

SIXTH EDITION

BIOMEDICAL ETHICS

UNIVERSITY OF WOLVERHAMPTON LEARNING & INFORMATION SERVICES		
ACC NO. 2428450	CLASS 999	
CONTROL NO. 0073976456		
DATE -7. JAN. 2008	SITE WL	

THOMAS A. MAPPE
Frostburg State University

DAVID DEGRAZIA
George Washington University



Boston Burr Ridge, IL Dubuque, IA Madison, WI New York San Francisco St. Louis
Bangkok Bogotá Caracas Lisbon London Madrid
Mexico City Milan New Delhi Seoul Singapore Sydney Taipei Toronto

McGraw-Hill Higher Education

A Division of The McGraw-Hill Companies



BIOMEDICAL ETHICS, SIXTH EDITION

Published by McGraw-Hill Higher Education, an imprint of The McGraw-Hill Companies, Inc., 1221 Avenue of the Americas, New York, NY 10020. Copyright © 2006, 2001, 1996, 1991, 1986, 1981 by The McGraw-Hill Companies, Inc. All rights reserved. No part of this publication may be reproduced or distributed in any form or by any means, or stored in a database or retrieval system, without the prior written consent of The McGraw-Hill Companies, Inc., including, but not limited to, in any network or other electronic storage or transmission, or broadcast for distance learning.

Some ancillaries, including electronic and print components, may not be available to customers outside the United States.

This book is printed on acid-free paper.

2 3 4 5 6 7 8 9 0 QPF/QPF 0 9 8 7 6

ISBN-13: 978-0-07-297644-1

ISBN-10: 0-07-297644-6

Editor-in-chief: *Emily J. Barrasse*

Publisher: *Lyn Uhl*

Senior sponsoring editor: *Jon-David Hague*

Editorial coordinator: *Allison Rona*

Marketing manager: *Zina Craft*

Project manager: *Judy Ludowitz/Carlisle Publishers Services*

Production supervisor: *Tandra Jorgensen*

Design manager: *Preston Thomas*

This book was set in 10/12 Times Roman by Carlisle Communications and printed on 45# New Era Matte by Quebecor World, Fairfield.

Library of Congress Cataloging-in-Publication Data

Biomedical ethics / [edited by] Thomas A. Mappes, David DeGrazia.

-- 6th ed.

p. ; cm.

Includes bibliographical references.

ISBN 0-07-297644-6

1. Medical ethics. 2. Bioethics. I. Mappes, Thomas A.

II. DeGrazia, David.

[DNLM: 1. Bioethics. 2. Ethics, Medical. W 50 B6153 2005]

R724.B49 2005

174.2--dc22

2005017337

not only from the vantage point of moral philosophy (the principal vantage point in the collection of readings in this text) but also from the vantage point of moral theology. Whereas philosophical arguments are constructed without presupposing the truth of any religious claims, that is, without reliance on religious *faith*, theological arguments are generally constructed within a faith framework. There is yet a third—and most significant—way in which biomedical ethics is interdisciplinary, and that is by reference to the disciplines of medicine and biology. Medical judgments and the findings of biology often play a crucial role in ethical deliberations. (The findings of the social sciences can be relevant as well.) It is also important to recognize that the *experience* of health-care professionals and biomedical researchers is often essential to ensure that ethical discussions retain firm contact with the concrete realities that permeate the practice of medicine and the pursuit of biomedical research.

Although the issues of biomedical ethics are essentially normative, they are intertwined with both conceptual issues and factual (i.e., empirical) issues. For example, suppose we are concerned with the ethical acceptability of intervention for the sake of preventing a person from committing suicide. Our basic concern is with a normative question; however, we must face the problem of clarifying the nature of suicide, a conceptual issue. For example, if a Jehovah's Witness, on the basis of religious principle, refuses a life-saving blood transfusion, is the resultant death to be classified as a suicide? In addition to facing conceptual perplexities, we are also faced with an important factual question: Do those who typically attempt suicide really want to die? Presumably psychologists have important things to tell us on this score. In the end, of course, we want to reach an ethical conclusion. However, ethical deliberations must proceed in the light of conceptual structures and factual beliefs. In the case of some issues in biomedical ethics, underlying factual issues are especially prominent. For example, in addressing the normative question of whether it is ever morally permissible to use children as research subjects, it is important to consider a factual question: To what extent can therapeutic techniques be developed for children in the absence of research employing children as research subjects? In the case of other issues in biomedical ethics, associated conceptual issues command special attention. For example, one could hardly discuss the normative issue of whether it is appropriate to transplant vital organs from brain-dead patients without closely examining the concept of death.

It is helpful to approach the literature of biomedical ethics with an eye toward distinguishing conceptual, factual, and normative issues. Furthermore, with regard to normative issues, which are the central issues of biomedical ethics, one cannot hope to situate argumentation in biomedical ethics properly without some awareness of the various types of ethical theory developed in general normative ethics. Such theories provide the frameworks within which many of the arguments in biomedical ethics are formulated.

RECENTLY DOMINANT ETHICAL THEORIES

An ethical theory provides a framework that can be used to determine what is morally right and morally wrong regarding human action in general, or what is morally good and morally bad regarding human character in general. The discussion of ethical theories in this section is restricted in two ways. First, consideration is limited to theories of right and wrong action, as opposed to theories of good and bad character (which fall naturally under the head-

ing of virtue ethics, an approach that is explicated in a later section of this chapter). Second, consideration is limited to those theories of right and wrong action that commanded the most attention in the twentieth century. These recently dominant theories are frequently reflected in arguments advanced in biomedical ethics.

An ethical theory—as discussed in this section—provides an ordered set of moral standards (in some cases, simply one ultimate moral principle) that is to be used in assessing what is morally right and what is morally wrong regarding human action in general. A proponent of any such theory puts it forth as a framework with which a person can correctly determine, on any given occasion, what he or she (morally) ought to do.

THE CRITICAL ASSESSMENT OF COMPETING ETHICAL THEORIES

Since a number of competing ethical theories may be identified, the question that immediately arises is what criteria are relevant to an assessment of these competing theories. There is no easy way to answer this very fundamental and very controversial question, but let us start with those considerations whose relevance is unlikely to be disputed. Any theory in any field is rightly expected to be internally consistent. Thus, a theory can be faulted on the basis of inconsistency. In a similar vein, any theory is surely flawed to the extent that it is either unclear or incomplete. In addition, a theory should be as simple as it can be without entailing a failure to satisfy other relevant criteria, such as clarity and completeness.

If the above considerations are relevant to a critical assessment of theories in any field, we must yet identify considerations relevant to our particular concern, the critical assessment of (normative) ethical theories. Responsive to this task, it is suggested that the following criteria embody the two most important considerations. (1) The implications of an ethical theory must be largely reconcilable with our experience of the moral life. (2) An ethical theory must provide effective guidance where it is most needed, that is, in those situations where substantial moral considerations can be advanced on both sides of an issue. In embracing the priority of criteria 1 and 2 we are saying that an adequate ethical theory must achieve two major goals. An adequate ethical theory must accord with the moral life as we experience it, and it must function heuristically by guiding us when we are confronted with moral perplexity. An ethical theory should, on the one hand, make sense out of the moral life by exhibiting the basic features of our ordinary moral thinking. On the other hand, it should illuminate our moral judgment precisely where it is experienced to falter—in the face of moral dilemmas.

There is certainly no suggestion here that the standards embodied in criteria 1 and 2 can be applied in some mechanical fashion to assess the relative adequacy of a proposed ethical theory. Intellectual judgments on these matters are necessarily complex and subtle. In saying, for example, that an adequate ethical theory must accord with our experience of the moral life, we certainly do not want to insist that each and every divergence from the verdict of “commonsense morality” must be interpreted as counting against an ethical theory. Perhaps we would be better advised to revise our moral judgment in light of the theory. (In empirical science, fact-theory mismatches are sometimes resolved not by modifying the theory but by reinterpreting the facts in the light of the theory.) In embracing criterion 1 we undoubtedly commit ourselves to a point of view incompatible with the acceptance of an ethical theory that is revisionary in some wholesale sense, but we do not commit ourselves to the view that “commonsense morality” is sacrosanct. If an ethical theory successfully captures the basic features of our ordinary moral thinking, it will, of

course, be true that its implications in large measure accord with our ordinary moral thinking. If the theory, however, cannot be reconciled with a relatively smaller range of our ordinary moral judgments, we may decide to interpret this disharmony as the product of some inadequacy in "commonsense morality" rather than as an inadequacy in the proposed theory.

TELEOLOGICAL VERSUS DEONTOLOGICAL THEORIES

With the introduction of criteria 1 and 2, we are now prepared to undertake a survey of alternative ethical theories. Our immediate concern is the identification, articulation, and critical consideration of those ethical theories that are the most prominent theories in general normative ethics—commanding the most attention in the twentieth century—and frequently reflected in argumentation advanced in biomedical ethics. In a later section, under the heading of "Alternative Directions and Methods," some additional theoretical perspectives that are important in biomedical ethics are presented.

