

# THE ETHICS OF CONSEQUENCES

CHAPTER

5

## Utilitarianism

In the 1970s, Americans became increasingly aware of the carcinogenic effects of asbestos. Becoming aware of its cancer-causing consequences was no easy matter, for asbestos often does not cause cancer until twenty or more years after exposure to the asbestos dust. During World War II, a number of American factory and dock workers handled large quantities of asbestos, which was needed for the war effort. In the 1960s, a disproportionately large percentage of them were coming down with lung cancer, and researchers began to realize that this was the result of their earlier exposure to asbestos.

By the time the harmful effects of asbestos exposure were discovered, it was being widely used throughout the United States for insulation and brake linings and pads. Asbestos was in office buildings, school buildings, nursing homes, and private residences. Because of its harmful nature, it was difficult and costly to remove. Workers needed special training and special protective clothing and breathing apparatus before they could safely work with it. One of the questions we faced as a country was precisely what to do about this problem.

To answer this question, we had to look at the consequences of either leaving the asbestos in place (perhaps with an appropriate warning) or of removing it (at high cost). When we looked at the consequences, we had to add up the potential costs and benefits of the various courses of action in various types of situations. The clearest case was in elementary schools. Children exposed to asbestos might come down with lung cancer in their thirties; in

addition, many of those who worked at those schools were comparatively young. Requiring the removal of asbestos from elementary schools was an easy decision: Whatever the costs of removing it, the potential damage was so great that few would disagree with regulations requiring its removal from schools. Should we also require its removal from all individual homes? Here, the decision was more difficult, for the costs were higher in proportion to the benefits. For every elementary school in the country, there are probably dozens—perhaps hundreds—of private residences insulated with asbestos. The work involved in removing asbestos from all those residences would be considerably more than the work required for the schools, and the enforcement task would be much greater as well. Moreover, the cost would presumably be borne by the individual residents or owners, who may be neither able nor willing to shoulder such a burden. Whereas a young couple with children would probably be willing to do whatever is necessary to remove the asbestos from their home, a retired couple in their eighties on a fixed income would presumably be more reluctant to spend a lot of money for this purpose. A utilitarian might well decide to leave this decision in each individual's hands. Finally, should we mandate the removal of asbestos from nursing homes and other such facilities, where the majority of people will not live an additional thirty years to be affected by asbestos? What about the welfare of those who are younger and work in such facilities or visit them regularly? Utilitarians have to look closely at the consequences in each type of case.

Utilitarianism begins with one of the most important moral insights of modern times and couples it with a powerful metaphor that underlies our moral life. The insight is that *consequences count*; indeed, it goes one step further than this and claims that *only* consequences count. This puts it in sharp contrast to Kant's moral philosophy, which—as we shall see in the next chapter—places almost exclusive emphasis on the motives behind an action. Utilitarianism goes to the other extreme and maintains that the **morality** of an action is determined solely through an assessment of its consequences. It is for this reason that we call utilitarianism a *consequentialist* moral doctrine; morality, for the utilitarian, is solely a matter of consequences. Utilitarianism is not the only consequentialist doctrine we have seen. Ethical egoism is also consequentialist, but it demands that we consider consequences only insofar as they affect our own individual well-being. Utilitarianism demands that we consider the impact of the consequences on *everyone* affected by the matter under consideration. The morally right action, the one we ought to perform, is the one that produces the greatest overall positive consequences for everyone.

Once utilitarians have claimed that morality is solely a matter of consequences, they need to address several questions. First, they need to specify *the yardstick* or criterion in terms of which consequences are measured. Typically, utilitarians claim that we ought to do whatever produces the greatest amount of utility. But then, utility must be defined. Pleasure, happiness, and preference satisfaction are the three most common candidates for the definition of utility. Second, utilitarians need to indicate *how the consequences can be measured*; that is, they need to provide an account of how the yardstick can be applied to

measure utility. Third, utilitarians must address the question of how high their standards are, of *how much utility we must strive for*. How much, in other words, is enough utility? Although utilitarianism has often been stated in terms of maximizing utility, some have recently suggested that a less stringent and more attainable standard of expectations should be assumed. Fourth, utilitarians must indicate *what types of things are to be judged* in terms of their consequences. The three most common candidates here are acts, rules, and social policies. Finally, utilitarians must answer the question of *whom these are consequences for*. Clearly, they are not just the consequences for the individual agent; that would be ethical egoism. Do utilitarians take into account the consequences for all human beings or just for some subset, such as those in our own country? Do they take into account the consequences for future generations as well as the present one? Do they take into account the consequences for all sentient beings, animals as well as human beings; the natural environment as well as our constructed world; or just the human population?

Let's see how utilitarians have taken this basic insight about the moral significance of consequences and elaborated it into a formal theory of ethics.

## DEFINING UTILITY

Utilitarians claim that the only thing that counts morally is whatever produces the greatest amount of utility, or the greatest overall positive consequences. Yet what is the proper yardstick of utility? What has intrinsic value? Historically, utilitarians have taken pleasure and happiness as the measure of consequences. More recent versions of utilitarianism have turned either to higher ("ideal") goods or to preferences as the measure of consequences. Each of these four measures of intrinsic value has its strengths and its limitations.

### BENTHAM AND PLEASURE

Originally, utilitarianism became influential with the work of Jeremy Bentham (1748–1832), who defined utility in terms of pleasure and pain. According to Bentham, we should act in such a way as to *maximize pleasure and minimize pain*. This position is known as **hedonistic utilitarianism**. Notice that this is very different from straightforward hedonism, which recommends maximizing one's own pleasure and minimizing one's own pain. Hedonistic utilitarianism recommends maximizing the overall amount of pleasure and minimizing the overall amount of pain.

### MILL AND HAPPINESS

Bentham's philosophy quickly came under attack as "the pig's philosophy" because of what seemed to be its crude emphasis on sensual, bodily pleasures. John Stuart Mill (1806–1873), Bentham's godson, proposed a major reformulation of the utilitarian position by arguing that utility should be defined in terms of happiness rather than pleasure. Mill's standard seemed to be a definite

advance over Bentham's, for it was based on a higher standard than mere pleasure. This is called **eudaimonistic utilitarianism**. (The word *eudaimonistic* comes from the Greek word for "happiness," *eudaimonia*.) To see why this is the case, let's consider some of the differences between pleasure and happiness as the standard of utility.

#### PLEASURE VERSUS HAPPINESS

The differences between pleasure and happiness are significant. We tend to think of pleasure as primarily bodily or sensual in character. Eating, drinking, and having sex come immediately to mind as model cases of pleasure. Happiness, on the other hand, is usually less immediately tied down to the body. We might initially characterize it as belonging more to the mind or spirit than to the body.

Second, pleasure generally seems to be of shorter duration than happiness. This stems from the nature of pleasure itself. Pleasure, at least in the eyes of many psychologists and philosophers, is the enjoyable feeling we experience when a state of deprivation is replaced by a state of satiation or fulfillment. For example, pleasure is what we feel when we drink a nice cool glass of water to quench our thirst. Yet this gives us an insight into the reason pleasures are short-lived. Once we are satiated, we no longer experience the object as pleasurable. Once we are no longer thirsty, drinking water becomes less pleasurable. Happiness, on the other hand, seems to lie in the realization of certain goals, hopes, or plans for one's life. Insofar as these goals are intrinsically rewarding, we do not tire of them in the same way we may tire of certain pleasures.

Third, happiness may encompass both pleasure and pain. Indeed, we could easily imagine someone saying that his or her life is happy but still acknowledging painful moments. A good example of this is a woman giving birth to a long-hoped-for child. She may experience quite a bit of pain during and after the delivery, but she may still feel happy. Conversely, we can imagine someone experiencing pleasure but not feeling happy. Think of someone smoking crack, which directly stimulates the brain's pleasure center. They might take pleasure from it as they inhale deeply, but they could feel very unhappy with their life, career, marriage, and so forth.

Finally, there is more of an evaluative element in our notion of happiness than there is in our idea of pleasure. In reading the preceding example, many nonsmokers might have been repulsed at the idea of taking pleasure in smoking crack cocaine, especially in the morning. Yet this is not a reason for doubting that some smokers do find pleasure in it. We may want to distinguish between good and bad pleasures, between harmless and harmful ones, but we do not doubt that the bad pleasures are still pleasures. With happiness, on the other hand, we build in an evaluative component. We are likely to question whether people are genuinely happy in a way that we do not question whether they are genuinely feeling pleasure.

The problem with weighing consequences is that it is much easier to weigh pleasure than happiness or ideal goods, yet pleasure is the least suitable standard. *The closer we move toward a suitable standard of utility, the less able we are to subject it to quantification.*

#### OTHER ACCOUNTS OF UTILITY

Pleasure and happiness are not the only possible standards of utility, and the twentieth century saw attempts to redefine the standard of utility in terms of ideal goods such as freedom, knowledge, and justice (G. E. Moore) and individual preferences (Kenneth Arrow). These versions, *ideal utilitarianism* and **preference utilitarianism**, respectively, provide variations on the utilitarian theme. We can summarize these various versions of utilitarianism in the following way.

#### THE MEASURES OF UTILITY

Type of Utilitarianism	Standard of Utility	Number of Intrinsic Goods	Main Proponent
Hedonistic	Pleasure	One	Bentham
Eudaimonistic	Happiness	One	Mill
Ideal	Justice, Freedom, etc.	Many	Moore
Preference	Preference	Indeterminate	Arrow

No single candidate has emerged as the sole choice among philosophers for the standard of utility. The disagreement among philosophers over this issue seems to reflect a wider disagreement in our own society. If consequences count, we have still to decide what yardstick to use in measuring them. The attraction of preference theory in this context is that it permits this multiplicity of standards, all of which are expressed as preferences.

Indeed, there may be a distinct advantage to allowing a *multiplicity* of different types of factors underlying utility. This is a type of pluralism within a specific moral theory. Utilitarianism has sometimes been criticized for being too narrow, for reducing all our considerations in life to a single axis of utility, usually either pleasure or happiness. There is much to be said for a fuller, suppler theory that permits us to recognize that consequences need to be measured according to several yardsticks. The difficulty with such a move, however, is that it then makes utilitarianism a more complex doctrine, which is more difficult to apply in practice. Furthermore, utilitarians who go in this direction then need to specify the relationship among the different kinds of yardsticks. When one alternative ranks high on the yardstick of happiness, for example, and another course of action is high on the scale of justice, which takes precedence? Finally, it threatens to rob utilitarianism of its chief

advantage—namely, that it offers a clear method for calculating the morality of actions, rules, and social policies. Utilitarians who opt for a multiplicity of yardsticks must address questions such as these.

## APPLYING THE MEASURE

Once we accept a standard of utility, we are still faced with the task of specifying how that standard is to apply to the world in which we live. This is a thorny but crucial issue for utilitarians because one of the major attractions of utilitarianism is that it promises precision in the moral life. If it can't deliver on this promise because it can't be applied precisely, then it loses a significant factor for preferring it to competing moral theories.

