# Philosophical Ethics

What constitutes an ethical choice? Does an ethical choice maximize happiness? Goodness? Does it follow some other universal principle? Is a universally applicable principle even possible? Ethical philosophers have long wrestled with these questions.

This section examines three ethical models: utilitarian, Kantian, and Nicomachean. In studying the following selections, consider how they might be applied to computer technologies. Following Aristotle's thinking, for example, can one conclude that hacking is ethical when it stems from a rational challenge conquered by a clever programmer? (See also Spafford, "Are Hacker Break-ins Ethical?")

## 1

## The Best Action is the One with the Best Consequences

Of those actions available to you, you are morally obliged to choose that action which maximizes total happiness (summed over all affected persons) according to utilitarian ethical theory. The utilitarian model is particularly useful in illuminating instances when many people are affected in different ways by an action; for example, a utilitarian analysis may be useful in deciding what the laws ought to be on copyright (see National Research Council, "Music: Intellectual Property's Canary in the Digital Coal Mine") and privacy (see Garfinkel, "Privacy in a Database Nation").

### *John Hospers*

Once one admits that one's own personal good is not the only consideration, how can one stop short of the good of everyone—"the general good"? This conclusion, at any rate, is the thesis of the ethical theory known as *utilitarianism*. The thesis is simply stated, though its application to actual situations is often extremely complex: whatever is intrinsically good should be promoted, and, accordingly, our obligation (or duty) is always to act so as to promote the greatest possible intrinsic good. It is never our duty to promote a lesser good when we could, by our action, promote a greater one; and the act which we should perform in any given situation is, therefore, the one which produces more intrinsic good than any other act we could have performed in its stead. In brief, the main tenet of utilitarianism is the maximization of intrinsic good.

The description just given is so brief that it will almost inevitably be misleading when one attempts to apply it in actual situations unless it is spelled out in greater detail. Let us proceed at once, then, to the necessary explanations and qualifications.

1. When utilitarians talk about right or wrong acts, they mean—and this point is shared by the proponents of all ethical theories—*voluntary* acts. Involuntary acts like the knee jerk are not included since we have no control over them: once the stimulus has occurred the act results quite irrespective of our own will. The most usual way in which the term "voluntary act" is defined is as follows:[1] an act is voluntary if the person *could* have acted differently *if* he had so chosen. For example, I went shopping yesterday, but if I had chosen (for one reason or another) to remain at home, I would have done so. My choosing made the difference. Making this condition is not the same as saying that an act, to be voluntary, must be *premeditated* or that it must be the outcome of *deliberation,* though voluntary acts often are planned. If you see a victim of a car accident lying in the street, you may rush to help him at once, without going through a process of deliberation; nevertheless your act is voluntary in that if you had chosen to ignore him you would have acted differently. Though not premeditated, the action *was* within your control. "Ought implies can," and there is no ought when there is no can. To be right or wrong, an act must be within your power to perform: it must be performable as the result of your choice, and a different choice must have led to a different act or to no act at all.

2. There is no preference for immediate, as opposed to remote, happiness. If Act A will produce a certain amount of happiness today and Act B will produce twice as much one year hence, I should do B, even though its effects are more remote. Remoteness does not affect the principle at all: happiness is as intrinsically good tomorrow or next year as it is today, and one should forego a smaller total intrinsic good now in favor of a larger one in the future. (Of course, a remote happiness is often less certain to occur. But in that case we should choose A not because it is more immediate but because it is more nearly certain to occur.) . . .

3. Unhappiness must be considered as well as happiness. Suppose that Act A will produce five units of happiness and none of unhappiness and Act B will produce ten units of happiness and ten of unhappiness. Then A is to be preferred because the *net* happiness—the resulting total after the unhappiness has been subtracted from it—is greater in A than in B: it is five in A and zero in B. Thus the formula "You should do what will produce the greatest total happiness" is not quite accurate; you should do what will produce the most *net* happiness. This modification is what we shall henceforth mean in talking about "producing the greatest happiness"—we shall assume that the unhappiness has already been figured into the total.