In contemporary discussions, ethical theories are often grouped into two basic, and mutually exclusive, classes—*teleological* and *deontological*. Any ethical theory that claims the rightness and wrongness of human action is *exclusively* a function of the goodness and badness of the consequences resulting directly or indirectly from that action is a teleological theory. Consequences are all-important here. A deontological theory maintains, in contrast, that the rightness and wrongness of human action is *not exclusively* (in the extreme case, not at all) a function of the goodness and badness of consequences. Accordingly, a theory is deontological (rather than teleological) if it places limits on the relevance of teleological considerations. Thus, an ethical theory in which the moral rightness and wrongness of human action is construed as totally independent of the goodness and badness of consequences would be only one kind, albeit the strongest or most extreme kind, of deontological theory.

The most prominent teleological ethical theory is the theory known as "utilitarianism." The adequacy of utilitarianism and the issue of its proper explication continue to be significant concerns in contemporary discussions of ethical theory. For this reason, and especially because much argumentation in biomedical ethics is based on utilitarian reasoning, utilitarianism warrants our detailed attention. However, it should first be noted that utilitarianism is not the only ethical theory that is correctly categorized as teleological. One other notable teleological theory is the theory known as "ethical egoism." The basic principle of ethical egoism can be phrased as follows: *A person ought to act so as to promote his or her own self-interest.* An action is morally right if, when compared with possible alternatives, its consequences are such as to generate the greatest balance of good over evil *for the agent.* (The impact of action on other people is irrelevant except as it may indirectly affect the agent.) Ethical egoism is a teleological theory precisely because, by the terms of the theory, the rightness and wrongness of human action is exclusively a function of the goodness and badness of consequences.

Ethical egoism is an enormously problematic theory, one whose implications seem to be intensely at odds with our ordinary moral thinking. Under certain conditions, ethical egoism leads us to the conclusion that it is a person's moral obligation to perform an action that is flagrantly antisocial in nature. Consider this example. Mr. A loves to set buildings on fire; nothing makes him happier than watching a building burn. He recognizes that arson destroys property and subjects human life to serious risk, but he happens to be a thoroughly unsympathetic person, one whose well-being is not negatively affected by the misfortune

of others. Of course, it is not in A's self-interest (and thus would not be A's moral obligation) to burn down a building if there is a good chance that he will be caught. (The punishment for arson is severe.) However, if A is very clever and if it is virtually certain that he will not be caught, ethical egoism seems to imply that arson is the morally right thing for him to do.

Another problematic feature of ethical egoism is that it cannot be publicly advocated without inconsistency. Suppose that Ms. B embraces ethical egoism. Accordingly, she considers it her moral obligation always to act in such a way as to promote her individual self-interest. Should she now publicly advocate ethical egoism, that is, encourage others to adopt the view that each person's moral obligation is to act in such a way as to promote his or her individual self-interest? No. Since it is to *her* advantage that others *not* act egoistically, it follows that it would be immoral for her to publicly advocate ethical egoism.

In reducing morality to considerations of personal prudence, it can be argued, ethical egoism destroys the very sense behind morality. Morality, it would seem, functions (at least in part) to restrict the pursuit of personal self-interest. It is not that morality prohibits the pursuit of personal self-interest; rather, it places limits on this pursuit. In "collapsing" morality into prudence, ethical egoism does not accord with a commonly experienced phenomenon of the moral life, the tension between self-interest and morality, between "what would be best for me" and "what is the morally right thing."

In fairness to ethical egoism, it must be noted that its proponents have sometimes devised ingenious arguments in an attempt to minimize the sort of difficulties just discussed. However, ethical egoism is not widely defended in contemporary discussions of ethical theory, and it surely plays an insignificant role in discussions of biomedical ethics. It has been introduced primarily as a notable instance of a teleological yet nonutilitarian theory. Attention will now be focused on utilitarianism.

In its classical formulation, utilitarianism is found most prominently in the works of two English philosophers, Jeremy Bentham (1748–1832) and John Stuart Mill (1806–1873). In contemporary discussions, a distinction is made between two kinds of utilitarianism—*act-utilitarianism* and *rule-utilitarianism*. Although it is somewhat controversial whether a significant distinction can be maintained between these two versions of utilitarianism, it is presumed for the sake of exposition that two distinct utilitarian ethical theories can indeed be articulated.³

ACT-UTILITARIANISM

Human action typically takes place within the fabric of our social existence. Thus, an action performed by one person often affects not only the agent but also the lives of many others. Consider a man who refuses to stop smoking even though he suffers from emphysema. He will not be the only one to suffer the consequences; certainly those who care about him will also. His refusal to give up smoking, since it has the effect of further damaging his health, also produces a higher level of anxiety among the members of his family. Among the other detrimental consequences of his continuing to smoke is the negative impact on any nonsmokers in the vicinity when he smokes: annoyance, displeasure, and the like. However, the various consequences of a single action are seldom uniformly good or uniformly bad. In addition to the bad consequences already indicated, there are also a number of good consequences that result from the refusal to stop smoking. Most notably, the emphysema patient continues to derive the satisfaction associated with cigarette smoking. In

addition, it is likely that his continuing to smoke will make him less irritable around others. When the various consequences of a single action are fully analyzed, more often than not we find ourselves confronted with a mixture of good and bad. For example, if a person throws a late-night party, it is true that those in attendance may have a very good time, but it is also true that the neighbors may lose out on some much needed sleep.

The basic principle of act-utilitarianism can be stated as follows: *A person ought to act so as to produce the greatest balance of good over evil, everyone considered.* Act-utilitarianism stands in vivid contrast to ethical egoism, which directs a person always to act so as to produce the greatest balance of good over evil *for oneself* (i.e., the agent). The act-utilitarian is committed to the proposition that the interests of everyone affected by an action are to be weighed in the balance along with the interests of the agent. Everyone's interests are entitled to an impartial consideration. According to the act-utilitarian, an action is morally right if, when compared with possible alternatives, its likely consequences are such as to generate the greatest balance of good over evil, everyone considered. If we refer to the net balance of good over evil (everyone considered) that is likely to be produced by a certain action as its (overall) *utility*, then we can say that act-utilitarianism directs a person always to choose that alternative that has the greatest utility. Thus, we can express the basic principle of act-utilitarianism as follows: A person ought to act so as to maximize utility.

For the act-utilitarian, calculation is a paramount element in the moral assessment of action. The question is always this: What is the utility of each of my alternatives in this particular set of circumstances? However, any system of utilitarian calculation must ultimately be anchored in some conception of intrinsic value (i.e., that which is good or desirable in and of itself). The act that will maximize utility (by our definition) is the act that is likely to produce the greatest balance of good over evil, everyone considered. However, what is to count as "good" and what as "evil" in our calculations? The answers provided within the framework of classical utilitarianism reflect a so-called hedonistic theory of intrinsic value. According to Bentham, only pleasure (understood broadly to include any type of satisfaction or enjoyment) has intrinsic value; only pain (understood broadly to include any dissatisfaction, frustration, or displeasure) has intrinsic disvalue. According to Mill, only happiness has intrinsic value; only unhappiness has intrinsic disvalue. To what extent there is substantive disagreement between Bentham and Mill on this matter is a complex question that cannot be dealt with here. It should be mentioned, however, that many contemporary utilitarian thinkers have embraced more elaborate and nonhedonistic theories of intrinsic value.⁴ Nevertheless, for the sake of exposition, we shall presume that a hedonistic theory of intrinsic value, in the spirit of Bentham and Mill, underlies utilitarian calculation.

In the spirit of act-utilitarianism, in order to determine what I should do in a certain situation, I must first attempt to delineate alternative paths of action. Next, I attempt to foresee the consequences (sometimes numerous and far-reaching) of each alternative action. Then I attempt, in each case, to evaluate the consequences and to weigh the good against the bad, considering the impact of my action on everyone whom it is likely to affect. Such a reckoning will reveal the act that is likely to produce the greatest balance of good over evil, and this act is the morally right act for me in my particular circumstances. (If it appears likely that two competing actions would produce the same balance of good over evil, then either action will qualify as morally correct.) In some situations, it is true that no matter what I do, more evil (pain or unhappiness) will come into the world than good (pleasure or happiness). In such unfortunate situations, according to the act-utilitarian, the morally right act is the one that will bring the least unfavorable balance of evil and good into the world.

Act-utilitarianism can rightly be understood as a form of "situation ethics." The act-utilitarian has no sympathy for the notion that certain kinds of actions are intrinsically wrong, that is, wrong by their very nature. Rather, a certain kind of action (e.g., lying) may be wrong in one set of circumstances yet right in another. The circumstances in which an action is to be performed are relevant to its morality (i.e., its rightness or wrongness) because the consequences of the action will vary depending upon the circumstances. Thus, the morality of action is a function of the situation confronting the agent—"situation ethics."