### THE SCALE METAPHOR

One of the things that makes utilitarianism attractive is the root metaphor that underlies much of its language about measuring consequences. Utilitarianism is grounded in a root metaphor that has tremendous intuitive appeal to many of us: the metaphor of the scale. The very notion of weighing consequences presupposes that consequences are the kinds of things that can be placed on a scale. This metaphor pervades our everyday discourse about deciding among competing courses of action. Consider the types of things we often say.

- On *balance*, I'd rather go to the movies.
- When I *weigh* the alternatives, going to Hawaii looks best, even if it is hot at this time of year.
- Buying me an extra nice present *balanced* out the fact that they forgot to send it until a week after my birthday.
- Nothing can *outweigh* all the grief that the hit-and-run accident caused us.
- Only the death penalty can right the *scales* of justice.
- This is a *weighty* choice with *heavy* consequences.

Some scale metaphors have specifically monetary overtones, as though things were weighed in terms of their dollar value.

- I'm going to *pay him back* for all the grief he's given me over the years.
- How can I ever *repay* you for your kindness?
- That was a *costly* mistake.
- I'll be forever in your *debt*.

In these expressions, we see the way money metaphorically plays the role of the measure in terms of which consequences are assessed.

To weigh consequences, the utilitarian needs some kind of measure in terms of which utility is determined. We have examined various candidates for this yardstick: pleasure, happiness, ideals, and preferences. But we also need some way of marking off the units to be measured. Scales are often marked off in terms of ounces or grams. Yardsticks are usually marked off in terms of

inches. How do we mark off units of utility? One way of doing this is by assigning cardinal numbers (1, 2, 3, etc.) to pleasure or happiness. In the following section, we will consider this approach by using the stipulative concepts of hedons and dolors. The other way is to assign things ordinal (1st, 2nd, 3rd, etc.) or co-ordinal rankings relative to one another. This is the approach that preference utilitarians take. Let's look at both.

### HEDONS AND DOLORS

Once they agree upon the kind of yardstick, utilitarians must arrive at some consensus about how the individual units on the measuring stick are to be marked off. Utilitarians sometimes refer to units of pleasure or happiness as **hedons** and units of displeasure, suffering, or unhappiness as **dolors**. (The word *hedon* comes from the Greek word for "pleasure"; this is the same root that *hedonism* comes from. The word *dolor* comes from the Latin word *dolor*, which means "pain.") Particular things, at least in textbook examples, are then often assigned some number. A good corned beef sandwich, for example, may be five hedons, while a pleasant vacation to the Bahamas may be six thousand hedons. A visit to the dentist could be one hundred dolors, and the death of a close friend several thousand dolors.

Such a system may at first seem artificial, but utilitarians would argue that this apparent artificiality is not a serious problem. Although there may be no absolute scale in which going to the dentist is a one hundred, it may be the case that having a close friend die is roughly twenty or thirty times worse than going to the dentist. It is this relationship of relative suffering (or pleasure) that the utilitarian seeks to capture in assigning numerical values to various consequences. This is, the utilitarian further argues, something that we do quite naturally in our everyday lives. The utilitarian **calculus** is but a refinement and formalization of that everyday activity of assigning relative values to various occurrences—of ranking them in relation to each other according to the amount of pleasure or pain they yield. But as soon as cardinal utilitarians admit this, they are on their way to becoming preference utilitarians.

### THE DECISION PROCEDURE

How do utilitarians go about deciding on the moral worth of an action, granting that the consequences of an action can be specified in terms of hedons and dolors? They claim that, in any given situation, we must to the best of our ability (1) determine the consequences of the various courses of action open to us, (2) specify the hedons and dolors associated with each alternative, and then (3) perform the course of action that results in the greatest total amount of pleasure (i.e., hedons minus dolors). Imagine, for example, you are a utilitarian in the process of deciding between two pieces of proposed legislation about medical aid for the elderly. There are three possibilities open to you: to vote for a bill to reduce medical aid, to vote for a bill that would increase such aid, or to vote against both and thus effectively vote for keeping things the same.

## 136 CHAPTER 5

Reducing medical benefits for the elderly may result in ten hedons apiece for 100 million people and two hundred dolors for 20 million people, resulting in an overall utility of 3 billion dolors. Keeping benefits the same may result in twenty hedons apiece for 20 million people and three dolors for 100 million people, with a total overall utility of 2 billion, 600 million hedons. Finally, increasing benefits may result in ninety hedons apiece for 20 million people and twenty dolors for 100 million people, with a total overall utility of 200 million hedons. Thus, from a utilitarian point of view, we would be obligated to keep the benefits the same because both the other possible courses of action have less overall utility.

## HOW MUCH UTILITY IS ENOUGH?

One of the difficulties utilitarians have faced centers on the question of how much utility we are obligated to produce. The usual answer has been that we ought to do whatever produces the greatest overall amount of utility. When we compare competing courses of action, we should choose the best one, the one that maximizes utility.

### MAXIMIZING UTILITY

It is easy to understand the initial plausibility of this answer. Imagine we are considering whether to pass a particular piece of social legislation. We naturally consider how we can produce the greatest amount of good. If we are weighing the alternatives impartially, there is no attraction toward anything else; there is no pull toward doing less than the best. If some particular special interest group desired that we do less than the best, their wishes would be factored into the utilitarian equation along with everyone else's. There would, however, be no reason for giving them special weight. If we thought they deserved more weight, then presumably, it would be for some reason that would be recognizable within the utilitarian framework. It is only when we have partial interests of our own that there is a desire to do less than the best.

Yet the picture changes rapidly when we consider what it would be like to make decisions about our own personal lives in the same way. Here, personal desires and wants have a much more prominent role. Yet traditional versions of utilitarianism seem to demand that we do the maximum even in these more personal situations. The utilitarian always tries to produce the greatest overall amount of utility. This is in sharp contrast, for example, to Kantian ethics, which we shall discuss in detail in the next chapter. Kant's position just states negatively that a particular action is morally forbidden, whereas the utilitarian tells us positively that we must choose the specific course of action that maximizes utility. Thus, utilitarianism is an extremely demanding moral doctrine because it demands that we sacrifice our own pleasure, happiness, or preference satisfaction for the greater good—that is, for social utility.



### SUPEREROGATION

Philosophers have noted an interesting implication of the utilitarian view that we must always try to maximize utility. In many moral philosophies, it is often possible to act in a **supererogatory** fashion—that is, to go beyond the demands of duty and thus do something exceptionally meritorious. For the utilitarian, however, this is impossible. One is always obligated to do what yields the greatest amount of utility, and it is precisely this that constitutes duty. Thus, there is nothing above the call of duty. For the utilitarian, there is no room for supererogatory actions, for duty is so demanding that there is nothing above it that is greater.

This is a cause for concern among some philosophers, for they believe that a moral theory that has no room for supererogation must be mistaken in some important way. Our everyday moral intuitions tell us that sometimes people do something that is above the call of duty. But if utilitarianism is correct, this is impossible. Duty always calls for the maximum, so it is impossible to do anything above its call.


### CONSEQUENCES OF WHAT?

The utilitarian maintains that we ought to prefer whatever produces the greatest overall utility, and this is determined by weighing the consequences. But the consequences *of what?* Utilitarians have given at least three different answers to this question that are not necessarily mutually exclusive: acts, rules, and practices.

### ACT UTILITARIANISM

The first and most common version of utilitarianism says that we should look at the consequences of each individual action when attempting to determine its moral worth. This position, which is called **act utilitarianism**, maintains that we should always perform the *action* that will maximize utility, which will produce the greatest overall utility. Act utilitarianism is a tremendously powerful doctrine. One of its main attractions is that it seems to allow us to avoid rule worshiping, to deal with exceptions on the merits of the individual case.

**THE ADVANTAGE OF FOCUSING ON ACTS** One of the principal attractions of act utilitarianism is that it deals with individual decisions on a case-by-case basis. There is no such thing as an exception for the act utilitarian because every case is judged on its individual merits. This is in sharp contrast to a rule-oriented morality such as Kant's, which might demand in some specific situation that we act in a way that would cause more harm than good. As we shall see in Chapter 6, Kant held very strict views on lying. At times, he seems to believe that we are *never* permitted to tell a lie, even if lying would result in saving innocent human life. The standard example is of the Gestapo asking if you have seen any Jews running away from them. If you know the location of any such Jews, and if you



John Scott-Railton

## TALKING ABOUT ETHICS: A CONVERSATION WITH PETER RAILTON ABOUT UTILITARIANISM

**Peter Railton** is a contemporary philosopher at the University of Michigan who has written extensively in moral philosophy from a utilitarian perspective. He is the author of numerous articles on utilitarian theory and a coeditor of *Moral Discourse and Practice: Some Philosophical Approaches* (1996).

**Hinman:** Throughout your work, you have continually come back to the utilitarian standpoint as the perspective that best illuminates the moral life. What first drew you to utilitarianism?

**Railton:** Strange to say, but looking back it seems I was in the right place at the right time. If I'd grown up in the fifties, morality for me would probably mean social convention and coercion—I'd have to rebel against it, a "rebel without a cause." Growing up in the sixties, morality for me was the very thing for challenging social convention and resisting coercion—I could rebel *with* it, a rebel with a cause. I was one of those shaggy student protesters you see in the black-and-white photos getting dragged by the hair by a helmeted policeman. This experience taught me a lot. It was my first real moral lesson, and it had two parts. First, morality isn't something that just happens—if it happens at all, it's something people make happen. Second, society can be way off base—on civil rights, foreign policy, social inequality—and the thanks you will get for protesting this can just as easily be a billy club and a threatening letter from the dean of students as a pat on the back for good moral alertness. Morality has a critical edge, it doesn't take the status quo or "what seems right" for granted, and it isn't always easy. Now my second lesson . . .

**Hinman:** Wait a minute. You just said morality has a critical edge, but it sounds as if you call this a moral *lesson* because it fits your preconceptions.

**Railton:** Fair point. That's why my second lesson was important. This time it was my turn to be taken by surprise by moral protest. I remember the first meetings when women students began to protest their unequal treatment, even *within* the activist movement, and when some people began to raise questions about the treatment of animals and the environment. I thought, "Come on, get serious. We've got to worry about people getting bombed in Southeast Asia or dying in coal mines, and you're talking about humane treatment for chickens or equal treatment for middle-class college women." Well, people eventually got it through my head that I was wrong and was wrapping myself in righteousness to avoid facing an uncomfortable truth: I had accepted gross inequalities in the treatment of women

and had been pretty much completely oblivious to the pain and suffering human societies inflict on animals. This taught me two more things. First, *I* could be as off base as anyone—so worrying about the morality of an act or practice is quite different from asking, “What do I think of it?” or “How do I feel about it?” It’s a lot more like belief: “Here’s what I now think. But what are my grounds—and how would I get more evidence?” And second, that a key element in morality’s critical force is overcoming the one-sidedness and partiality that come to us so naturally. We build mental and social barriers between “us” and “them”—other races, cultures, religions, species, and so on—which keep us from seeing or feeling what is really happening to them.