4. It is not even accurate to say that you should always do what leads to



the greatest *balance* of happiness over unhappiness, for there may be no such balance in any alternative open to the agent: he may have to choose between "the lesser of two evils." If Act A leads to five units of happiness and ten of unhappiness and Act B leads to five units of happiness and fifteen of unhappiness, you should choose A, not because it produces the most happiness (they both produce an equal amount) and not because there is a greater balance of happiness over unhappiness in A (there is a balance of unhappiness over happiness in both), but because, although both A and B produce a balance of unhappiness over happiness, A leads to a *smaller* balance of unhappiness over happiness than B does. Thus we should say, "Do that act which produces the greatest balance of happiness over unhappiness, or, if no act possible under the circumstance does this, do the one which produces the smallest balance of unhappiness over happiness." This qualification also we shall assume to be included in the utilitarian formula from now on in speaking of "producing the greatest happiness" or "maximizing happiness."

5. One should not assume that an act is right according to utilitarianism simply because it produces more happiness than unhappiness in its total consequences. If one did make this assumption, it would be right for ten men collectively to torture a victim, provided that the total pleasure enjoyed by the sadists exceeded the pain endured by the victim (assuming that pain is here equated with unhappiness and that all the persons died immediately thereafter and there were no further consequences). The requirement is not that the happiness exceed the unhappiness but that it do so *more* than any other act that could have been performed instead. This requirement is hardly fulfilled here: it is very probable indeed that the torturers could think of something better to do with their time.

6. When there is a choice between a greater happiness for yourself at the expense of others, and a greater happiness for others at the expense of your own, which should you choose? You choose, according to the utilitarian formula, whatever alternative results in the greater total amount of *net* happiness, precisely as we have described. If the net happiness is greater in the alternative favorable to yourself, you adopt this alternative; otherwise not. Mill says, "The happiness which forms the utilitarian standard of what is right in conduct, is not the agent's own happiness, but that of all concerned. As between his own happiness and that of others, utilitarianism requires him to be as strictly impartial as a disinterested and benevolent spectator."[2] To state this in different language, you are not to ignore your own happiness in your calculations, but neither are you to consider it more important than anyone else's; you count as one, and only as one, along with everyone else. Thus if Act A produces a total net happiness of one hundred, and Act B produces seventy-five, A is the right act even if you personally would be happier in consequence of B. Your choice should not be an "interested"

one; you are not to be prejudiced in favor of your own happiness nor, for that matter, against it; your choice should be strictly *dis*interested as in the case of an impartial judge. Your choice should be dictated by the greatest-total-happiness principle, not by a *your*-greatest-happiness principle. If you imagine yourself as a judge having to make a decision designed to produce the most happiness for all concerned *without* knowing which of the people affected would be *you,* you have the best idea of the impartiality of judgment required by the utilitarian morality.

In egoistic ethics . . . your sole duty is to promote your own interests as much as possible, making quite sure, of course, that what you do will make you really happy (or whatever else you include in "your own interest") and that you do not choose merely what you *think* at the moment will do so; we have called this policy the policy of "*enlightened* self-interest." In an *altruistic* ethics, on the other hand, you sacrifice your own interests completely to those of others: you ignore your own welfare and become a doormat for the fulfillment of the interests of others. . . . But the utilitarian ethics is neither egoistic nor altruistic: it is a *universalistic* ethics, since it considers your interests equally with everyone else's. You are not the slave of others, nor are they your slaves. Indeed, there are countless instances in which the act required of you by ethical egoism and the act required by utilitarianism will be the same: for very often indeed the act that makes you happy will also make those around you happy, and by promoting your own welfare you will also be promoting theirs. (As support for this position, consider capitalistic society: the producer of wealth, by being free to amass profits, will have more incentive to produce and, by increasing production, will be able to create more work and more wealth. By increasing production, he will be increasing the welfare of his employees and the wealth of the nation.) Moreover, it is much more likely that you can effectively produce good by concentrating on your immediate environment than by "spreading yourself thin" and trying to help everyone in the world: "do-gooders" often succeed in achieving no good at all. (But, of course, sometimes they do.) You are in a much better position to produce good among those people whose needs and interests you already know than among strangers; and, of course, the person whose needs and interests you probably know best of all (though not always) is yourself. Utilitarianism is very far, then, from recommending that you ignore your own interests.