The situational character of act-utilitarianism is reflected in the act-utilitarian attitude toward moral rules. Among the "commonsense rules of morality" are the following: "do not kill," "do not injure," "do not steal," "do not lie," "do not break promises." According to the act-utilitarian, these rules are to be understood merely as rules of thumb. They are, for the most part, reliable guides for human action, especially relevant when time constraints undermine the possibility of careful calculation. In most circumstances, acting in accordance with a moral rule is the way to maximize utility, but in some cases this is not so. In these latter cases, whenever there is good reason to believe that breaking a moral rule will produce a greater balance of good over evil (everyone considered), the right thing to do is to break it. In such a case, it would be wrong to follow the rule. Lying is usually wrong, breaking promises is usually wrong, killing is usually wrong; however, whenever circumstances are such that there is good reason to believe that breaking a certain moral rule will maximize utility, the rule should be broken. Of course, the act-utilitarian insists, one must be cautious in concluding that any given exception to a moral rule is indeed justified. One must be wary of rationalization and not allow one's own interests to weigh more heavily than the interests of others in the utilitarian calculation. Most importantly, one must not be simpleminded in a consideration of the likely consequences of breaking a moral rule. Indirect and long-term consequences must be considered as well as direct and short-term consequences. Lying on a certain occasion may seem to promote most effectively the interests of those immediately involved, but perhaps the lie will provide a bad example for less reflective people, or perhaps it will contribute to a general breakdown of trust among human beings. In this same vein, one prominent contemporary act-utilitarian emphasizes the significance of the long-term, indirect consequences of promise breaking, while at the same time exhibiting the underlying act-utilitarian attitude toward moral rules:

The rightness or wrongness of keeping a promise on a particular occasion depends only on the goodness or badness of the consequences of keeping or of breaking the promise on that particular occasion. Of course part of the consequences of breaking the promise, and a part to which we will normally ascribe decisive importance, will be the weakening of faith in the institution of promising. However, if the goodness of the consequences of breaking the rule is *in toto* greater than the goodness of the consequences of keeping it, then we must break the rule. . . .³

Act-utilitarianism has often been criticized on the grounds that, due to the extensive sort of calculations it seems to demand, it cannot function as a useful guide for human action. In the spirit of this criticism, the following questions are asked: How can I possibly predict all the consequences of my actions? How am I to assign weights to the various kinds of human satisfactions—for example, the pleasure of eating a candy bar versus the aesthetic enjoyment of the ballet? How am I to weigh the anxiety of one person against the inconvenience of another? Besides, how am I supposed to have time to do these extensive calculations? Act-utilitarians, in response to such questions, usually appeal rather directly to

"common sense." They say, typically, that there is no escape from a consideration of probabilities in rational decision making; predict as best you can and weigh as best you can, considering the time you have available for deliberation. All that can be expected is that you come to grips with the likely consequences of your alternatives in a serious-minded, sensible way and then act accordingly.

Examples of Act-Utilitarian Reasoning in a Biomedical Context The following examples are provided in an effort to exhibit act-utilitarian reasoning as it might arise in a biomedical context. It is not claimed that an act-utilitarian must necessarily reach the conclusion suggested in each case. It is claimed only that an act-utilitarian might plausibly reach the stated conclusion.

(1) A severely impaired newborn, believed to have no realistic chance of surviving more than a few weeks, has contracted pneumonia. (The treatment of impaired newborns is discussed in Chapter 5.) A physician, in conjunction with the parents of the infant, must decide whether to fight off the pneumonia with antibiotics, thereby prolonging the life of the infant. The alternative is simply to allow the infant to die. It seems clear that the interests of all those immediately involved are best served by deciding not to treat the pneumonia. Surely the infant has nothing to gain, and something to lose, by a slight extension of a pain-filled life. The parents, whose suffering cannot be eradicated whatever action is taken, nevertheless will find some relief knowing that their child's suffering has ended. In addition, hospital resources can be better utilized than to prolong the dying process of an infant who cannot benefit from further treatment. However, there may be decisive consequences of allowing death to occur that are indirect and long-term. Perhaps allowing this infant to die will contribute to a breakdown of protective attitudes toward infants in general. No, the risk of this untoward consequence seems minimal. Withholding antibiotics, thereby allowing the infant to die, is the right thing to do in this case.

(2) A biomedical researcher, on the basis of animal studies she has conducted, believes that a certain drug therapy has great promise for the treatment of a particular kind of cancer in human beings. (The use of animals in research is discussed in Chapter 4.) At present, however, her primary concern is to establish an appropriate dosage level for human beings: there have been several troublesome side effects exhibited by the animals who received large doses of the drug. Over the years, the researcher has found that students at her university are very willing to volunteer as research subjects in experiments that can be identified as presenting only minimal risks to themselves. They are, however, understandably reluctant to volunteer for experiments that seem to present more substantial risks. The researcher in this case cannot honestly say that there are no substantial risks for research subjects. She expects, in particular, that perhaps 30 to 40 percent of the research subjects will have to contend with very prolonged nausea. However, if she is honest in conveying this information to potential research subjects, it is unlikely that they will volunteer in sufficient numbers. (The ethics of experimentation on human subjects is discussed in Chapter 4.) Perhaps, she reasons, it is justifiable in this case to withhold information about the risk of very prolonged nausea. After all, it is very likely that numerous people will eventually derive great benefit from the therapeutic technique under study. Surely this likely benefit far outweighs the short-term discomfort of a much smaller number. But consider the very real possibility that the deception would come to light. If those who routinely volunteer as research subjects are given a reason to distrust those conducting the experiments, the overall research effort on campus will be negatively affected. Moreover, publicity about the decep-

tion would create a major public relations scandal for the university, forcing it to devote valuable time, energy, and money to repairing its reputation. These seem to be decisive considerations. In this case, then, deception would be wrong. (If there were no realistic chance of the deception's being discovered, it seems that the conclusion would be different.)

(3) The setting is in the 1960s, when kidney dialysis machines are scarce, and it is not possible for all who need them to be given access. A hospital administrator or perhaps a committee has been charged with the responsibility of deciding, in essence, whose lives will be saved. (Such decisions are often referred to as "microallocation decisions.") On a particular occasion, when there is room for one more patient, there are two candidates in great need. One of the candidates, a civic-minded woman of 40, is married and the mother of four children. The other candidate, an unmarried man of the same age, is known to be a drifter and an alcoholic. It seems clear, at first glance, that the consequences of saving the woman's life are far superior to those of saving the man's life. Her husband, her children, and the community in general would be negatively affected in very substantial ways by her death. However, is it not problematic to accord a person access to a scarce medical resource on the basis of his or her social role? If a precedent of this sort is set, will not those whose lives are less "socially useful" become somewhat anxious and fearful? On the other hand, perhaps this negative consequence will be balanced by a positive consequence; that is, people will be more inclined to become "socially useful." It still seems clear that the woman in this case should have priority over the man.

Critical Assessment of Act-Utilitarianism Act-utilitarianism arguably fares poorly when measured against a previously identified standard: The implications of an ethical theory must be largely reconcilable with our experience of the moral life. In a number of ways, it can be argued, act-utilitarianism clashes with our experience of the moral life. This perceived failure to accord with our ordinary moral thinking is reflected in the following well-known objections to act-utilitarianism.

(1) Act-Utilitarianism Confronts Individuals with an Overly Demanding Moral Standard. We are accustomed to thinking that at least some of our decisions are matters of "mere prudence," rightly decided on the basis of "what is best for me." Which major a college student should choose is a good example of a choice that we are inclined to consider essentially a nonmoral matter, a matter of "mere prudence." According to the act-utilitarian, however, a person is continually under a moral obligation to produce the greatest balance of good over evil, everyone considered. Whereas ethical egoism seems to wrongly "collapse" morality into prudence, it would seem that act-utilitarianism "expands" morality so as to destroy the realm of prudence. No aspect of a person's life can be considered merely a matter of prudence. Every decision is a moral decision, to be made on the basis of utilitarian calculation. However, no matter how noble it might be for a college student to decide his or her major on the basis of a utilitarian calculation, it would seem that one is certainly not under an obligation to proceed in this manner. Doing so, we would ordinarily say, is not one's duty but, rather, is something "above and beyond the call of duty." Act-utilitarianism, in directing a person always to act so as to maximize utility, seems problematically to imply that it is one's duty to act in a way that we ordinarily consider "above and beyond the call of duty."

(2) Act-Utilitarianism Does Not Accord with Our Experience of Particular, Morally Significant Relationships. In our experience of the moral life, we are continually aware of highly particular, morally significant relationships that exist between ourselves and others.