**Hinman:** This sounds like an argument for moral objectivity, but why utilitarianism—why not objectivity about *rights*? These struggles you describe were movements for rights—civil rights for African Americans, equal rights for women, animal rights.

**Railton:** That’s the way I tended to look at it, too. Something as important as having basic needs met or equal treatment should be a guarantee, a right. But then rights also seemed to be barriers to responding to basic needs. Serious protests inevitably lead to innocent people being harmed; rebellions, like the American Revolution or slave rebellions, always involve violence and violations of property rights that harm innocent people as well as oppressors. Could morality be saying that such protests or rebellions should never happen? Think where we’d be then! Or think about the way some people look at poverty: Do those people have a *right* to my help? That’s another kind of barrier. It seemed to me that there had to be something more basic than rights, something that explains where they come from and why they’re so important but also helps us to see how to resolve conflicts of rights—and to see when rights aren’t the whole story.

So by then I was all set for utilitarianism, or *consequentialism*, as I prefer to call it. Consequentialism begins by asking what sorts of things matter for their own sake in life—happiness and overcoming of pain, sure, but also freedom, knowledge, accomplishment. The way these things matter does not depend upon any particular rights or moral views we might have—virtuous and wicked pleasure are both pleasant, deserved and undeserved freedom are both sought after. Consequentialism then says that, at the most basic level, our moral evaluations should reflect how fully these intrinsic values are realized, not giving anyone or any group special weight. You can’t say, “Well, that person’s happiness is coming from homosexual activity, which I despise, so it doesn’t count” or “This calf—or for that matter, this fetus—would grow up and have certain good or bad experiences, but since it’s not a person and doesn’t have rights, that doesn’t count.” Should we permit homosexual marriages? Abortion? Eating the meat or eggs of free-range hens? The basic question for the consequentialist involves comparing the kinds of lives that would be enjoyed overall if we permitted these things, or if we prohibited them. This is a very hard question—we have to look at all

(continued)

## 140 CHAPTER 5

*(continued)*

kinds of evidence. “Moral intuitions,” however strong, don’t settle it. My strong hunch as of now is that we should permit all three, but I could be wrong. In any event, I can’t tell whether I’m right or wrong without a lot of evidence about long-term consequences.

**Hinman:** But you’ve just said we should count everyone equally. Now, you’re a middle-class parent in a prosperous society. Shouldn’t you, as a serious consequentialist, really think that people in your position ought to spend a lot less of their time and resources on themselves and their own families and a lot more on those in greater need? Don’t you yourself deserve pretty severe moral criticism precisely for being so partial?

**Railton:** There’s no question that I fall well short of doing the most possible to make the world a better place. But for the consequentialist, it’s a further question whether people like me are obliged to dramatically reduce their involvement in their own families or should be condemned and punished. We have to ask: What would things really be like if we all went around damning each other for things like giving special attention to one’s own children? In some Quaker communities in colonial North America, people were shunned, investigated, and excluded for things like merrymaking when God’s serious work was to be done. In some really earnest Quaker communities, a substantial fraction of the population might be under suspicion at a given time—and the rest were busy being selfdenying and feeling either self-righteous or nervous. This is not an uplifting spectacle, and as far as I can see, not a whole lot of good came from it. Modern U.S. Quakers, by contrast, run the American Friends Service Committee, which helps people all over the world have decent medical care, civil rights, self-sustaining agriculture, improve the status of women, etc. Many Quakers regularly devote time to this activity, and some even go off for several years to work for the AFSC overseas after college. But these Quakers also have parties, hobbies, active family lives, vacation homes. I see a lot of good coming out of that, for them, the people around them, and people around the world. I’m not a religious person, but I believe we should take some lessons from the Quakers about how much more can be accomplished when people work together rather than engage in solo acts of charity; how long-term moral activity—because it *does* often involve some sacrifice—needs to be supported; and how one can easily distort one’s own moral understanding and risk losing the point of morality if one becomes too censorious and moralistic. The U.S. has tremendously moralistic attitudes toward food and eating and among the highest rates in the world of eating disorder and obesity to show for it.

**Hinman:** But if the motives from which one acts (say, personal motives within the family) are quite different from the justification one has in acting (say, that this is part of a morally good life when viewed impersonally), isn’t that schizophrenic, or alienating?

---

**Railton:** There's no way to avoid differences between motives and justifications, unless we're going to say that motives are self-justifying. A soldier may be motivated to fight bravely to the death by love of his home country, but does this justify his act? What if his country is engaged in a cruel and unjust war? He may well believe that his country is just and worth fighting for. But if he's sane, he will also think that if he's wrong, and his country is waging a terrible war of conquest, he shouldn't be out there laying down his life for that.

Being intrinsically dedicated to something, or someone, doesn't mean being unconditionally committed. Should my wife really be unconditionally committed to me, regardless of how rotten I turn out to be, how lousy I am as a father to our children or a companion to her? I'd think her emotion toward me was something more like obsession than love if that were the case. And I want her love, not her obsession. Is that alienating?

**Hinman:** You're beginning to sound pretty complacent to me: All consequentialism seems to ask of us is to be conditionally mindful of the larger world and to put in some volunteer time at the AFSC on the weekend, except when there's a big family picnic (with barbecued free-range chicken, of course!). That's a long way from your radical starting point.

**Railton:** It is. And in truth I cannot paint too "user friendly" a picture of any moral theory I'd really respect. A moral theory that lets us off the hook too easily can't be taking seriously the real, but avoidable, suffering that goes on in the world. I'm sure my current life falls short—I don't even put in the weekly hours at the AFSC, and I buy factory chicken when the freerange is sold out! I don't think of my complacent life as irrational, but I do think there are lots of good moral reasons for acting that I'm not attending to. I say to myself: I'll do more when the kids are more self-sufficient. But that's probably wrong—I'd be doing a better job of raising them if they saw in me a more regular commitment to broad social issues.

**Hinman:** Enough on your weaknesses. What about consequentialism's weaknesses?

**Railton:** Lots of the well-known criticisms of consequentialism don't seem to work very well—often they turn on consequentialist considerations ("Life would be miserable if . . ."). But there are many weaknesses. Maybe the biggest isn't talked about much. People tend to see consequentialism as too definite, too simple a theory to capture all the vague and multilayered phenomena of morality. There's something to this criticism, but I want to emphasize something else: Consequentialism may be too *indefinite* and too *complex*. We're used to hearing about act versus rule utilitarianism. But really there are dozens of kinds of consequentialism—cooperative utilitarianism, motive utilitarianism, conscience utilitarianism, total versus average utilitarianism, actual versus expected value utilitarianism, and on and on. It depends upon what you are evaluating. The motives with the best consequences, for example, may not be the ones that always

---

(continued)

## 142 CHAPTER 5

*(continued)*

lead to the acts with the best consequences (since motives have consequences other than the acts they produce—for example, they shape how we experience things). Robert Adams and Donald Regan have done beautiful work showing the variety of consequentialism.

Morality is about how life goes, overall—not just about action and its consequences. The problem is that no one yet knows how to combine all these different evaluative focuses and criteria together into a unified picture. The underlying consequentialist intuition is that things should be arranged so as to bring about the best outcomes for all affected. But which things and what kinds of outcomes? The result is a rather severe lack of determinacy at the very heart of the theory. All is not lost, though. We may be able to identify what matters most and to decide how much determinacy does—and doesn't—matter.

must follow a rule that prohibits ever telling a lie, then you must tell the Gestapo the location of the fleeing Jews. The probable result is that innocent people who might otherwise have escaped their captors will be killed.

An act utilitarian, on the other hand, would have no difficulty lying in that particular situation if lying would produce the greatest overall amount of utility. It is simply a matter of doing the calculations, which in this case seem relatively straightforward. Moreover, such an approach accords with our basic moral intuitions. Common-sense morality tells us that it would at least be permissible to lie in such a situation—perhaps that it would even be required of us.

Despite its sensitivity to particular cases, or perhaps because of it, act utilitarianism is open to a number of objections. Let's briefly consider three: (1) the objection that it is too time-consuming to calculate the consequences of each individual action; (2) a parallel objection that it is too difficult to predict the consequences of individual actions, especially the long-term consequences; and (3) the argument that act utilitarianism opens the door to abuses, especially to abuses of justice, because of its neglect of general moral rules.

**TIME TO CALCULATE** Act utilitarians seem to be faced with quite a challenge. If they are going to weigh the consequences of each individual action they perform, they will probably spend a large and disproportionate portion of their lives just calculating consequences. That, critics charge, hardly seems like the best way to spend one's time. Furthermore, there may well be situations when there simply is no time to calculate. When we see an out-of-control cement truck careening toward a pedestrian about to cross the street, we can hardly stop to calculate the hedons and dolors before trying to pull the pedestrian to safety.

Act utilitarians have an answer to such criticisms. We can, they maintain, live most of our lives on the basis of *rules of thumb*, which summarize past experience in such situations. Indeed, many of our general moral rules in society are precisely

of this type. They express our collective social wisdom about what generally produces the best consequences for everyone. Act utilitarians have no difficulty with generally following such rules, but they insist that we be clear about the moral status of those rules. They are not absolute; rather, they are simply convenient summaries of past evaluations of individual acts. If, in some particular situation, we have reason to question whether the rule of thumb will produce the best consequences, then it is entirely appropriate to call that rule of thumb into question in that instance. So an act utilitarian may accept the rule of thumb, “Don’t lie,” but reconsider it in the Gestapo example.

Such an approach, act utilitarians maintain, gives them the best of both worlds: They are able to recognize the advantages of generally relying on rules in the moral life without being caught up in the rule worship that seems to characterize those who consider such rules absolute.

**THE LIMITS OF PREDICTION** Act utilitarians face a second difficulty, which they share with other versions of utilitarianism but which may be more acute for act utilitarians. How accurately can we predict consequences, especially long-term consequences of individual actions? Think, for example, of deciding which college to attend. It is often a difficult decision that involved endless comparisons. Yet for many of us, the most influential consequences—such as the individuals we met, the friends we developed, the people we fell in love with—were ones we could never have predicted.

Most act utilitarians are willing to agree with critics about the limits of our predictive powers, but they reframe that insight in such a way that it no longer counts against act utilitarianism. Our predictive powers are limited, they concede, but this is a difficulty with life, not with act utilitarianism. The proper response to this limitation is not to reject act utilitarianism but simply to recognize that this is part of the human condition. The best we can do is try to increase our ability to foresee consequences. It is unrealistic, however, to hope that we can eliminate uncertainty completely. The moral life contains an ineluctable element of uncertainty, which can be reduced but never eliminated.

It is impossible to imagine that we could live without some basic belief in the general predictability of the human as well as the natural world. If we turn on the shower in the morning, we expect the water to be there and at the proper temperature—and it usually is. We don’t expect that tomato juice will come out of the showerhead, just as we do not expect that food will cook if we put it in the refrigerator. We would quickly go crazy if there was not a significant degree of predictability in the world. Similarly, in human affairs, we have general expectations about how people will behave. If we ask people questions, we generally expect that they will answer. If we give someone \$1 million, we generally expect that the person will be delighted about it. If a close friend dies, we expect to be sad. We could hardly act in the social world if we did not believe that there is a general predictability to human behavior. This predictability does not have to be complete. We can make mistakes in our predictions and be surprised, but the very idea of being surprised presupposes expectations. Predictability is necessary for living.