It is only when your interests cannot be achieved except at the cost of sacrificing the *greater* interests of others that utilitarianism recommends self-sacrifice. When interests conflict, you have to weigh your own interest against the general interest. If, on the one hand, you are spending all your valuable study time (and thus sacrificing your grades and perhaps your college degree) visiting your sick aunt because she wants you to, you would probably produce more good by spending your time studying. But on the other hand, if an undeniably greater good will result from your sacrifice, if, for instance, your mother is seriously ill and no one else is available to care for her, you might have to drop out of school for a semester to care for her. It might even, on occasion, be your utilitarian duty to sacrifice your very life for a cause, when the cause is extremely worthy and requires your sacrifice for its fulfillment. But your must first make quire sure that your sacrifice will indeed produce the great good intended; otherwise you would be throwing your life away uselessly. You must act with your eyes open, not under the spell of a martyr complex.



. . . . .

7. The general temper of the utilitarian ethics can perhaps best be seen in its attitude toward moral rules, the traditional dos and don'ts. What is the utilitarian's attitude toward rules such as "Don't kill," "Don't tell lies," "Don't steal"?

According to utilitarianism, such rules are *on the whole* good, useful, and worthwhile, but they *may* have exceptions. None of them is sacrosanct. If killing is wrong, it is not because there is something intrinsically bad about killing itself, but because killing leads to a diminution of human happiness. This undesirable consequence almost always occurs: when a man takes another human life, he not only extinguishes in his victim all chances of future happiness, but he causes grief, bereavement, and perhaps years of misery for the victim's family and loved ones; moreover, for weeks or months countless people who know of this act may walk the streets in fear, wondering who will be the next victim—the amount of insecurity caused by even one act of murder is almost incalculable; and in addition to all this unhappiness, every violation of a law has a tendency to weaken the whole fabric of the law itself and tends to make other violations easier and more likely to occur. If the guilty man is caught, he himself hardly gains much happiness from lifelong imprisonment, nor are other people usually much happier for long because of his incarceration; and if he is not caught, many people will live in fear and dread, and he himself will probably repeat his act sooner or later, having escaped capture this time. The good consequences, if any, are few and far between and are overwhelmingly outweighed by the bad ones. Because of these prevailingly bad consequences, killing is condemned by the utilitarian, and thus he agrees with the traditional moral rule prohibiting it.

He would nevertheless admit the possibility of exceptions: if you had had the opportunity to assassinate Hitler in 1943 and did not, the utilitarian would probably say that you were doing wrong in *not* killing him. By not killing him, you would be stealing the death of thousands, if not millions, of other people: political prisoners and Jews whom he tortured and killed in concentration camps and thousands of soldiers (both Axis and Allied)

whose lives would have been saved by an earlier cessation of the war. If you had refrained from killing him when you had the chance, saying "It is my duty never to take a life, therefore I shall not take his," the man whose life you saved would then turn around and have a thousand others killed, and for his act the victims would have you to thank. Your conscience, guided by the traditional moral rules, would have helped to bring about the torture and death of countless other people.