We are related to particular individuals in a host of morally significant ways, such as spouse to spouse, parent to child, creditor to debtor, promisor to promisee, employer to employee, teacher to student, physician or nurse to patient. In view of such relationships, it is ordinarily thought, we have special obligations—obligations that restrict the effort to maximize utility. Parents, we are strongly inclined to say, are obligated to care for their children even if there is good reason to think that the time and energy necessary for this task would maximize utility if redirected to some other task. In the same way, by virtue of the special relationship that exists between a physician and a patient, would it not be wrong for a physician to make decisions regarding a patient's treatment in the manner of an act-utilitarian? For a physician to compromise the interests of an individual patient in an effort to maximize utility surely seems wrong. W. D. Ross, who has vigorously pressed this overall line of criticism against act-utilitarianism, asserts that the "essential defect of the . . . theory is that it ignores, or at least does not do full justice to, the highly personal character of duty."⁶

(3) *Act-Utilitarianism Does Not Accord with Our Conviction That Individuals Have Rights.* The notion of rights plays an important part in our ordinary moral thinking, but act-utilitarianism seems incapable of accommodating this notion. Moreover, in certain circumstances, the action that would maximize utility (and thus the right action, according to the act-utilitarian) is one that we are inclined to consider seriously immoral precisely because it entails the violation of some person's right. For example, it seems that act-utilitarianism would allow an innocent person to be unjustly punished, as long as circumstances were such as to make this line of action the one that would generate the greatest balance of good over evil. Suppose extreme social unrest has been created by a wave of unsolved crimes. The enraged crowd will violently erupt, bringing massive evil into the world, unless the authorities punish someone (anyone) in an effort to appease the appetite for vengeance. So act-utilitarianism seems to allow the unjust treatment of a person as a scapegoat, as a mere means to a social end. But surely an innocent person has a right not to be punished, and it is by reference to this right that the wrongness of scapegoating is most naturally understood. Similarly, "the common moral opinion that painless undetected murders of old unhappy people are wicked, no matter what benefits result"⁷ can be thought to rest on the contention that people, however old and unhappy, nevertheless have a *right* to life. It is often asserted against act-utilitarianism that it is a defective theory because it allows "the end to justify the means." At least part of the sense behind this charge can be made out in reference to the notion of rights. Certain means of achieving a desirable social end are simply wrong because they entail the violation of a person's right. Contrary to act-utilitarianism, such means cannot be justified by the end.

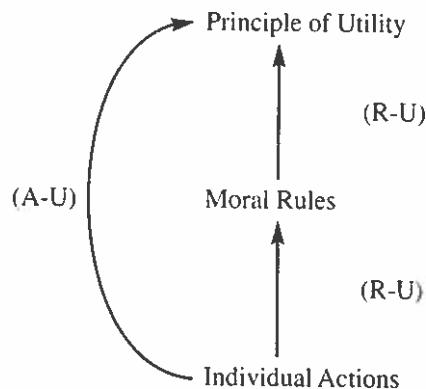
Act-utilitarians have responded in two ways to the overall claim that the theory cannot be reconciled with our ordinary moral thinking. Some say, in essence, "so much the worse for our ordinary moral thinking." In their view, we must simply overhaul our collective moral consciousness and embrace the mind-set of the act-utilitarian. Most act-utilitarians, however, do not adopt this revisionary stance. Rather, they seek to demonstrate that the clash between act-utilitarianism and our ordinary moral thinking is not nearly so severe as the above criticisms suggest. They argue that, when act-utilitarianism is properly applied, when all the significant long-term, indirect consequences are taken into account, the theory does not give rise to conclusions that seem so patently objectionable. It is very doubtful, however, that this strategy of argument can completely rescue act-utilitarianism from its difficulties.

Perhaps act-utilitarianism fares better when measured against the second of our previously identified standards: An ethical theory must provide effective guidance where it is most needed. At the very least, it must be said in favor of act-utilitarianism that it provides a reasonably clear decision procedure, a sense of direction, for the resolution of moral dilemmas. In the face of moral considerations that incline our judgment in conflicting ways, act-utilitarianism counsels us to analyze the likely consequences of alternative actions in order to determine the alternative that will maximize utility. Still, however well act-utilitarianism might be thought to fare with regard to our second standard—and even that is debatable—it seems to encounter significant problems when measured against our first standard. Indeed, in contemporary times, most utilitarian thinkers have rejected act-utilitarianism in favor of a theory known as rule-utilitarianism.⁸

RULE-UTILITARIANISM

The basic principle of act-utilitarianism has previously been formulated as follows: A person ought to act so as to produce the greatest balance of good over evil, everyone considered. In contrast, the basic principle of rule-utilitarianism can be formulated as follows: *A person ought to act in accordance with the rule that, if generally followed, would produce the greatest balance of good over evil, everyone considered.* If the demand to produce the greatest balance of good over evil, everyone considered, is referred to as the principle (standard) of utility, then the principle of utility is the basic ethical principle in both the act-utilitarian and the rule-utilitarian systems. However, in the act-utilitarian system, determining the morally correct action is a matter of assessing alternative actions directly against the standard of utility, whereas in the rule-utilitarian system determining the morally correct action involves an *indirect* appeal to the principle of utility. In the spirit of rule-utilitarianism, a moral code is first established by reference to the principle of utility. That is, a set of valid moral rules is established by determining which rules (as opposed to conceivable alternatives), if generally followed, would produce the greatest balance of good over evil. In rule-utilitarianism, individual actions are morally right if they are in accord with those rules.

The difference between act-utilitarian reasoning and rule-utilitarian reasoning can be represented schematically as follows:



Act-utilitarian reasoning embodies a single-stage procedure; rule-utilitarian reasoning, a two-stage procedure. Because the act-utilitarian is committed to assessing individual actions strictly on the basis of utilitarian considerations, act-utilitarianism is often referred to as "extreme" or "unrestricted" utilitarianism. Because the rule-utilitarian is committed to developing a moral code (a set of moral rules) on the basis of utilitarian considerations and then assessing individual actions, not on the basis of utilitarian considerations but on the basis of accordance with the moral rules that have been established, rule-utilitarianism is often referred to as "restricted" utilitarianism.

For the act-utilitarian, moral rules have a very subordinate status. They are merely "rules of thumb," providing some measure of practical guidance. For the rule-utilitarian, moral rules assume a much more fundamental status, indeed a theoretical primacy. Only in reference to established moral rules can the moral assessment of individual actions be carried out. Thus, the first and most crucial step for the rule-utilitarian is the articulation of a set of moral rules, themselves justified on the basis of utilitarian considerations. Underlying this task is the question of which rules (as opposed to conceivable alternatives), if generally followed, would produce the greatest balance of good over evil, everyone considered. That is, which rules, if adopted or recognized in our moral code, would maximize utility?

As a first approximation of a set of moral rules that could be justified on the basis of utilitarian considerations, consider the "commonsense rules of morality," such as "do not kill," "do not steal," "do not lie," "do not break promises." It is not difficult to think of such rules as resting upon a utilitarian foundation. Surely the consequences of the adoption of the rule "do not kill" are dramatically better than the consequences of the adoption of the rule "kill whenever you want." If the latter rule were generally followed, society would give way to anarchy. Similarly, the consequences of the adoption of the rule "do not steal" are dramatically better than the consequences of the adoption of the rule "steal whenever you want." If the former rule is generally followed, individuals will enjoy an important measure of personal security. If the latter rule were adopted by a society, anxiety and tension would dominate social existence. As for lying and promise breaking, if people felt free to engage in such behavior, the numerous advantages that derive from human trust and cooperation would evaporate. However, the rules thus far exhibited as having a utilitarian foundation are essentially prohibitions. Are there not also rules of a more positive sort that could also be justified on the basis of utilitarian considerations? It would seem so. Consider rules such as "come to the aid of people in distress" and "prevent innocent people from being harmed." It surely seems that human welfare would be enhanced by the adoption of such rules as part of the overall fabric of our moral code.

According to the rule-utilitarian, an individual action is morally right when it accords with the rules or moral code established on a utilitarian basis. However, the account of moral rules thus far presented is too simplistic. In order to be plausible, the rules that constitute the moral code must be understood as incorporating certain exceptions. The need to recognize justified exceptions is perhaps most apparent when we remember that moral rules, if stated unconditionally, can easily come into conflict with each other. When an obviously agitated person waves a gun and inquires as to the whereabouts of a third party, it may not be possible to act in accordance with both the rule "do not lie" and the rule "prevent innocent people from being harmed." Indeed, it is precisely this sort of situation that inclines us to consider incorporating an exception into our rule against lying. Suppose we say, "Do not lie *except* when necessary to prevent an innocent person from being seriously harmed." When the possibility of a justified exception is raised, the rule-utilitarian employs

the following decision procedure. The question is posed, "Would the adoption of the rule with the exception have better consequences than the adoption of the rule without the exception?" If so, the exception is a justified one; the rule incorporating the exception has greater utility than the rule without the exception. In the face of our proposed exception to the rule against lying, the rule-utilitarian would probably conclude that it does constitute a justified exception. The adoption of the rule "do not lie *except* when necessary to protect an innocent person from being seriously harmed" would seem to preserve essentially all the social benefits provided by the adoption of the rule "do not lie," while bringing about an additional social benefit, an increased measure of personal security for potential victims of assault.