THE LIMITS IMPOSED BY JUSTICE Critics see a third difficulty with act utilitarianism that is potentially much more disturbing than the preceding two objections. Act utilitarianism, they charge, opens the door to potential abuse—to condoning, perhaps even requiring, acts that contradict our everyday moral intuitions, especially intuitions about justice.

Let us begin by considering a bizarre but real example. In the 1970s, when a Beverly Hills widow named Sandra West died, she stipulated in her will that she be buried in her favorite 1969 red Ferrari, “with the seat slanted comfortably,” next to her deceased husband, who had been a Texas millionaire. Her brother-in-law, Sol West III, had promised to carry out her wishes. (Just to be sure, Mrs. West set up her will in such a way that Sol West would not inherit her promised \$3 million if he did not follow her wishes.) Clearly, this was a colossal waste of money—and of a good car. What would be wrong with disregarding that provision of the will, give her a regular burial, and use the money saved to do some good among the living? Clearly, one difficulty is that there are laws against such things, and one would be caught. But imagine what would happen if we embellished the example a little bit. Imagine the woman told only you about her will and imagine you promised her that it would be carried out. Imagine, finally, you know the location of her earlier will, which is exactly the same except that it contains no proviso about being buried in the Ferrari. What would be wrong with burning the most recent will and substituting the earlier one, which would leave all her money to a worthy charitable organization?

The obvious answer to this question is that you promised the woman to do what she asked. Yet keeping a promise in itself has no value for act utilitarians; its value depends on the consequences of the particular promise. In this case, who would benefit from your keeping the promise? Clearly, the woman would not—she’s already dead. Whatever benefit she might have derived from the thought of being buried occurred before her death. The institution of promising, including other people’s confidence that their wishes will be honored after their deaths, would neither suffer nor benefit because no one else knows her wishes in this matter. The numbers seem to come out clearly in favor of breaking the promise. Moreover, if the numbers do come out this way, then not only are we *permitted* to break the rules; if we are optimizing utilitarians, we are *obligated* to break them.

Consider a second, more difficult example. Imagine you are the police chief in a small town that has been terrorized for months by a child rapist. Imagine you discover through some unusual set of circumstances that the rapist has died in a freak accident, but there is no way you could convince the public that the person was indeed the rapist. The threat is past, but the public still lives in fear because you cannot convince them that the rapist is actually dead. Now imagine you have arrested someone whom you could frame for the rapes, a hobo with tuberculosis who has only six months to live. What would be wrong with framing him for the rapes? There would be no danger that the real rapist would be free to continue his rapes; you are certain the rapist is dead. The public would be reassured, feeling that their town was once again safe. The



man being punished would die soon anyway, and he might actually receive better medical care in jail than on the street. The act utilitarian would seem justified in convicting the man, even though he had not committed these crimes, because his conviction would result in the greatest overall utility.

### RULE UTILITARIANISM

Examples like these last two are disturbing to utilitarians, for they seem to suggest that act utilitarianism is too open to abuse and too likely to justify actions that conflict with justice or other values that intuitively are accepted by common-sense morality. Yet many think that utilitarianism is still basically correct in its emphasis on consequences and in its standard of utility. The problem, they conclude, lies in the fact that utilitarianism is looking at the consequences of each individual act. Instead, we should look at the overall consequences of adopting a *rule* that everyone should act in a particular way under certain types of circumstances. This approach is known as **rule utilitarianism**, *which claims that we ought to act in accordance with those rules that will produce the greatest overall amount of utility for society as a whole.*

Clearly, rule utilitarianism has a much better chance of dealing with the types of examples described here than act utilitarianism does. It would be much more difficult to imagine how one could justify a *rule* that supports breaking promises to the dead or convicting innocent people of crimes. Such rules simply do not maximize utility. By insisting that we justify rules instead of individual acts, the rule utilitarian seems to avoid certain injustices contained in act utilitarianism.

Act utilitarians disagree. They maintain that rule utilitarians are caught on the horns of a dilemma. Rule utilitarians must maintain either that their rules are without qualifications or else that their rules do have very specific qualifications. If their rules are without qualifications, then they are rule worshipers according to act utilitarians, for they say we should follow the rule at any cost, even when it produces bad consequences. Act utilitarians would admit that we should generally keep our promises to the dead, but it would simply be rule worshiping to keep the promise about being buried in the Ferrari in the circumstances described. On the other hand, if rule utilitarians are willing to make qualifications in cases when following the rule would produce bad consequences, they are really covert act utilitarians. If, for example, they would be willing to break their promise to the dead in the Ferrari case, then they would really be thinking like an act utilitarian. They would be willing to admit that they can disregard the rule in any individual case when the overall utility clearly demands it. Thus, rule utilitarianism seems to collapse back into act utilitarianism.

**JUSTICE AND RULE UTILITARIANISM** As we have seen, one of the principal concerns of critics of utilitarianism has centered on the possibility that act utilitarianism might require us to perform acts that clearly violate our common-sense moral expectations about justice. Yet it does not seem that rule utilitarianism is

entirely immune to such criticisms either. Imagine a society in which 10 percent of the population is enslaved but provided with many of the basic physical comforts of life such as good housing, nutritious food, and reasonable working hours. Imagine further that the slaves are given plenty of free entertainment (television, movies, music, etc.), much of which stresses the joys and rewards of being a good slave. As a result of this ongoing indoctrination and the tolerable physical conditions, the slaves do not feel a tremendous amount of discontent about their state in life. The masters, on the other hand, feel great that they have slaves, and it makes their lives significantly happier. Let's say that living in accord with this rule about slavery causes the slaves generally to feel about eight units of displeasure (eight dolors) apiece in their condition. Let's say that the masters experience two units of pleasure (two hedons) apiece as a result of being masters and having slaves. Recall that the slaves compose only 10 percent of the population. Here's the picture we get.

	How Much Pleasure Apiece?	How Many People?	Total
Slaves	8 dolors	100	-800
Masters	2 hedons	900	+1800

Rule utilitarianism would seem to justify slavery under these conditions, but our clear moral belief is that slavery is unjust. There are two possible conclusions we could draw from this. We could argue that since utilitarianism leads to morally unacceptable conclusions (in this instance, the justification of slavery), it is unacceptable.

The other possible conclusion involves the modification, rather than the rejection, of utilitarianism. We could imagine that some people might agree to such a societal arrangement if they knew in advance that they would be among the masters rather than the slaves. However, would anyone agree to it if they did not know which of the two classes they would be in? Some philosophers—John Rawls in *A Theory of Justice* (1971) is the most influential of these—have argued that we need to set up limits on the range of possible rules that could be adopted. These limits would be determined by considerations of justice, by consensual agreements, by human rights, and the like. Within these limits, utilitarian justifications of particular rules would be permitted, but no justification would be allowed that violated those limits. It is, then, debatable whether proponents of such positions would still count, or count themselves, as utilitarians. Some contemporary theorists who accept this view, such as Samuel Scheffler, continue to insist that they are consequentialists.

**MOTIVES AND THE DOMAIN OF APPLICABILITY** There is another aspect to this continuing controversy between act and rule utilitarians, and it centers on the role

of motives in the moral life. Although there has been some attempt (discussed later) recently to develop an account of motive utilitarianism, utilitarianism has long been open to the objection that it ignores motives. Typically, philosophers give hypothetical examples of this. Imagine an assassin, motivated purely by resentment and envy, is trying to kill the president of a small Middle Eastern country during a public ceremony. The assassin's bullet misses the president but strikes the ground, causing oil to gush forth. The newly discovered oil proves to be a main source of revenue for the previously impoverished country, and soon the entire country prospers. Utilitarians would have to maintain that this act was good because the consequences were good. Yet this seems to contradict our basic moral intuition that the assassin's motives of resentment and envy should count in our evaluation of the action.

Examples of this do not need to be far-fetched. Consider the attempted military coup in Russia in 1991, in which old-guard communists tried to take over the government from Gorbachev. The coup failed, and in the end, reformist leaders, especially Boris Yeltsin, emerged as far more powerful than they had been before. Presuming that this result will eventually bring more utility than any alternative, do we want to say that the hard-liners' attempted coup was a good act? Clearly, it was not *intended* to strengthen Yeltsin's position. The fact that it had this consequence was purely accidental. Indeed, it was intended to produce exactly the opposite consequences. It seems as though our moral evaluation of that action should somehow take account of the motive behind the act.

It is here that the distinction between act and rule utilitarianism has particular relevance. Although both act and rule utilitarianism ignore motives, this seems to be less of a problem for rule utilitarianism. When, for example, we want to assess the moral correctness of proposed governmental legislation, we may wish to set aside any question of the motives of the legislators. After all, good laws may be passed for the most venal of political motives, and bad legislation may be the outcome of quite good motives. Instead, we can concentrate solely on the question of what *effects* the legislation may have on the people. When we make this shift, we are not necessarily denying that individual motives are important on some level, but rather, we are confining our attention to a level on which those motives become largely irrelevant. This is particularly appropriate in the case of policy decisions by governments, corporations, or groups. In such cases, there may be a diversity of different motives that one may want to treat as essentially private matters when assessing the moral worth of the proposed law, policy, or action. Therefore, rule utilitarianism's neglect of motives intuitively makes the most sense when we are assessing the moral worth of some large-scale policy proposed by an entity consisting of more than one individual.

Imagine again that the government is debating proposed legislation to change health-care benefits for the elderly. We could examine the motives of the individual legislators and lobbyists who support or oppose the new bill, but this might be of little help in determining whether the bill itself was a good one or not. We might discover, for example, that a large percentage of the people

who either support or oppose the legislation are motivated by a concern with political advancement or financial gain. Some may be acting out of principle, but the number may be roughly evenly split between the two camps. Looking at these individual motivations, we may simply decide that such motivations are just a matter of personal concern. They may tell us whether Senator X is morally well motivated in supporting this legislation, but they give us no insight into the question of whether the proposed law itself will prove to be a good or a bad one. Thus, we may decide simply to set aside the question of individual motives and turn to the law itself. Once we do this, one of the ways we may assess its moral worth is by looking at what results or consequences would follow from adopting the proposed law. We then assess these consequences in terms of some yardstick—for instance, in terms of the amount of pleasure and pain they may cause or in terms of the amount of happiness and unhappiness that could result from their adoption. This is roughly what the rule utilitarian seeks to do: to assess the moral worth of a policy or rule in terms of the consequences that it will probably have.

