latter perspective we might judge such treatment to have been unjust, and deserving of moral sanction.

No moral justification proceeds directly from these empirical results (attempting to do so would commit the naturalistic fallacy), but they do reveal something about the general character

⁶ See Nel Hoddings, *The Challenge to Care in Schools* (New York: Teachers College Press, 1992).

of our moral intuitions. Specifically, they suggest that at times our moral intuitions respond most strongly to the good or bad effects of actions upon people--effects that can be the object of our empathy--and at other times intuitions respond more strongly to how actions constitute unfair or unjust treatment. When the former intuitions are prominent, we can say as a general rule that a teleological approach, and more specifically utilitarian theory, which takes into account the effects of actions, is likely to provide a justificatory basis that captures more accurately the substance of our intuitions. On the other hand, when concerns of justice are paramount, a deontological approach is likely to provide surer guidance.

Within the deontological framework, the decision of which of the three principles we considered--the universal law and end-in-itself formulations of the Categorical Imperative and the contractarian principle--provides surer guidance can require some reflective consideration of the reasons behind our belief that some given act was unfair or unjust. If these reasons center upon the view that the agent who performs this action could not possibly want everyone to do the same, and thus takes a liberty that he or she would not want others to take, then the universal law formulation provides a straightforward approach to moral justification. If we believe that an action exploits or uses another person, the end-in-itself formulation provides a way of offering a clear moral justification. Finally, if our sense is that another person breaks an implicit trust with us based on reasonable expectations of how people should act in social relationships where mutual interests are involved, contractarian theory is likely to provide the surest guidance.

It is helpful when considering these issues to make trial applications of the different ethical theories. Often when we consider in some detail how different ethical theories would define the morally relevant facts concerning a given moral question, it becomes clear that these facts as defined by one ethical theory accord more closely with what we consider to be most significant with regard to the moral question than alternative ethical theories.

Again, it should be stressed that these are at best rules of thumb. Ethical thought, like any complex human endeavor, is a creative process where progress can require some considerable effort mixed with some creative insight. No set of rules can be offered that assures outcomes that in the long run will be found acceptable. But the history of ethical discussion and debate does demonstrate progress, and there is no reason not to expect that further progress can and will be made on the ethical issues that at present trouble and perplex us.