Examples of Rule-Utilitarian Reasoning in a Biomedical Context (1) A substantial problem in biomedical ethics (discussed in Chapter 2) is whether it is ever right for a physician to lie to a patient, saying that the patient's illness is not terminal when it is believed to be so. The rule-utilitarian would conceptualize this issue as raising the possibility of a justified exception to the rule against lying. (Notice that an act-utilitarian, in contrast, would insist on assessing every individual case on its own utilitarian merits.) Suppose we consider incorporating into the rule against lying an exception to this effect: "*except* when in the judgment of a physician it would be better for a patient not to know that his or her illness is believed to be terminal." Would the adoption of the rule incorporating this exception have better consequences than the adoption of the rule without the exception? The correct answer to this question is perhaps arguable, but it would seem that the rule-utilitarian would conclude that the proposed exception is an unjustified one. It is perhaps true that adoption of a rule incorporating the proposed exception would result in many patients' being spared (at least temporarily) the distress that accompanies knowledge of one's impending death. On the other hand, it seems that this gain would be dwarfed by the distress and anxiety that would emerge from the erosion of trust within the confines of the physician-patient relationship. Whether a more limited exception could be formulated to a rule-utilitarian's satisfaction remains an open question.

(2) Another substantial problem in biomedical ethics (discussed in Chapter 6) has to do with the morality of mercy killing. Suppose a terminally ill patient, in great pain, requests that a physician terminate his or her life by administering a lethal dose of a drug. Such a case raises the issue of voluntary (active) euthanasia. The rule-utilitarian would conceptualize this issue (and other issues, such as suicide and abortion) as raising the possibility of a justified exception to our rule against killing. Notice that at least one exception to our rule against killing is relatively uncontroversial. Killing in self-defense is justifiable, according to the rule-utilitarian, because although the adoption of the rule "do not kill" has dramatically better consequences than the adoption of the rule "kill whenever you want," adoption of the rule "do not kill *except* in self-defense" has still better consequences. As for voluntary (active) euthanasia, perhaps we should say that strong rule-utilitarian arguments can be advanced on both sides of the issue. Rule-utilitarian proponents of voluntary (active) euthanasia emphasize that social acceptance of this practice would result in great benefits—the primary one being that many dying people would be able to escape an extension of an anguished dying process. On the other side of the issue, however, we find, among a number of important concerns, insistence that availability of the lethal dose would create a climate of fear and anxiety among the elderly. Will dying people not come to feel that their families, to whom they have become a burden, expect them to ask for the lethal dose?

(3) A final illustration of rule-utilitarian reasoning in a biomedical context can be presented in reference to the principle of medical confidentiality (discussed in Chapter 3). This principle, which has an obvious basis in a rule-utilitarian structure, demands that information revealed within the context of a therapeutic relationship be held confidential. If patients could not rely on this expectation, they would be reluctant to communicate information that is essential to their proper treatment. Still, are there not justifiable exceptions to the principle of medical confidentiality? Suppose, for example, a patient reveals to his or her therapist an intention to kill or injure a third party. Is it not incumbent upon the therapist to break medical confidentiality in an effort to ensure protection for the third party? The situation just described is the basis of the *Tarasoff* case considered in Chapter 3, and rule-utilitarian arguments on both sides of the issue can be found in the judicial opinions presented. There is an obvious benefit associated with the recognition of an exception to medical confidentiality based on the interests of innocent third parties; namely, threatened people will sometimes be saved from injury and death. On the other hand, it is argued, emotionally disturbed patients are likely to become more inhibited in communicating with therapists; thus, their treatment will be inhibited, and a greater incidence of violence against innocent people will result.

Critical Assessment of Rule-Utilitarianism Rule-utilitarianism, it would seem, goes some distance toward alleviating the perceived difficulties of act-utilitarianism. Although act-utilitarians have charged rule-utilitarians with "superstitious rule-worship,"⁹ it is act-utilitarianism rather than rule-utilitarianism that seems to clash with our ordinary moral thinking on this score. Rule-utilitarianism seems to fare at least somewhat better than act-utilitarianism when measured against the standard that the implications of an ethical theory must be largely reconcilable with our experience of the moral life.

Whereas act-utilitarianism seems to confront individuals with an overly demanding moral standard, placing each of us under a continuing obligation to maximize utility with each of our actions, rule-utilitarianism may prove to be far less demanding of individuals. It requires only that individuals conform their actions to the rules that constitute a utilitarian-based moral code, which may not include any rules that prove overly demanding. Rule-utilitarianism also seems to accord reasonably well with our experience of particular, morally significant relationships. We commonly perceive ourselves as having special obligations arising out of our various morally significant relationships, and we think of these obligations as incompatible with functioning in the manner of an act-utilitarian. For example, parents have a special obligation to care for their children, physicians have a special obligation to act in the interests of their patients, and so forth. Such special obligations can be understood as having a rule-utilitarian foundation, as deriving from rules that, if generally followed, would maximize utility. Thus, rule-utilitarianism seems to remedy another perceived difficulty of act-utilitarianism.

It is less clear that rule-utilitarianism is capable of providing a complete remedy for another perceived difficulty of act-utilitarianism, that is, its inability to provide an adequate theoretical foundation for individual rights. Surely rule-utilitarianism does not lead us as easily as does act-utilitarianism to conclusions that are incompatible with our ordinary moral thinking about the rights of individuals. For example, in suggesting that the painless murder of an old, unhappy (but not suicidal) person is the right thing as long as it can be done in complete secrecy, act-utilitarianism seems to clash violently with our conviction that such an action is patently objectionable, inasmuch as it constitutes a violation of a person's right to life. Rule-utilitarianism, in contrast, would never lead us to the conclusion that

this sort of killing is morally legitimate. Surely the consequences of adopting the rule "do not kill *except* in the case of old, unhappy people who can be killed in complete secrecy" are dramatically worse than the consequences of adopting the rule without such an exception. If the rule with the exception were adopted, the lives of elderly people would be filled with anxiety and fear; moreover, people attempting to follow such a rule would no doubt sometimes kill old, *happy* people, mistakenly believing them to be unhappy. In addition to rescuing utilitarian thinking from such obvious clashes with our ordinary moral thinking, rule-utilitarianism suggests a way of accommodating the notion of individual rights. Just as our special obligations can be understood as deriving from rules in a utilitarian-based moral code, so, too, can an individual's rights be understood in this fashion. A person's right to life, for example, can be understood as a correlate of our utilitarian-based rule against killing. Of course, whatever exceptions are properly incorporated into our rule against killing will factor out as limitations on a person's right to life. Whether rule-utilitarianism in this manner can provide an adequate theoretical foundation for individual rights is a very controversial matter. Its critics charge that it cannot.

Closely related to the claim that rule-utilitarianism does not provide an adequate theoretical foundation for individual rights is the somewhat broader claim that rule-utilitarianism fails to provide an adequate theoretical grounding for what we take to be the obligations of justice. This broader criticism, which is also vigorously advanced against act-utilitarianism, is perhaps the principal residual difficulty confronting rule-utilitarianism. Critics of rule-utilitarianism allege, for example, that the theory is compatible with the blatant injustice of enslaving one segment of a society's population or at least discriminating against this segment. The idea is that social rules discriminating against an explicitly identified minority group might maximize utility by bringing about more happiness in the advantaged majority than unhappiness in the disadvantaged minority. Rule-utilitarians are inclined to argue in response to this line of criticism that, when the consequences of adopting "unjust rules" are completely analyzed, it is never true that their adoption can be justified on utilitarian grounds. Rather, the rule-utilitarian contends, "the rules of justice" rest on a secure utilitarian foundation. Whether rule-utilitarianism, in this manner, can adequately be reconciled with the perceived obligations of justice is a matter of contemporary debate.

Rule-utilitarianism also seems to fare reasonably well when measured against the second of our suggested standards: An ethical theory must provide effective guidance where it is most needed. In a dilemma, where one moral rule, or principle, inclines us one way and another moral rule, or principle, inclines us another way, the rule-utilitarian instructs us to establish relative priority by considering the consequences of incorporating appropriate exceptions into the rules that are in conflict. The dilemma is to be resolved by adoption of a rule that will maximize utility. Although this decision procedure sometimes entails very complex factual analysis and deliberation, it does seem to provide us a substantial measure of explicit guidance. Since rule-utilitarianism also seems to be reasonably harmonious with our ordinary moral thinking, it is an ethical theory that cannot easily be dismissed.

KANTIAN DEONTOLOGY

The most prominent of the classical deontological theories is that developed by German philosopher Immanuel Kant (1724–1804). Kantian deontology continues to command substantial attention in contemporary discussions of ethical theory and, importantly, is the underlying framework of much argumentation in biomedical ethics. In both of these

respects, Kantian deontology is similar to utilitarianism and, like utilitarianism, warrants our detailed attention.

Kant sees utilitarianism as embodying a radically wrong approach in ethical theory. He emphasizes the need to avoid the "serpent-windings" of utilitarian thinking and refers to the principle of utility as "a wavering and uncertain standard." There is indeed a single, fundamental principle that is the basis of all moral obligation, but this fundamental principle is *not* the principle of utility. The supreme principle of morality, the principle from which all of our various duties derive, Kant calls the "categorical imperative."