#### PRACTICES

Some philosophers have gone a step further than rule utilitarians by suggesting that utilitarian considerations have relevance in justifying the existence of certain types of practices, even though utilitarianism may not provide a proper basis for deciding particular acts within that practice. One can consider this as a type of rule utilitarianism, but it is important to notice that there is a significant difference between rules and practices. Rules are more specific than practices, and a practice may encompass numerous rules. Stamp collecting, for example, is a practice, and it contains many specific action-guiding rules about what types of stamps to buy, when to sell, and so on. Practices include rules but contain more as well. They are often embodied in specific institutions (philatelic societies) and in patterns of social interaction (e.g., stamp collectors' conventions) that go beyond any specific set of rules.

John Rawls, a contemporary philosopher whose *Theory of Justice* is one of the most influential recent works in ethics, has suggested that we may justify the *practice* of punishment as a whole through utilitarian arguments. A society without institutions and practices of punishment would produce less overall utility than one that contained such institutions and practices. Rawls avoids the problems raised by utilitarian justifications of specific acts of punishment by arguing that specific punishments be determined on the basis of retributive considerations, not on utilitarian grounds. The specific punishment would depend on the severity of the offense, not the utility of imposing the punishment. Specific individuals, in other words, would be punished because they deserved it, not because of the consequences produced by punishing them.

The merit of a suggestion like Rawls's is that it allows us to combine both utilitarian and Kantian insights. Utilitarian reasoning justifies the existence of the institution of punishment, and Kantian considerations of desert and retribution determine the nature and severity of specific acts of punishment.

When we summarize these various positions, we get the following chart.

#### CONSEQUENCES OF WHAT?

Type of Utilitarianism	Consequences of What?	Main Principle
Act	Each act	Perform the act that will produce the greatest overall amount of utility
Rule	Rules	Follow the rule that will produce the greatest overall amount of utility
Practice	Practices	Support the practices that will produce the greatest amount of utility

#### CONSEQUENCES FOR WHOM?

Utilitarianism is a type of **consequentialism**. It makes moral judgments on the basis of consequences. However, we need to ask: consequences for whom? Initially, that would seem to be a rather straightforward question with an equally straightforward answer: for people. Few things, though, are that easy. Let's examine three areas in which there is some controversy about who should count in determining consequences. The first of these relates to nonhumans. To what extent, if at all, should the suffering of animals count in our calculations of utility? Second, are all human beings included in our calculations of utility, or should subgroups (e.g., our nation or our family) be given special weight? Third, to what extent should we take into account the consequences for future generations? Let's look at each of these issues.

#### THE SUFFERING OF ANIMALS

Utilitarianism, at least in some of its principal forms, is dedicated to the reduction of suffering. Certainly, human beings suffer, but suffering does not appear to be limited to human beings. Animals suffer, even if their suffering is not exactly the same type as ours. How much, if any, weight should utilitarians give to animal suffering?

**PETS** Many of us, whether utilitarians or not, accord at least some value to the suffering of our pets. In fact, some of us will even go to great lengths, in terms of time as well as money, to preserve or restore the well-being of a pet. Although some of this may be a selfish concern for our own well-being if we were to lose the pet, in at least some cases, there is no reason to doubt that it is also a concern for the animal itself. We *care* about our pets, and consequently, their suffering counts. Yet the difficulty of saying this, for a utilitarian, is that this apparently violates the impartiality of the utilitarian outlook. From an impartial moral point of view, the suffering of one dog does not count any more

## 150 CHAPTER 5

than the same amount of suffering of any other dog. The only relevant difference is that in the case of pets, we need to add the suffering of their owners (and others who care about them) when the pets are in distress. There is a legitimate element in the utilitarian calculus, but it should be recognized that this element has nothing directly to do with the suffering of animals.

**PAINLESS KILLING** It is also not clear that a utilitarian perspective prevents the killing of animals as long as this is done in a way that does not cause them suffering. Take raising mink for their fur as an example. If the mink are bred and raised in conditions of comfort (a very large, enclosed area instead of individual cages) and if they are killed painlessly without feeling fear in advance, then it would seem that there is little a utilitarian could find objectionable in terms of suffering. We should realize, of course, that actual conditions rarely correspond to this example.

#### THE BORDERS OF OUR GROUP

In theory, utilitarianism maintains that we be impartially concerned with overall utility. There is nothing in the theory to suggest that we draw a line at national borders or at some other point. Indeed, the history of utilitarian thought is quite interesting in this respect. It flourished in Britain precisely during the period of the Empire and was particularly influential in shaping—and being shaped by—Britain’s rule of India. This reveals the potentially nondemocratic side of utilitarianism and suggests the importance of our answer to the question, “Who decides what the utility is?” The British were quite willing to rule India for the sake of the greatest utility—as long as they were the ones who decided what the utility was.

Today, the situation is not quite the same. Many of us, especially those in affluent countries, often seem to draw the boundaries of utility at the customs booth. However, there are few good arguments to justify this, except the claim that the citizens of a country are in the best position to decide what is best for that country. Yet we are faced with vast social and economic inequalities between countries, and there is little in utilitarianism to justify being concerned with only the welfare of one’s own country.

#### FUTURE GENERATIONS

The interests of future generations present a perplexing problem. On the one hand, they do not yet exist, and depending on our actions, they might never exist. Yet they are affected by the consequences of our actions, just as we (although presumably to a lesser extent, given differences in technological effectiveness) have been affected by our ancestors’ actions. When utilitarians compute consequences, should they take into account the consequences for future generations of people who are as yet unborn? If so, to what extent? These remain vexing questions not only for utilitarians but for many other moral philosophers as well.

## RACE, ETHNICITY, AND UTILITARIAN THEORY

What role, if any, do considerations of race and ethnicity have in utilitarian theory? Let's consider three related areas: race, ethnicity, and impartiality; the issue of group-specific consequences; and the question of who weighs the consequences.

### RACE, ETHNICITY, AND IMPARTIALITY

As we have seen, utilitarianism is, at heart, an impartial moral doctrine, and as such, it does not give any *special* weight to the concerns of any particular group, whether racial, ethnic, or cultural. But in many ways, its impartiality is also its potential strength for minority groups with little power, for utilitarianism *when properly applied* says that their suffering and unhappiness count just as much as the suffering and unhappiness of those who do hold the power and influence in society.

Strict adherence to utilitarian impartiality alone could bring significant advantages to minority groups, but this is not always true. Consider a typical situation in which the interests of minority groups have not counted on a par with those of the majority group. Imagine the planning of a new highway for which private lands have to be appropriated. Often, the lands appropriated for such projects are those that belong to poorer groups with less political influence. Does this violate utilitarian principles? Utilitarianism states that everyone's suffering is of equal weight (presuming it is of equal intensity). This means that the suffering a poor person of color experiences when uprooted is of equal value to the suffering a rich, white corporate executive experiences when uprooted, again presuming both have equally intense feelings about being relocated.

### GROUP-SPECIFIC CONSEQUENCES

One of the interesting issues in utilitarian theory is whether utilitarianism recognizes consequences that are harmful to racial, ethnic, or cultural groups as such and not just to the individual members of the group.

**THE EXAMPLE OF HATE SPEECH LEGISLATION** Hate speech legislation provides an example of this type of consideration. One of the proposed justifications for banning hate speech has been that it is typically demeaning and harmful to minority groups as groups, even in instances when we cannot show that some specific individual was directly harmed by the speech. In at least some versions, such protection is explicitly restricted to minority groups. The majority group, which in U.S. society is comprised of white males, is not offered the same protection.

For this type of argument to work, we need both empirical and normative premises. On the empirical side, we need premises that show a given behavior—such as hate speech—is in fact harmful to a specific group. This is primarily the domain of the social sciences, and there is extensive data to support claims of this sort.

On the normative side, we need a premise to the effect that the continued existence of a particular group is of value to society. Alternatively, we might substitute some premise to the effect that all groups are of (equal?) value to society and that there is a value in encouraging diversity as such. The potential difficulty with this argument is that it promotes skinheads and neo-Nazis as much as any other groups. Supporters of such an argument would presumably want to claim that there is a value in diversity as such. They might well point to a similar issue in environmental ethics, where some would maintain that we ought to encourage biological diversity on utilitarian grounds. We can never predict, they maintain, what might ultimately become biologically valuable to our survival. Consequently, we should strive to preserve all species as biological “money in the bank” against some future disaster. Similarly, professing epistemic ignorance, some may maintain that we should encourage the continued existence of all groups in society because they may ultimately preserve resources for society’s maintenance and flourishing that would otherwise be lost.

Given such premises, it seems possible in principle to develop sound utilitarian justifications for special treatment of racial, ethnic, or cultural groups.

#### WHO CALCULATES THE CONSEQUENCES?

One of the attractions of utilitarianism is that it promises impartiality and objectivity grounded in quantification. If everything can be translated into units of utility, then there is an objective basis for deciding between competing courses of action. Ultimately, the numbers decide for us.

Yet as the inhabitants of India well knew when the Empire applied its utilitarian logic to them, it makes a difference who is doing the calculating. When one group in society does the calculations for another group, it is all too easy for those calculations to become miscalculations. What mattered to the native inhabitants of India was very different from what the British thought mattered to them, as Gandhi was to prove.

Similar considerations have certainly applied in the United States, where we have seen time and again that one group (usually white, upper middle class, and male) has made decisions for other groups (including ethnic minorities). Even when it has done so with good will, it has often been wrong.

This is, at least in part, an epistemological point. Many would argue that those most directly affected by consequences are in the best position to estimate the importance of those consequences for themselves. There is also a moral and a psychological point here. The moral one is that those who will bear consequences should have a voice in determining those consequences. The psychological point is that people are more likely to bear onerous consequences when they themselves have had a voice in choosing them.

Thus, considerations of race, ethnicity, and culture have an important place in a utilitarian framework in regard to the calculator. All other things being equal, it is better that identifiable racial and ethnic groups be represented among the calculators of utility and that one group not assume the role of calculator for all other groups.



## ISSUES IN UTILITARIANISM

We have developed a fairly comprehensive picture of utilitarian ethical theory in the preceding pages, but there remain several key issues that need to be discussed if we are to round out our understanding of utilitarianism and evaluate its strengths and weaknesses. We focus on four of those issues here: (1) the difficulties of reasoning about matters of life and death, (2) the role of emotions in the utilitarian view of the moral life, (3) the limits of personal responsibility, and (4) the place of personal integrity in the utilitarian's world.

### WEIGHING MATTERS OF LIFE AND DEATH

We are sometimes hesitant about utilitarianism because it seems to weigh everything, even human life. Some things, many of us want to argue, cannot be put on the utilitarian scale. Kant, for example, clearly maintained that human beings were priceless; they could not be assigned a monetary value in the way that mere physical objects could. Indeed, as we shall see in the next chapter, this was for Kant one of the principal differences between people and things: Things have a price tag, but people do not. Yet to put human life into the utilitarian balance seems to come perilously close to placing a price tag on it.

At the beginning of this chapter, we discussed the example of removing asbestos from buildings. Some of you may have felt that human life is too precious to put into the balance and that any price is worth paying to save people's lives. Yet this is open to two rejoinders. First, sometimes we simply do not have an unlimited amount of funds; sometimes we have to make choices that result in letting some people die. But there is a second reply that is possible as well: Even though we say that human life is priceless, we don't really act that way. Let me give an example.