Does the utilitarian's willingness to adopt violence upon occasion mean that a utilitarian could never be a pacifist? Not necessarily. He *might* say that *all* taking of human life is wrong, but if he took this stand, he would do so because he believed that killing *always* leads to worse consequences (or greater unhappiness) than not killing and *not* because there is anything intrinsically bad about killing. He might even be able to make out a plausible argument for saying that killing Hitler would have been wrong: perhaps even worse men would have taken over and the slaughter wouldn't have been prevented (but then wouldn't it have been right to kill *all* of them if one had the chance?); perhaps Hitler's "intuitions" led to an earlier defeat for Germany than if stabler men had made more rationally self-seeking decisions on behalf of Nazi Germany; perhaps the assassination of a bad leader would help lead to the assassination of a good one later on. With regard to some Latin American nations, at any rate, one might argue that killing one dictator would only lead to a revolution and another dictator just as bad as the first, with the consequent assassination of the second one, thus leading to revolution and social chaos and a third dictator. There are countless empirical facts that must be taken into consideration and carefully weighed before any such decision can safely be made. The utilitarian is not committed to saying that any one policy or line of action is the best in any particular situation, for what is best depends on empirical facts which may be extremely difficult to ascertain. All he is committed to is the statement that when the action is one that does not promote human happiness as much as another action that he could have performed instead, then the action is wrong; and that when it does promote more happiness, it is right. Which particular action will maximize happiness more than any other, in a particular situation, can be determined only by empirical investigation. Thus, it is possible that killing is always wrong—at least the utilitarian could consistently say so and thus be a pacifist; but *if* killing is always wrong, it is wrong not because killing is wrong per se but because it always and without exception leads to worse consequences than any other actions that could have been performed instead. Then the pacifist, if he is a consistent utilitarian, would have to go on to show in each instance that each and every act of killing is worse (leads to worse consequences) than any act of refraining from doing so—even when the man is a trigger-happy gunman who will kill dozens of people in a crowded street if he is not killed first.



That killing is worse in every instance would be extremely difficult—most people would say impossible—to prove.

Consider the syllogism:

*The action which promotes the maximum happiness is right.*

*This action is the one which promotes the maximum happiness.*

Therefore, *This action is right.*

The utilitarian gives undeviating assent only to the *first* of these three statements (the major premise); this statement is the chief article of his utilitarian creed, and he cannot abandon it without being inconsistent with his own doctrine. But this first premise is not enough to yield the third statement, which is the conclusion of the argument. To know that the conclusion is true, even granting that the major premise is, one must also know whether the second statement (the minor premise) is true; and the second statement is an empirical one, which cannot be verified by the philosopher sitting in his study but only by a thorough investigation of the empirical facts of the situation. Many people would accept the major premise (and thus be utilitarians) and yet disagree among themselves on the conclusion because they would disagree on the minor premise. They would agree that an act is right if it leads to maximum happiness, but they would not agree on whether this action or that one is the one which *will* in fact lead to the most happiness. They disagree about the empirical facts of the case, not in their utilitarian ethics. The disagreement could be resolved if both parties had a complete grasp of all the relevant empirical facts, for then they would know *which* action *would* lead to the most happiness. In many situations, of course, such agreement will never be reached because the consequences of people's actions (especially when they affect thousands of other people over a long period of time, as happens when war is declared) are so numerous and so complex that nobody will ever know them all. Such a disagreement will not be the fault of ethics, or of philosophy in general, but of the empirical world for being so complicated and subtle in its workings that the full consequences of our actions often can not be determined. Frequently it would take an omniscient deity to know which action in a particular situation was right. Finite human beings have to be content with basing their actions on estimates of probability.

According to utilitarianism, then, the traditional moral rules are justified for the most part because following them will lead to the best consequences far more often than violating them will; and that is why they are useful rules of thumb in human action. But, for the utilitarian, this is *all* they are—rules of thumb. They should never be used blindly, as a pat formula or inviolable rule subject to no exceptions, without an eye to the detailed consequences in each particular situation. The judge who condemned a man to