Although our present objective is an exposition of Kantian deontology, the enormous complexity of Kant's moral philosophy is a formidable obstacle to any concise exposition of the structure of Kant's ethical system. In particular, we are faced with the problem that Kant formulated the basic principle of his system, the categorical imperative, in a number of different ways. Although Kant insists that his various formulations are all equivalent, this contention is explicitly denied by many of his expositors and critics. Thus, if we are to provide a coherent account of Kantian deontology, mindful of the need to provide an account that is especially useful in dealing with issues in biomedical ethics, it seems advisable to settle on a favored formulation of the categorical imperative. Since two of Kant's formulations of the categorical imperative are especially prominent, it will suffice for our purposes to choose a favored formulation from these two.

According to what we will call the "first formulation," the categorical imperative tells us: "Act only on that maxim through which you can at the same time will that it should become a universal law."¹⁰ According to what we will call the "second formulation," the categorical imperative tells us: "Act in such a way that you always treat humanity, whether in your own person or in the person of any other, never simply as a means, but always at the same time as an end."¹¹ The first formulation of the categorical imperative has often been compared to the Golden Rule ("Do unto others as you would have them do unto you"), and it may be true that these principles, when suitably interpreted, have roughly the same implications. At any rate, Kant apparently considered the first formulation to be the most basic of all his formulations, yet despite this fact, and despite the fact that ethical theorists have tended to pay more attention to the first formulation than the second, it is the second formulation that we take to have greater promise for the task at hand. Two major reasons can be advanced for choosing to exhibit the structure of Kant's ethical system in reference to the second formulation of the categorical imperative. First, the second formulation embodies a central notion—respect for persons—that is somewhat easier to grasp and apply than the more formalistic notion of universalizability, which is the core element of the first formulation. Second, when argumentation in biomedical ethics reflects a Kantian viewpoint, it is almost always couched in terms of the second formulation rather than the first.

Kantian deontology is an ethics of respect for persons. In Kant's view, every person, by virtue of his or her humanity (i.e., rational nature) has an inherent dignity. All persons, as rational creatures, are entitled to respect, not only from others but from themselves as well. Thus, the categorical imperative directs each of us to "act in such a way that you always treat humanity, whether in your own person or in the person of any other, never simply as a means, but always at the same time as an end." From this fundamental principle, according to Kant, a host of particular duties can be derived. The resultant system of duties includes duties to self as well as duties to others. In each of these cases, "perfect duties" must be distinguished from "imperfect duties," thus generating a fourfold classification of duties: (1) perfect duties to self, (2) imperfect duties to self, (3) perfect duties to others, and

(4) imperfect duties to others. Although the distinction between perfect and imperfect duties is not a transparent one, its structural importance in the Kantian system is hard to overemphasize. Perfect duties require that we do or abstain from certain acts. *There are no legitimate exceptions to a perfect duty.* Such duties are binding in all circumstances, because certain kinds of action are simply incompatible with respect for persons, hence strictly impermissible. Imperfect duties, by contrast, require us to pursue or promote certain goals (e.g., the welfare of others). However, action in the name of these goals must never be at the expense of a perfect duty. One of Kant's most prominent commentators relates the distinction between perfect and imperfect duties to the categorical imperative in the following way: "We transgress perfect duties by treating any person *merely* as a means. We transgress imperfect duties by failing to treat a person as an end, even though we do not actively treat him as a means."¹²

Our discussion of Kant's fourfold classification of duties begins with a consideration of perfect duties to others. A transgression in this category of duty occurs whenever one person treats another person merely as a means. It is strictly impermissible for person A to treat person B merely as a means because such treatment is incompatible with respect for B as a person. Notice that Kant does not claim that it is morally wrong for one person to use another as a means. His claim is that it is morally wrong for one person to use another *merely* as a means. In the ordinary course of life, it is surely unavoidable (and morally unproblematic) that each of us in numerous ways uses others as means to achieve our various ends. A college teacher uses students as a means to achieve his or her livelihood. A college student uses instructors as a means of gaining knowledge and skills. Such human interactions, presumably based on the voluntary participation of the respective parties, are quite compatible with a principle of respect for persons. However, respect for persons entails that each of us recognizes the rightful authority of other persons (as rational beings) to conduct their individual lives as they see fit. We may legitimately recruit others to participate in the satisfaction of our personal ends, but they are used merely as a means whenever we undermine the voluntary or informed character of their consent to interact with us in some desired way. Person A coerces person B at knifepoint to hand over \$200. A uses B merely as a means. If A had requested of B a gift of \$200, leaving B free to determine whether or not to make the gift, A would have proceeded in a manner compatible with respect for B as a person. Person C deceptively rolls back the odometer of a car and thereby manipulates person D's decision to buy the car. C uses D merely as a means. C has acted in a way that is strictly incompatible with respect for D as a person.

In the Kantian system, among the most notable of our perfect duties to others are (1) the duty not to kill an innocent person, (2) the duty not to lie, and (3) the duty to keep promises. Murder (the killing of an innocent person), lying, and promise breaking are actions that are intrinsically wrong. However beneficial the consequences of such an action might be in a given circumstance, the action is strictly impermissible. (Notice the anti-utilitarian character of Kant's thinking.) The murderer exhibits obvious disrespect for the person of the victim. The liar, in misinforming another person, violates the respect due to that person as a rational creature with a fundamental interest in the truth. A person who makes a promise issues a guarantee upon which the recipient of the promise is entitled to rely in his or her future planning. The person who breaks a promise shows disrespect for another by undermining the effort to conduct the affairs of one's life. By murdering, lying, or breaking a promise, an agent uses another person merely as a means to the agent's own ends.

According to Kant, each person has not only perfect duties to others but also perfect duties to self. The categorical imperative demands that no person (including oneself) be treated merely as a means. It is no more permissible to manifest disrespect for one's own person than to do so for the person of another. Kant insists, for example, that each person has a perfect duty to self to avoid drunkenness. Since drunkenness undermines a person's rational capacities, it is incompatible with respect for oneself as a rational creature. Kant believes that individuals debase themselves in the effort to achieve pleasure via inebriation. Inebriates treat themselves merely as a means (to the end of pleasure). But surely the foremost example of a perfect duty to self in the Kantian system is the duty not to commit suicide. To terminate one's own life, Kant insists, is strictly incompatible with respect for oneself as a person. In eradicating one's very existence as a rational creature, a person treats oneself merely as a means (ordinarily to the end of avoiding discomfort or distress). Suicide is an action that is intrinsically wrong, and there are no circumstances in which it is morally permissible.

In addition to the notion of perfect duties (both to self and others), the Kantian system also incorporates the notion of imperfect duties. Whereas perfect duties require, in essence, strict abstention from those actions that involve using a person merely as a means, imperfect duties have a very different underlying sense. Imperfect duties require the promotion of certain goals. In broad terms, there are two such goals—an agent's personal perfection (i.e., development) and the happiness or welfare of others. Respect for oneself as a person requires commitment to the development of one's capacities as a rational being. Thus Kant spoke of an imperfect duty to self to develop one's talents. The sense of this duty is that, by and large, it is up to each individual to decide which talents to cultivate and which to deemphasize. But a person is not free to abandon the goal of personal development. Although the duty to develop one's talents requires no *specific* actions, it does require each individual to formulate a plan of life that embodies a commitment to the goal of personal development.

Before discussing Kant's final category of duty, imperfect duty to others, it will prove helpful to introduce the notion of *beneficence*.¹³ If one acts in such a way as to further the happiness or welfare of another, then one acts beneficently. (A benevolent person is one who is inclined to act beneficently.) Beneficence may be contrasted with *nonmaleficence*, which is ordinarily understood as the noninfliction of harm on others. One who harms ("does evil" to) another acts in a maleficent fashion. One who *refrains* from harming others acts in a nonmaleficent fashion. One who acts, in a more positive way, to contribute to the welfare of others acts in a beneficent fashion. Beneficence is a generic notion that can best be understood as including the following types of activity: (1) preventing evil or harm from befalling someone, (2) removing an evil that is presently afflicting someone, and (3) providing benefits ("doing good") for someone. Although it is sometimes difficult to decide which of these categories is the most appropriate classification for a particular beneficent action, the following examples seem relatively straightforward. Pushing someone out of the path of an oncoming car is an example of the first type of activity. Curing a patient's disease is an example of the second. Giving someone a \$100 gift is an example of the third.

According to Kant, respect for other persons requires not only that we avoid using them merely as a means (by the observance of our perfect duties to others) but also that we commit ourselves in some general way to furthering their happiness or welfare. Thus, Kant considers what we will call the "duty of beneficence" to be an imperfect duty to others. As with the duty to develop our talents, an imperfect duty to self, the duty of beneficence requires no *specific* actions. One does not violate the duty of beneficence by refusing to act beneficently

in any individual case where the opportunity arises. What is required instead of specific actions is that each person incorporate into his or her lifeplan a commitment to promote the well-being of others. Individuals are free to choose the sorts of actions they will embrace in an effort to further the well-being of others (e.g., contributing to the relief of famine victims); they are not free to abandon the general goal of furthering the well-being of others.