**THE SPEED LIMIT EXAMPLE** During the energy crisis of the 1970s, a national speed limit of 55 mph was instituted to conserve gasoline. One of the side effects of this change was that hundreds of lives a year were saved and thousands of traffic injuries were averted. Let's imagine that additional deaths and injuries could be avoided if we reduced the speed limit even further. At what point do we draw the line? When do we say that even if this means that a certain number of people will die or be injured in automobile accidents, we still want to keep the speed limit from being lowered?

My suspicion is that virtually all of us would draw the line at some point, even if it is considerably lower than the present limit. (Many would actually like the limit higher, and many drive as though it were.) My point here is not to argue in favor of any particular speed limit but rather to illustrate that most of us in fact are willing at some point to put human lives into the equation. Notice too that these may be innocent lives; there is no guarantee that those killed or injured in traffic accidents are necessarily the ones at fault. We should not be too quick to judge utilitarians harshly solely because they are willing to put human lives into the equation. Most of us are willing to do the same, even if we are reluctant to admit it.

**TERRORISM, TORTURE, AND CONSEQUENTIALISM** We are all familiar with the scenario. Terrorists have hidden a bomb in a crowded public place, but we don't know where. The clock is ticking, and we know the bomb will explode in ninety minutes. We have captured one of the terrorists, but he isn't talking. If we could only torture him, so the argument goes, we could find the bomb, disarm it, and save thousands of innocent lives. Wouldn't we be justified in torturing the terrorist, at least in this one case?

Alan Dershowitz, a Harvard Law School professor, has argued that torture ought to be permitted in such cases and that a judge's warrant should be required to ensure that the correct conditions are met. His reasoning is purely consequentialist: the harm saved (thousands of innocent lives saved) far outweighs the harm to the terrorist. Indeed, viewed in these stark consequentialist terms, it seems to be an easy choice: the suffering of one guilty man versus the lives of thousands of innocent men, women, and children.

But this type of reasoning also reveals some of the pitfalls of consequentialist thinking. First, the example presumes we know that the alleged terrorist is indeed responsible for the threat and knows where the bomb is, but in real life, there is usually an element of uncertainty. Only in retrospect do things appear so clearly, but in real life, we do not have the luxury of living retrospectively. We must live it prospectively, with uncertainty about how things will work out. Second, the example presumes that torture works. In actual practice, the information obtained under torture is often quite unreliable and either intentional deception meant to buy time or an attempt to say whatever the torturer wants simply to stop the torture.

This example points to a fundamental weakness of consequentialism. Taken by itself, it can justify actions that most of us would condemn. Some have even maintained that in certain circumstances consequentialism could justify slavery of a small minority if this brought sufficiently good consequences to the majority. Many philosophers have concluded that certain limits need to be imposed on consequentialist calculations. Human rights provide one of the most important of those limits for many of us. No matter what the consequentialists say, basic human rights ought not to be violated.

### THE ROLE OF EMOTIONS IN THE UTILITARIAN MORAL LIFE

Recently, some philosophers have criticized utilitarianism because of the role—or lack of a role—that it gives to emotions in the moral life. There are two aspects that merit our attention here. The first relates to the question of whether we ought to accord more weight to some emotions than others; the second issue centers on the nature of our relationship to our own emotions in utilitarianism. The latter issue, as we shall see, leads directly into the question of personal integrity.

**WHICH EMOTIONS ARE GIVEN WEIGHT?** Utilitarianism obviously gives moral weight to the emotions. If, for example, I am contemplating stealing some money from a friend, one of the things I would have to take into account as a

utilitarian is the suffering my friend would undergo at the loss of the money and the possibility of even greater suffering if he discovered that a friend had stolen it. I would also, of course, have to take into account my own feelings, possibly quite positive ones with a high hedonic value, if I were to succeed in stealing the money. Let's assume that my friend would have only negligibly negative feelings at having the money stolen and only slightly more negative feelings if he discovered that I was the thief. In addition, assume that I would certainly get great pleasure and happiness from successfully taking the money. Under these conditions, it would appear—all other things being equal, which they never are in real life—that my feelings could tip the balance in favor of committing the theft. What the utilitarian seems unable to do is distinguish between what we could call good and bad feelings or, more precisely, morally justified and morally unjustified emotions.

Examples of this problem need not be far-fetched, as the preceding example was. Consider the problem of racism in U.S. society. Certainly, there are many courses of action—such as genuine integration of schools—that would at least initially bring about intensely negative feelings by a large segment of the population. It is easy to imagine, if we were for the moment to grant the utilitarian premise that feelings can be placed on some hedonic scale, that integration would cause more negative feeling in the oppressors than it would cause positive feeling in the oppressed. Yet is this, again assuming other things equal, a sufficient reason for not seeking to eliminate racism from the society? (This, of course, does not even touch on the question of *who* weighs these feelings. One suspects that it would be much easier for a white person to sympathetically weigh the amount of fear and displeasure that whites experience at the thought of integration, whereas presumably a black person would be much more sensitive to the pain suffered by blacks. It is far from evident that there is a neutral standpoint here.)

The difficulty that the utilitarian faces is obvious: If all feelings are of equal value—and if all that distinguishes them is their sign (negative or positive), quantity (how many people experience the feeling), and intensity (how strong the feeling is for each)—then morally good and justifiable feelings will have no greater weight than morally evil and unjustifiable feelings of the same sign, quantity, and intensity. Yet we want to say, or at least our everyday moral intuitions suggest, that some types of feelings should be given greater weight than others.

The utilitarian seems at first to have an answer to this, but it is just the appearance of an answer. It seems that one could differentiate among various emotions on the basis of their overall social utility. Altruistic feelings, for example, may have greater social utility than discriminatory feelings. It may thus be beneficial from a utilitarian point of view to encourage courses of action that promote the development of altruistic feelings and reduce discriminatory feelings. Yet this will hardly do, for it really deals only with the question of the feelings that result from particular courses of action, not the feelings that might provide reasons for such actions. We could imagine, in the integration case, for example, that the results might not bring about a greater amount of altruistic

feelings. They could produce more bad feelings than good, but it might still be the right thing to do if we think that some of the negative feelings it produces are unjustified.

Thus, the first problem with the utilitarian account of the emotions is that it fails to provide any adequate way of discriminating between justified and unjustified feelings, between good and bad emotions, and is committed to giving both the same weight in its utilitarian calculations.

**MY RELATION TO MY OWN EMOTIONS AND CONVICTIONS** The second major problem with the utilitarian account of the relationship between morality and emotions centers around the question of how I am related to my own emotions, deeply held beliefs, and reasoned commitments. Bernard Williams develops this objection in some detail, and here, I only summarize the main elements of his argument.

Imagine the following kind of case, which is one of the type that Williams describes. A chemist—let's call him Harold—with a wife and children is out of work. He and his family are beginning to suffer significantly as a result of his unemployment. During the past few years, he has become increasingly convinced that all war, especially chemical and biological warfare, is immoral. An old friend offers him a job as a chemist in a firm developing and producing chemical warfare weapons. Furthermore, he is told by his old friend that, if he does not take the job, it will in all probability be given to a younger chemist he knows who is both a better chemist and quite committed to the development of such weapons. The moral quandary Harold faces is this. He is morally opposed to war and does not want to participate in developing weapons of war. His deepest moral commitments tell him not to take the job. However, he knows that if he takes the job, he will fulfill his obligations to his family, and at the same time, he will not be doing anything that will result in chemical warfare being more advanced than if he refused the job. Indeed, if Harold turns the job down, he will probably be helping the chemical warfare industry because his position will be filled by someone better and more enthusiastic than he is. Thus, the problem this example initially poses is this: What weight should Harold give to his own strong feelings against the morality of chemical warfare?

The standard utilitarian answer to this question recognizes Harold's commitments but only in a limited way. Clearly, these deeply held feelings and beliefs are one of the factors that Harold must take into consideration, but no special consideration must be given to the fact that they are *his* emotions and beliefs. He must also take into consideration everyone else's emotions and beliefs, and these presumably must be given equal weight. Indeed, in this situation, Harold might even be obliged to give less weight to his own feelings and convictions. After all, his own feelings are probably the ones he has greatest control over, so he may be obliged to try to change them to positive feelings about warfare. One thing is clear: His emotions do not deserve special weight simply by virtue of the fact that they are *his* emotions. Indeed, the utilitarian might well want to argue that this is the very essence of morality: impartiality. By not giving special weight to his own feelings, Harold is simply assuming the moral point of view. To be moral is to be impartial.

Yet critics such as Williams have suggested that there is still a serious problem here—namely, that utilitarian morality sometimes demands that a person give up his or her most deeply held feelings and convictions. So, if it is precisely these kinds of things that make life worth living for an individual, then utilitarian morality may demand that the individual give up his or her very reason for existing. Something is seriously wrong, Williams argues, if morality makes this kind of demand on an individual.

The reason utilitarianism falls short of the mark here is that it fails to recognize any special relationship between the agent and his or her own feelings and deeply held beliefs. It fails to recognize that these are my own in some unique way. When the utilitarian contemplates the consequences of an action (including the feelings it may create), there is nothing significant about the fact that some of the consequences may be *mine* in a special fashion. All feelings and commitments are taken into account simply as a group, summed up, and then the “bottom line” dictates the decision we should make.

#### THE LIMITS OF PERSONAL RESPONSIBILITY

In an intriguing example, Bernard Williams presents the following situation in which we must make a moral choice.

Jim finds himself in the central square of a small South American town. Tied up against the wall are a row of twenty Indians, most terrified, a few defiant, in front of them several armed men in uniform. A heavy man in a sweat-stained khaki shirt turns out to be the captain in charge and, after a good deal of questioning of Jim which establishes that he got there by accident while on a botanical expedition, explains that the Indians are a random group of the inhabitants who, after recent acts of protest against the government, are just about to be killed to remind other possible protesters of the advantages of not protesting. However, since Jim is an honored visitor from another land, the captain is happy to offer him a guest’s privilege of killing one of the Indians himself. If Jim accepts, then as a special mark of the occasion, the other Indians will be let off. Of course, if Jim refuses, then there is no special occasion, and Pedro here will do what he was about to do when Jim arrived, and kill them all. Jim, with some desperate recollection of schoolboy fiction, wonders whether if he got hold of a gun, he could hold the captain, Pedro, and the rest of the soldiers to threat, but it is quite clear from the set-up that nothing of that kind is going to work: any attempt at that sort of thing will mean that all the Indians will be killed, and himself. The men against the wall, and the other villagers, understand the situation, and are obviously begging him to accept. What should he do?

The utilitarian, Williams maintains, has a clear and easy answer to this question. Within the utilitarian perspective, there is no question that Jim should shoot the one prisoner so that the others could go free. It is not only the right thing to do, but it is the *obviously* right thing to do. When we add up the hedons and dolors in even the most cursory way, we see clearly that the alternative that will produce the greatest overall amount of utility is for Jim to pull the trigger. Nineteen lives would be saved. Everyone wants him to do it.