death in the electric chair for stealing $1.95 (as in the case in Alabama in 1959) was probably not contributing to human happiness by inflicting this extreme penalty, even though he acted in accordance with the law of that state. The utilitarian would say that if a starving man steals a loaf of bread, as in Victor Hugo's *Les Miserables,* he should not be condemned for violating the rule "Do not steal"; in fact he probably did nothing morally wrong by stealing in this instance because the effects of not stealing would . . . have meant starvation and preserving a life (the utilitarian would say) is more important to human happiness than refraining from stealing a loaf of bread—especially since the man stole from one who was far from starving himself (the "victim" would never have missed it). He is probably blameless furthermore because the whole episode was made possible in the first place by a system of laws and a social structure which, by any utilitarian standard, were vicious in the extreme. (But see the effects of lawbreaking, below.)

Moral rules are especially useful when we have to act at once without being able adequately to weigh the consequences; for *usually* (as experience shows) better—i.e., more-happiness-producing—consequences are obtained by following moral rules than by not following them. If there is a drowning person whom you could rescue, you should do so without further investigation; for if you stopped to investigate his record, he would already have drowned. True, he might turn out to be a Hitler, but unless we have such evidence, we have to go by the probability that the world is better off for his being alive than his being dead. Again, there may be situations in which telling a lie will have better affects than telling the truth. But since, on the whole, lying has bad effects, we have to have special evidence that this situation is different before we are justified in violating the rule. If we have no time to gather such evidence, we should act on what is most probable, namely that telling a lie in this situation will produce consequences less good than telling the truth.

The utilitarian attitude toward moral rules is more favorable than might first appear because of the hidden, or subtle, or not frequently thought of, consequences of actions which at first sight would seem to justify a violation of the rules. One might consider *all* the consequences of the action and not just the immediate ones or the ones that happen to be the most conspicuous. For example: the utilitarian would not hold that it is *always* wrong to break a law, unless, he had good grounds for saying that breaking the law *always* leads to worse consequences than observing it. But if the law is a bad law to begin with or even if it is a good law on the whole but observing the law in this particular case would be deleterious to human happiness, then the law should be broken in this case. You would be morally justified, for example, in breaking the speed law in order to rush a



badly wounded person to a hospital. But in many situations (probably in most) in which the utilitarian criterion at first *seems* to justify the violation of a law, it does not really do so after careful consideration because of the far-flung consequences. For example, in a more typical instance of breaking the speed law, you might argue as follows: "It would make me happier if I were not arrested for the violation, and it wouldn't make the arresting officer any the less happy, in fact it would save him the trouble of writing out the ticket, so—why not? By letting me go, wouldn't the arresting officer be increasing the total happiness of the world by just a little bit, both his and mine, whereas by giving me a ticket he might actually decrease the world's happiness slightly?"

But happiness would be slightly increased only if one considers only the immediate situation. For one thing, by breaking the speed limit you are endangering the lives of others—you are less able to stop or to swerve out of the way in an emergency. Also those who see you speeding and escaping the penalty may decide to do the same thing themselves; even though you don't cause any accidents by your violation, *they* may do so after taking their cue from you. Moreover, lawbreaking may reduce respect for law itself; although there may well be unjust laws and many laws could be improved, it is usually better (has better consequences) to work for their repeal than to break them while they are still in effect. Every violation decreases the effectiveness of law, and we are surely better off having law than not having it at all—even the man who violently objects to a law and complains bitterly when he's arrested will invoke the law to protect *himself* against the violations of others. In spite of these cautions, utilitarianism does not say that one should *never* break a law but only that the consequences of doing so are far more often bad than good; a closer look at the consequences will show how true their reasoning is.

**Notes**

1. This term is most precisely defined by G. E. Moore in chapter 1, "Utilitarianism," of his book *Ethics.* [New York: Henry Holt and Company, 1912]. For the clearest and most rigorous statement of utilitarianism in its hedonistic form, see chapters 1 and 2 of [Moore's] book.
2. J. S. Mill, *Utilitarianism.* [ed. Oskar Piest (New York: Bobbs-Merrill, 1957; originally published 1863.], chap. 2.