Since the duty of beneficence is an imperfect duty in the Kantian system, action in the name of beneficence must never be taken at the expense of a perfect duty. For example, it is impermissible to lie or break a promise in an effort to save a third party from harm. The same is true with regard to the imperfect duty to develop one's talents. For example, if one has resolved (quite properly) to develop one's creative powers, it is nevertheless impermissible to do so by "creatively" defrauding others.

The Kantian Framework in a Biomedical Context With our exposition of Kantian deontology now complete, we are in a position to exhibit some of the more important implications of this ethical theory in the realm of biomedical ethics. To begin with, the theory has an obvious relevance to the much discussed problem of whether or not a physician may justifiably lie to a patient (an issue discussed in Chapter 2). Since every person has a perfect duty to others not to lie, a straightforward implication of Kantian deontology is that a physician may *never* lie to a patient. If a patient diagnosed as terminally ill by a physician inquires about his or her prognosis, the physician may be much inclined to lie, motivated by a desire to protect the patient from the psychological turmoil that would accompany knowledge of his or her true condition; but action in the name of beneficence (an imperfect duty) may never be at the expense of a perfect duty. This same analysis is relevant to the use of placebos by physicians. Sometimes a patient becomes psychologically dependent on a certain medication. When the medication is discontinued, because the physician is convinced it is no longer needed and because its continued use represents a threat to health, the patient complains of the reemergence of symptoms. If such a patient is given a placebo, that is, a therapeutically inert but harmless substance, misrepresented as a medication, the patient may feel fine. Nevertheless, despite the fact that placebos may be capable of enhancing patient welfare, their use is morally impermissible, at least in cases involving an explicit lie.

Kantian deontology has some very important and very direct implications for the ethics of experimentation with human subjects (a topic discussed in Chapter 4). Since it is morally wrong for any person to use any other person merely as a means, it follows that it is morally wrong for a biomedical researcher to use a human research subject merely as a means. From this consideration it is but a short step to the requirement of voluntary informed consent as a basic principle of research ethics. If a researcher is engaged in a study that involves human subjects, we may presume that the immediate "end" being sought by the researcher is the successful completion of the study. But notice that the researcher may desire this particular end for any number of reasons: the speculative understanding it will provide; the technology it will make possible; the eventual benefit to humankind; personal recognition in the eyes of the scientific community; a raise in pay; and so forth. This mixture of self-centered and benevolent motivations may be considered the researcher's less immediate ends. If researchers are to avoid using their research subjects merely as means (to the ends of the researchers), surely they must refrain from coercing the participation of their subjects and provide information about the research project (most notably, risks to the subjects) sufficient for the subjects to make a rational

decision with regard to their personal participation. Thus, respect for persons demands that researchers honor the requirement of voluntary informed consent.

Suppose a researcher explains to a potential research subject how important it is that he or she consent to participate. There is no question but that the research project at issue, if brought to a successful conclusion, will provide substantial benefit to humankind. Does the potential subject have a moral obligation to participate? Surely not. Within the framework of Kantian deontology, the duty of beneficence is an imperfect duty. A person must on occasion act beneficently, but there is no obligation to perform any *specific* beneficent action.

Critical Assessment of Kantian Deontology Are the implications of Kantian deontology largely reconcilable with our experience of the moral life? Can this theory provide effective guidance in the face of perceived moral dilemmas? These two questions reflect the criteria suggested earlier as most central to the assessment of the relative adequacy of an ethical theory.

Before indicating some of the ways in which Kantian deontology can be thought to be at odds with our ordinary moral thinking, it is important to emphasize that the theory does successfully account for crucial aspects of our experience of the moral life. To begin with, Kantian deontology provides an obvious foundation for the “commonsense rules of morality.” The wrongfulness of actions that fly in the face of these rules—actions such as killing, injuring, stealing, lying, breaking promises—can very plausibly be understood as flowing from the categorical imperative. The Kantian deontologist maintains that these actions are wrong because they involve treating another person merely as a means, and there is something very compelling about the notion of respect for persons as the core notion of morality.

Kantian deontology also seems to provide a secure foundation for the notion of individual rights, a notion that is very prominent in our ordinary moral thinking. Individual rights, in the Kantian system, are to be understood as the correlates of our perfect duties to others. (Imperfect duties, in contrast, do not generate rights.) For example, each of us has a perfect *duty* not to kill an innocent person; thus, every innocent person has a *right* not to be killed. More generally, every person has a right not to be used by another merely as a means. An innocent person has a right not to be punished, no matter how socially desirable the consequences might be in a certain set of circumstances. A potential research subject has a right not to be coerced or deceived into participation, even if the satisfactory completion of the study promises great benefit for humankind. In its insistence that individual rights cannot be overridden by “utilitarian” considerations, Kantian deontology achieves accord with our firmly entrenched (if somewhat vague) conviction that the end does not justify the means.

However, there are aspects of Kantian deontology that cannot be easily reconciled with our experience of the moral life. One very prominent difficulty has to do with the Kantian contention that keeping promises and not lying are both duties of perfect obligation. We are quite at home, in our ordinary moral thinking, with both a duty to keep promises and a duty not to lie, but it is the exceptionless character of these duties in the Kantian system that we find troublesome. Surely in extreme cases, we are inclined to say, these duties must yield to more weighty moral considerations. For example, if a person breaks a rather trivial promise (say, to return a book at a certain time) in order to respond to the needs of a person in serious distress, surely he or she has not acted immorally. Or again, if a person lies to a would-be murderer about the whereabouts of the intended victim, surely the liar has not (all things considered) acted immorally. The Kantian deontologist sees in such examples a clash

between a perfect duty and the imperfect duty of beneficence, and the Kantian teaching is that the former may never yield to the latter. But it would seem that a theory with such implausible implications stands in need of revision. Perhaps the problem is not only that Kantian deontology overstates the significance of certain "perfect" duties but also that it understates the significance of the duty of beneficence, at least that aspect of beneficence that has to do with preventing serious harm from befalling another or alleviating the serious distress of another.

In our everyday existence as moral agents, we are accustomed to the idea that we have a number of important duties to others. It is less clear that the Kantian notion of duties to self can be satisfactorily reconciled with our experience of the moral life. This is difficult territory. For one thing, the issue of suicide (discussed in Chapter 6) seems to confound our moral "common sense" in a way that blatant wrongs such as murder, rape, and slavery do not. Still, despite significant disagreement, suicide is considered by many to be morally wrong. But the issue is this: Do those who consider suicide morally wrong experience the duty not to commit suicide as a duty to self? It seems more likely that this duty is experienced as a duty to others (who may be negatively affected by one's suicide) or, in the case of religious believers, as a duty to God. (Whether a similar argument would be persuasive with regard to the imperfect duty to develop one's talents is unclear.)

It cannot be denied that Kantian deontology, to a substantial degree, is reconcilable with our experience of the moral life. On the other hand, it appears that the theory is attended with some significant and unresolved difficulties. How does Kantian deontology fare when measured against the second of our standards, the requirement that an ethical theory provide effective guidance in the face of moral dilemmas? Once again, it seems, the verdict is somewhat mixed.

It might be argued that Kantian deontology, by sorting our various duties into the categories of perfect and imperfect and assigning priority to perfect duties, provides us with a structure in terms of which moral dilemmas can be resolved. This is perhaps true to the extent that our perplexity can be analyzed in terms of perfect duties marshaled against imperfect duties, but even here it is difficult to overlook the fact that the priority of perfect over imperfect duties is itself a somewhat problematic feature of Kantian deontology. One is tempted to say that, even if the theory provides reasonably *clear* guidance, it sometimes fails to provide *correct* guidance.

W. D. ROSS'S THEORY OF PRIMA FACIE DUTIES

In a book entitled *The Right and the Good* (1930), English philosopher W. D. Ross proposed a deontological theory that has received considerable attention among ethical theorists. The point of departure for the development of Ross's theory is his concern to provide a defensible account of "cases of conscience," that is, situations that confront us with a conflict of duties. One perceived line of obligation pulls us in one direction; another perceived line of obligation pulls us in a contrary direction. We find ourselves unsettled and uncertain but cannot avoid a choice. Which duty takes precedence over the other? The parent of a young child has promised to attend a community meeting, but the child seems to need special attention. Since our social existence is complex, conflict-of-duty situations are a recurrent feature of our daily life. In the biomedical context, such situations are pervasive.

For understandable reasons, Ross maintains that neither the Kantian nor the utilitarian can provide an account of conflict-of-duty situations that harmonizes with what he calls

"ordinary moral consciousness." We have just considered the relevant deficiency in the Kantian approach. It is implausible to maintain that the duty of beneficence can never take precedence over the duty to keep promises or the duty not to lie. As for the utilitarian approach (and here it is clear that Ross has act-utilitarianism in mind), this theory's insistence that in reality we have only the one duty of maximizing utility clashes with our conviction that we have distinct lines of obligation to distinct people. In order to provide an adequate account of conflict-of-duty situations, Ross maintains, it is essential to introduce the notion of "prima facie duty." The Latin phrase *prima facie*, now commonplace in moral philosophy, literally means "at first glance." But the word *conditional* best expresses the sense of the phrase as Ross intends it. A prima facie duty is a conditional duty. A prima facie duty (as opposed to an absolute duty) can be overridden by another prima facie duty that in a particular set of circumstances is more stringent.