## 158 CHAPTER 5

The question that Williams raises in this example is an important one. He does not deny that shooting the one villager may be the best alternative, but he points out that in the utilitarian perspective Jim cannot attach any special weight to the fact that *he* is the one who pulls the trigger. Jim is equally responsible for the deeds that he does directly and the deeds performed by others that he could have prevented. Jim is responsible for Pedro's killing the twenty prisoners because he could have prevented it himself by killing one of the prisoners. This, Williams contends, is a confused notion of personal responsibility. We bear a special relationship to, and responsibility for, our own actions that utilitarianism fails to capture.

**THE SIAMESE TWINS** We see this same problem occurring in real life. In the summer of 2000, a couple in Malta had Siamese twins, Jodie and Mary. The family came to England for medical assistance. The doctors told the parents that if the twins were not surgically separated in the next few weeks, both would die. If the doctors operated to separate the twins, then the stronger twin had a good chance of surviving, although she would be severely disabled and have to undergo multiple surgeries. The parents, devout Catholics, did not want the surgery because they felt it would be killing one of the two twins. The British government, in a controversial decision, ordered the surgery against the parents' wishes; it also issued an order forbidding them from leaving England with the twins to avoid the surgery. The surgery was performed. This resulted in the death of the weaker twin; the stronger twin survived and is now doing better than expected.

From a utilitarian point of view, this decision should be an easy one: the possible life of one child versus the certain death of both. But for the parents, it was an agonizing decision precisely because of the issue of responsibility: They refused to make a choice that would result in the death of one of their children to save the life of the other. This is precisely the issue of responsibility that Williams discussed. The issue is not simply that one child dies but that one child dies as a result of the parents' decision.

#### INTEGRITY AND IMPARTIALITY

Two final characteristics of utilitarianism are particularly important to note here. First, utilitarians are not allowed to give any special weight to the fact that certain consequences may affect them personally. The popular image of utilitarians is often of people who are just concerned with achieving their own selfish aims and who then view everything else simply as a means to the attainment of those ends. Yet the picture we get of utilitarians from an ethical standpoint is quite different. Utilitarians are not allowed to give any special weight to the fact that some negative consequences will affect them quite personally. If, for example, we have a utilitarian legislator who will personally suffer if there is not an increase in medical benefits for the aged, the utilitarian legislator will still be required to vote against such an increase if that increase would yield less total utility than the alternatives. If, to take a second example,

we have a utilitarian gourmet who is contemplating either having dinner at Chez Panisse, an expensive restaurant in Berkeley, or donating the dinner money to a charity devoted to relieving hunger in the world, from a utilitarian standpoint there is only one morally right alternative: to give the money to help reduce hunger. Even though dinner at Chez Panisse may yield one hundred hedons for the gourmet, the same amount of money may well bring ten hedons for thirty people. Thus, utilitarians cannot give special weight to the fact that certain pleasures or displeasures are their own; they must be weighed just like everyone else's hedons and dolors.

**PETER SINGER AND WORLD HUNGER** An excellent example of how demanding utilitarianism can be in this area can be found in the work of Peter Singer. He sees the issue of world hunger from a utilitarian perspective, and doing so imposes strong obligations on those of us who live in more affluent countries. Singer begins with a simple question:

If I am walking past a shallow pond and see a child drowning in it, I ought to wade in and pull the child out. That will mean getting my clothes muddy, but this is insignificant when the death of the child would presumably be a very bad thing.

Indeed, more generally, we could say that if we can prevent something bad from happening without in the process giving up something of comparable importance, then we ought to do so. Singer then goes on to argue that we can eliminate some world hunger and absolute poverty without throwing ourselves into a state of equal deprivation; therefore, we ought to do so.

Indeed, we spend money on many luxuries in our society that we could do without. Some could drive reliable, moderately priced automobiles instead of high-performance vehicles; others could have healthful but plain meals at home instead of expensive (and often unhealthy) meals at fancy restaurants; some could forgo having the latest gadget, and perhaps substitute talent and the hard work of practice for high-priced equipment in sports and other areas. Indeed, if we think of this issue in purely impartial consequentialist terms, it seems clear that we should redirect our expenditures from luxury items to their more basic equivalents (a Toyota Camry instead of a BMW convertible) and then give the remainder to relief organizations. This raises interesting and important questions about how much a moral theory can demand of us.

**FUNDAMENTAL PROJECTS** When we hear that a concert pianist was involved in an auto accident that crushed her hands, that a famous painter has gone blind, or that a well-known baseball pitcher has suffered irreparable damage to his pitching arm, we are especially moved. Our heart goes out to such people because we realize how such an accident strikes at the very heart of who they are as persons. Indeed, we want to say that certain projects, commitments, and desires are closer than others to individuals' personal sense of identity (i.e., their idea of who they are). Those closest to the person's sense of his or her own identity comprise what we shall call that person's **fundamental projects**.

The issue of fundamental projects becomes even more vivid if we recall Williams's example of Jim in the South American town. The way Williams sets up the example, we are asked to imagine what Jim (a hypothetical character) would do. But let's change the example a bit. Imagine two different scenarios. First, instead of Jim, imagine a person whose whole life was devoted to peace and nonviolence. Mother Teresa or Martin Luther King Jr. come to mind as obvious examples. What should *they* do? Second, imagine that a mercenary soldier, for whom killing is a casual activity, arrives in the village instead of Jim. What should the soldier do? Most of us would give quite different answers to these questions depending on whether the visitor was someone like Mother Teresa or Dr. King or was like the mercenary. The reason for our different answers is precisely the issue of fundamental projects. To kill anyone would run counter to the whole sense of what Mother Teresa's or Dr. King's lives were about. It would not, however, contradict the mercenary's life at all (except perhaps in the fact that he isn't being paid for it). Utilitarianism seems to give insufficient recognition to this difference in fundamental projects.

## LIVING THE UTILITARIAN LIFE

When we are considering so many arguments for and against utilitarianism and drawing so many distinctions between various types of utilitarianism, it is easy to lose sight of what it means to live life as a utilitarian. Yet utilitarianism is a moral theory that was meant to be lived, and a consideration of what it would look like in practice can provide us with a good way of drawing together some of our conclusions about utilitarianism.

Two insights guide the utilitarian's life. The first is that *consequences count*. Consequently, *utilitarians will always want to know what actual effects their choices will have for real people* (and perhaps other sentient beings as well). They continually direct their attention to the basic facts of the moral life about who will be hurt and who will be made happy as a result of a particular decision. It may be very difficult at times to predict what the actual consequences of a particular decision will be, but utilitarians are committed to trying to make such predictions as accurately as possible. To the extent that they are not able to make such predictions accurately, it indicates that there is an unavoidable element of luck in the moral life. The presence of luck is a problem with the moral life but not an objection to utilitarianism as a moral theory.

Second, *utilitarians want the world to be a better place for everyone*. It is a benevolent moral doctrine; that is, it wishes people well and seeks to increase the amount of well-being in the world. Indeed, this is the whole point of morality for utilitarians: It produces a better, ultimately happier world. Ethical reflection is not something pursued in abstraction from the real pain and suffering of the world around us. The point of ethics is to help reduce that pain and suffering. Morality should make the world a better place for everyone.

Despite these strengths, many find that utilitarianism does not provide the complete story of the moral life. One of the dangers that many utilitarians,



especially act utilitarians, face is that their principles might require actions that violate the rights of small groups of individuals. Act utilitarianism alone cannot provide sufficient guarantees against the possibility of such abuse. There are two ways to respond to this difficulty. Some philosophers have opted for some version of rule utilitarianism that seems less susceptible to such difficulties. Others have suggested that *there must be a moral “floor” or minimum below which we cannot go*, even if utilitarian considerations seem to demand that we do so. To live only by utilitarian considerations, especially act utilitarian ones, is to open the door to possible abuses of the minority when such injustices yield high benefits for the majority. Ethical theories that emphasize the importance of human rights seem to offer a standard of value in the moral life that escapes from these dangers. Chapter 7 of this book is devoted to rights-based moral theories.

Another difficulty plagues utilitarian accounts of ethics. *Utilitarians ignore the importance of motives*. Although there are plenty of circumstances when motives do not matter, there are times—especially in personal relationships—when they are of crucial importance. This dimension of the moral life has been almost completely ignored by utilitarians because of their exclusive focus on consequences. Three quite different approaches to morality help in understanding this dimension of the moral life better than utilitarians alone have been able to do. First, Kantian accounts of morality, as we shall see in the next chapter, emphasize the importance of motives in the moral life, especially the importance of acting out of a motive of duty. Kantians see something morally admirable about acting for the sake of duty that utilitarians are unable to recognize. Second, some critics of morality have argued—as we shall see in Chapter 8—that all major moral theories fail to provide an adequate account of our moral motivations. Such theories, critics like Michael Stocker argue, produce a kind of motivational schizophrenia, a deep and pervasive split between our actual motives and the legitimate reasons within any particular moral theory. Such criticisms can be answered, I shall argue, only by a moral theory that focuses primarily on character. This, as we shall see in Chapter 9, is precisely the kind of theory that Aristotle offers.

## DISCUSSION QUESTIONS

1. Recall your rating of Ethics Inventory STATEMENT 16: “When I am trying to decide what the right thing to do is, I look at the consequences of the various alternatives open to me.”
  - a. What moral theory does this statement illustrate?
  - b. Has your rating of this item changed after reading this chapter? If so, in what way? If your rating has not changed, are your reasons for the rating any different now from when you first responded to this statement?
2. Recall your rating of STATEMENT 17: “The right thing to do is whatever is best for everyone.”

## 162 CHAPTER 5

- a. In what types of cases, if any, are we not justified in doing what will produce the greatest overall amount of good? Be specific.
- b. Has your rating of this item changed after reading this chapter? If so, in what way? If your rating has not changed, are your reasons for the rating any different now from when you first responded to this statement?
3. Recall STATEMENTS 19 and 20 about whether pleasure or happiness is the most important thing in life.
  - a. If you agreed with STATEMENT 19, what arguments do you now see that could be advanced against your position? How would you reply to these arguments?
  - b. Have your ratings of these items changed after reading this chapter? If so, in what way? If your ratings have not changed, are your reasons the your ratings any different now from when you first responded to these statements?
  - c. If you hold that neither pleasure nor happiness is of intrinsic value, what is? Explain.
4. Take a contemporary social issue such as kidney transplants that involves the allocation of scarce resources and discuss the ways various types of utilitarians would recommend that we deal with it. How would their recommendations differ from the recommendations of ethical egoists? Which of these traditions do you find more convincing? Why? If you disagree with either, what are your reasons for disagreement?
5. In recent decades, Americans have been reconsidering their treatment of persons with physical handicaps or disabilities in a number of different areas of life, including education and sports. The Americans with Disabilities Act (1990) has been a major factor in requiring equal access to public facilities for persons with disabilities. Imagine that a proposal has been put to your local school board to institute a limited sports program for students with physical impairments. The projected cost of running such a program would be approximately four times as much per student as is spent on the regular sports programs, although the number of students is much lower. How would a sophisticated utilitarian deal with this proposal?
6. Benjamin Franklin (1706–1790) said, “Honesty is the best policy.” Is this true only on utilitarian grounds (as Franklin thought)? Or do we need some other justification for it? When, if ever, do you think honesty if *not* the best policy? Give an example.
7. **MOVIE** In the movie *Saving Private Ryan*, Captain John Miller (Tom Hanks) and seven other men are sent on a mission to find and bring to safety Private James Francis Ryan (Matt Damon), whose three other brothers have already been killed in combat. One member of the unit asks Captain Miller, “Explain the math of this to me—risking the eight of us for one life.” What would a utilitarian say about the mission to save Private Ryan? Do you agree with the utilitarian analysis in this instance?