According to Ross, there are no absolute, or unconditional, duties (such as "Never lie"), only prima facie duties. But what is the basis of our prima facie duties? Both the utilitarian and the Kantian assert that our various duties have a unitary basis in a fundamental principle of morality. The utilitarian believes that our various duties can be derived from the principle of utility. The Kantian believes that our various duties can be derived from the categorical imperative. Ross, in vivid contrast, maintains that our various prima facie duties have no unitary basis. Rather, they emerge out of our numerous "morally significant relations," relations such as promisee to promiser, creditor to debtor, spouse to spouse, child to parent, friend to friend, citizen to the state, fellow human being to fellow human being. "Each of these relations is the foundation of a *prima facie* duty, which is more or less incumbent on me according to the circumstances of the case."¹⁴

In unproblematic circumstances, where we are bound by only one prima facie duty, this particular prima facie duty is our *actual* duty. In conflict-of-duty situations, where two (or more) prima facie duties compete for priority, only one of these duties, the more stringent one in the circumstances, can be our actual duty. We have, for example, both a prima facie duty to keep promises and a prima facie duty to assist those who are in need. According to Ross, when these two duties come into conflict, it is clear (in terms of our "ordinary moral consciousness") that the duty to keep promises is usually more incumbent upon us than the duty to assist those who are in need. However, if the promise is relatively trivial and the need of another is compelling—a matter of serious distress—then it is equally clear that the priority is reversed. In the difficult cases, Ross maintains, there is in principle no hard-and-fast rule to apply. In his view, the best anyone can do is to make a reflective, "considered decision" as to which of the competing prima facie duties has the priority in any given situation.

According to Ross, "there is nothing arbitrary about [our] *prima facie* duties. Each rests on a definite circumstance which cannot seriously be held to be without moral significance."¹⁵ Accordingly, he proposes the following division of our prima facie duties.

(1) *Duties of fidelity* include keeping promises, honoring contracts and agreements, and telling the truth. Duties in this class rest on a person's previous acts. In giving one's word to do something, a person creates the duty to do so. (Ross thinks that by entering a conversation, a person implicitly agrees to tell the truth.) Notice that a person's so-called role responsibilities can be identified as an important subclass of duties of fidelity. For example, a teacher has certain responsibilities as a teacher, a physician certain responsibilities as a physician, and a nurse certain responsibilities as a nurse. In taking on a certain social role, a person brings into existence various duties of fidelity. In addition, further duties of fidelity arise out of agreements (both explicit and implicit) that a person enters into while functioning in a professional capacity.

(2) *Duties of reparation* also rest on a person's previous acts. Any person, by wrongfully treating someone else, creates the duty to rectify the wrong that has been perpetrated. For example, if A steals a certain amount of money from B, A thereby brings into existence the duty to repay this amount. (3) *Duties of gratitude* rest upon previous acts of other persons, namely, beneficial services provided by them. If A has provided a good service for B when B was in need, B thereby stands under a duty to provide a good service for A when A is in need.

(4) *Duties of beneficence* "rest on the mere fact that there are other beings in the world whose condition we can make better."¹⁶ (5) *Duties of nonmaleficence* rest on the complementary fact that we can also make the condition of other beings worse. The duties in this category, which Ross recognizes as especially stringent, can be summed up under the heading of "not injuring others." The duty not to kill is an obvious example.

(6) *Duties of justice* "rest on the fact or possibility of a distribution of pleasure or happiness (or of the means thereto) which is not in accordance with the merit of the persons concerned."¹⁷ Benefits are to be distributed in accordance with personal merit, and existing unjust patterns of distribution are to be rectified. (7) *Duties of self-improvement* "rest on the fact that we can improve our own condition."¹⁸

Prima Facie Duties in a Biomedical Context Ross's framework of prima facie duties is helpful for conceptualizing many of the moral dilemmas that arise in a biomedical context. In analyzing such dilemmas as they arise from the point of view of health-care professionals, the category of duties of fidelity is especially important. Consider, for example, the physician-patient relationship (a topic discussed in Chapter 2). The social understanding or implicit agreement that underlies this relationship undoubtedly includes a number of important provisions. Among these are the provision that the physician is to act in the best medical interest of the patient and the provision that the physician is to keep confidential any personal information that comes to light within the context of the physician-patient relationship. In the very act of accepting a patient for treatment, a physician thereby incurs a number of important prima facie duties of fidelity.

Suppose a physician is convinced that lying to a patient is in the best medical interest of the patient. In Ross's scheme, the prima facie duty not to lie, itself a duty of fidelity, comes into conflict with another duty of fidelity, the prima facie duty to act in the best medical interest of the patient. Since neither duty is unconditional, in one case the duty not to lie might be more incumbent upon the physician, whereas in another case the duty to act in the best interest of the patient might be the more stringent duty. Suppose, in a different case, a physician is treating a patient suffering from a condition that renders the patient in his or her occupation a danger to others. In addition, suppose that the patient is a bus driver subject to blackouts. The patient is desperate to keep his or her job and refuses to divulge the problem to his or her employer. Should the physician break medical confidentiality and notify the patient's employer in an effort to ensure the public safety? In this case, the prima facie duty of beneficence comes into conflict with a duty of fidelity, the prima facie duty to keep medical confidentiality. (Justifiable exceptions to the duty to keep medical confidentiality are discussed in Chapter 3.)

Among the explicit role responsibilities of a typical hospital nurse is the obligation to follow a physician's orders in the treatment of patients. By the simple act of accepting employment in the hospital setting, a nurse thereby incurs, among other numerous duties of fidelity, the prima facie duty to obey a physician's orders. An important moral dilemma for the hospital nurse arises when, in the judgment of the nurse, following a physician's order

would be detrimental to the patient. (This dilemma is discussed in Chapter 3.) Thinking in terms of Ross's theory, we can structure the dilemma as follows. The prima facie duty to follow a physician's orders comes into conflict with two other prima facie duties. First, there is a relevant duty of nonmaleficence. A nurse should not act in a way that would, in effect, injure another person. Second, there is another relevant duty of fidelity, deriving from the fact that a nurse has an implicit contract or agreement with the patient to act in his or her best medical interest. Is the collective force of these two prima facie duties more incumbent upon the nurse than the prima facie duty to follow a physician's orders? Since the duty of nonmaleficence is recognized by Ross (and "ordinary moral consciousness") as especially stringent, it seems that, in most cases, at least where the potential harm to patients is significant, the nurse must conclude that it would be wrong to follow the physician's order.

Abstracting from any relevant role responsibilities on the part of health-care professionals, the issue of the moral justifiability of active euthanasia (discussed in Chapter 6) might be conceptualized, in accordance with Ross's scheme, as a moral dilemma involving the conflict between a duty of beneficence and a duty of nonmaleficence. A terminally ill person suffering unbearable pain could be understood to benefit from an immediate and painless death. Thus, we have on one hand a duty of beneficence—the prima facie duty to come to the assistance of a person in serious distress—and on the other hand a duty of nonmaleficence—the prima facie duty not to kill.

Critical Assessment of Ross's Theory Since Ross developed his theory of prima facie duties explicitly in reference to the promptings of "ordinary moral consciousness," it would be surprising if his theory could not be reconciled with our experience of the moral life. Indeed, let us put aside whatever worries might be expressed on this score, for there is a much more obvious deficiency in Ross's theory. Recall that we have asked not only that an ethical theory be largely reconcilable with our experience of the moral life but also that it provide us with effective guidance where it is most needed, in the face of moral dilemmas. And despite the fact that Ross's theory provides us with a helpful framework for conceptualizing our moral dilemmas, it provides us with virtually no substantive guidance for resolving them.

In the difficult cases, where two prima facie duties come into strong conflict, Ross holds that there are no principles we can appeal to in an effort to make an appropriate decision. The most we can do, in his view, is render a "considered decision" as to which duty is more incumbent upon us in a certain situation. Although it is fine to be told to make a considered decision, what exactly is worthy of consideration in reaching a decision? At this point, there is a strong argument for moving beyond Ross's theory. One plausible approach would identify *considerations of coherence* (within our overall system of moral convictions) as the relevant standard. (See the discussion "Reflective Equilibrium and Appeals to Coherence" later in this chapter.) If Ross's theory were supplemented with a coherence-based decision procedure, the advantages of thinking in terms of prima facie duties could be combined with a plausible methodology for mediating among conflicting duties.

THE PRINCIPLES OF BIOMEDICAL ETHICS

One prominent approach to problems in biomedical ethics has been articulated by Tom Beauchamp and James Childress in *Principles of Biomedical Ethics*, originally published in 1979. The basic idea is that problems can be appropriately identified, analyzed, and re-