8. Human life, philosopher Immanuel Kant tells us, is priceless. However, we often seem to put a price tag on human life. Is this always wrong? Why or why not? If it is ever morally permissible to do so, when is this allowed? Why? If we don't put a price tag on human life, how do we deal with
  - a. the allocation of scarce medical resources for which we have the money to save only some of the people?
  - b. jury awards in wrongful death suits?
9. Imagine you are a utilitarian who has \$10 million to spend on health care for infants. Which would be better: spending it on extensive prenatal care or high-technology neonatal intensive care units? To answer this question, what further questions would you have to ask about each alternative? Would you agree with the utilitarian solution to this question?
10. **MOVIE** In the movie *Gattaca*, genetically engineered individuals ("valids") are given preference over naturally born individuals ("in-valids"). Parents are urged to have their children genetically engineered. On what basis should we judge whether to use genetic engineering or not? Is it simply a matter of the consequences, or are there other relevant moral considerations?

## BIBLIOGRAPHICAL ESSAY

The *classic texts* for utilitarianism are those of Jeremy Bentham, John Stuart Mill, and Henry Sidgwick. Among *Bentham's* works, see in particular, *A Fragment on Government*, edited by J. H. Burns and H. L. A. Hart (London: Athlone Press, 1977), and his *The Introduction to the Principles of Morals and Legislation*, edited by J. H. Burns and H. L. A. Hart (London: Athlone Press, 1970). (Both are also available in other less expensive editions.) For excellent *introductions to Bentham's moral and political thought*, see John Dinwiddy, *Bentham* (Oxford: Oxford University Press, 1989), and Ross Harrison, *Bentham* (London: Routledge & Kegan Paul, 1984). Also see David Lyons, *In the Interest of the Governed* (Oxford: Clarendon Press, 1973); H. L. A. Hart, *Essays on Bentham: Jurisprudence and Political Theory* (Oxford: Clarendon Press, 1982); and Ross Harrison's "Bentham, Mill and Sidgwick," in *The Blackwell Companion to Philosophy*, edited by Nicholas Bunnin and E. P. Tsui-James (Oxford: Basil Blackwell, 1996), 627–42.

Many of *John Stuart Mill's* works are relevant, especially his *Utilitarianism* and *On Liberty*. These are available on the World Wide Web and in various bound editions, including several that also contain critical essays. See John Stuart Mill, *Utilitarianism: Text with Critical Essays*, edited by Samuel Gorovitz (Indianapolis, IN: Bobbs-Merrill, 1971); Mill's *Utilitarianism: Text and Criticism*, edited by James M. Smith and Ernest Sosa (Belmont, CA: Wadsworth, 1969); and *On Liberty: Annotated Text, Sources and Background*, edited by David Spitz (New York: Norton, 1975). For an excellent selection of Mill's writings on ethics, see *Mill's Ethical Writings*, edited by J. B. Schneewind (New York: Collier, 1965). Among the best books on Mill's philosophy is Fred Berger's *Happiness, Justice, and Freedom: The Moral and Political Philosophy of John Stuart Mill* (Berkeley: University of California Press, 1984). For a brief but comprehensive overview of Mill's thought, see Henry West, "Mill, John Stuart," in *Encyclopedia of Ethics*, edited by Lawrence C. Becker and Charlotte B. Becker (New York: Garland Publishing, 1992), vol. 2, 809–16.

Book 4 of *Henry Sidgwick's The Methods of Ethics*, 7th ed. (Indianapolis, IN: Hackett, 1981), is also a classic source of utilitarian thought. For a fine introduction to Sidgwick's thought and times, see J. B. Schneewind, *Sidgwick's Ethics and Victorian Moral Philosophy* (Oxford: Clarendon Press, 1977). For a brief introduction to Sidgwick's thought, see Marcus G. Singer, "Sidgwick," in *Encyclopedia of Ethics*, edited by Lawrence C. Becker and Charlotte B. Becker (New York: Garland Publishing, 1992), vol. 2, 1149–52.

**World Wide Web.** Many of the classical works in utilitarianism are available on the World Wide Web. An up-to-date list of such works is available on the utilitarianism page of *Ethics Updates* at [ethics.SanDiego.edu/theories/Utilitarianism/](http://ethics.SanDiego.edu/theories/Utilitarianism/). This also contains links to utilitarianism websites and other related resources.

Several of the editions of Mill's *Utilitarianism* and *On Liberty* contain excellent collections of *critical essays*. For an excellent collection of essays on Mill's *Utilitarianism*, see David Lyons, ed., *Mill's "Utilitarianism": Critical Essays* (Lanham, MD: Rowman & Littlefield, 1997); for essays on Mill's *On Liberty*, see Mill's *On Liberty*, edited by Gerald Dworkin (Lanham, MD: Rowman & Littlefield, 1997). A good starting point for studying utilitarianism is J. J. C. Smart and Bernard Williams, *Utilitarianism: For and Against* (Cambridge: Cambridge University Press, 1973), which contains an explication and defense of act utilitarianism by Smart and an interesting critique by Williams. The essay by Williams has been one of the most influential in raising the issues of moral alienation and integrity. One of the most nuanced and powerful replies to Williams and others on this issue is Peter Railton's "Alienation, Consequentialism, and the Demands of Morality," *Philosophy & Public Affairs*, 13, no. 2 (Spring 1984), 134–71. This essay, along with a number of other important pieces, has been reprinted in an excellent anthology edited by Samuel Scheffler, *Consequentialism and Its Critics* (Oxford: Oxford University Press, 1988); also see David O. Brink, "Utilitarian Morality and the Personal Point of View," *Journal of Philosophy*, 83 (1986), 417–38. The anthology edited by Bernard Williams and Amartya Sen, *Utilitarianism and Beyond* (Cambridge: Cambridge University Press, 1982), contains a number of perceptive articles. For a helpful anthology of essays on the place of rights in utilitarian moral theory, see *Utility and Rights*, edited by R. G. Frey (Minneapolis: University of Minnesota Press, 1984); Richard B. Brandt, *Morality, Utilitarianism, and Rights* (New York: Cambridge University Press, 1992); and *Rationality, Rules, and Utility: New Essays on the Moral Philosophy of Richard B. Brandt*, edited by Brad Hooker (Boulder, CO: Westview Press, 1993).

For a very reliable, brief *overviews of utilitarian thought*, see David Lyons, "Utilitarianism," in *Encyclopedia of Ethics*, edited by Lawrence C. Becker and Charlotte B. Becker (New York: Garland Publishing, 1992), vol. 2, 1261–68; Philip Pettit, "Consequentialism," in *A Companion to Ethics*, edited by Peter Singer (Oxford: Basil Blackwell, 1991), 230–40. R. G. Frey's "Act-Utilitarianism" and Brad Hooker's "Rule-Consequentialism" are in *The Blackwell Guide to Ethical Theory*, edited by Hugh LaFollette (Oxford: Blackwell, 2000), 165–82 and 183–204, respectively. Online, see Walter Sinnott-Armstrong, "Consequentialism," in *The Stanford Encyclopedia of Philosophy* (Winter 2006 ed.), Edward N. Zalta (ed.), at [plato.stanford.edu/entries/consequentialism/](http://plato.stanford.edu/entries/consequentialism/), which contains an excellent treatment of consequentialism in general and utilitarianism in particular. William H. Shaw's *Contemporary Ethics: Taking Account of Utilitarianism* (Oxford: Basil Blackwell, 1999) is an excellent and highly readable introduction to both the theory and the application of utilitarianism. Russell Hardin's *Ethics Within the Limits of Reason* (Chicago: University of Chicago Press, 1988) develops a defense of utilitarianism that relies heavily on game theory. Robert E.

Goodin provides a strong defense of a utilitarian approach to public policy issues in *Utilitarianism as a Public Philosophy* (Cambridge: Cambridge University Press, 1995). For a very perceptive discussion of *well-being* in relationship to utilitarianism, see James Griffin, *Well-Being: Its Meaning, Measurement and Moral Importance* (Oxford: Clarendon Press, 1986). For an overview of utilitarian thought that includes other traditions (including the Chinese), see Geoffrey Scarre, *Utilitarianism* (New York: Routledge, 1996). Also see *Recent Work on Intrinsic Value* (Library of Ethics and Applied Philosophy), edited by Toni Rønnow-Rasmussen and Michael J. Zimmerman (Dordrecht, The Netherlands: Springer, 2006).

Richard Brandt introduced *the distinction between act and rule utilitarianism* in his *Ethical Theory* (Englewood Cliffs, NJ: Prentice Hall, 1959). For a defense of rule utilitarianism, see Brad Hooker, *Ideal Code, Real World: A Rule-Consequentialist Theory of Morality* (Clarendon: Oxford University Press, 2000).

The discussion of utilitarianism often takes place within the context of a contrast with Kantian and other deontological accounts of morality. John Rawls's "Two Concepts of Rules," *Philosophical Review*, 64 (1955), 3–22, is an important attempt to reconcile partially these two traditions. Samuel Scheffler's *The Rejection of Consequentialism* (Oxford: Clarendon Press, 1982) provides a provocative rethinking of some of these issues, as does Michael Slote's *Common-Sense Morality and Consequentialism* (London: Routledge & Kegan Paul, 1985), which contains an extended discussion of satisficing consequentialism. David Gauthier, *Kantian Consequentialism* (New York: Oxford University Press, 1996), approaches this issue from the other side, exploring the role of consequences in Kant's thinking.

**Citations.** The example of Jim is found in Bernard Williams, "A Critique of Utilitarianism," in J. J. C. Smart and Bernard Williams, *Utilitarianism: For and Against* (Cambridge: Cambridge University Press, 1973), 98–99. On the British, utilitarianism, and India, see Ragavan Iyer, *Utilitarianism and All That: The Political Theory of British Imperialism* (Santa Barbara, CA: Concord Grove Press, 1983). The quotation from Singer is taken from Peter Singer, "Famine, Affluence, and Morality," *Philosophy and Public Affairs*, 1 (1972), 231.

**Further Information.** For additional information about the issues raised in this chapter, see the page on utilitarianism of my *Ethics Updates* site on the World Wide Web at [ethics.SanDiego.edu/theories/Utilitarianism/](http://ethics.SanDiego.edu/theories/Utilitarianism/).