Introduction to Bioethics

Elias Baumgarten, Ph.D.

Table of Contents

CHAPTER ONE: REASONING ABOUT BIOETHICS 1
Bioethical and Scientific Inquiry: Similarities 3
The rejection of feelings 3
The rejection of custom 4
The rejection of authority 5
Bioethical and Scientific Inquiry: Different Meanings of "Evidence" 6
Bioethics and Logic 7
CHAPTER TWO: BIOETHICS AND ETHICAL THEORY 11
Three levels of judgment: particular cases, principles, and theories 11
Two Approaches to Ethical Theory: Overview 12
Utilitarianism 13
What Consequences Are "Good"? 14
Two views of "happiness" 15
Hedonistic versus Preference Utilitarianism 16
Utilitarianism "in Practice" 17
Criticisms of Utilitarianism 18
Utilitarianism does not take into account rights and rules. 19
Utilitarianism does not take into account justice. 20
Consequentialism does not take into account the morally relevant difference between acts and omissions 22
Kantian Ethics 23
The Categorical Imperative 23
Difficulties of Kantian Ethics 25
Other Rule-Based and Rights-Based Approaches 26
Prima Facie Rules or Rights 26
Rule Utilitarianism 27
Which Ethical Theory Is the Right One? 28
APPENDIX: RISK-BENEFIT AND COST-BENEFIT ANALYSIS AS APPLIED UTILITARIANISM 30
Risk-Benefit Analysis in Clinical Practice 30
Cost-Benefit Analysis In Health Policy 31
Problems of Cost-Benefit Analysis: Quandaries of Quantifying Goodness 31

Chapter One: Reasoning about Bioethics

"This is just the kind of case Jack Kevorkian has in mind," Dr. Uberti said.

Dr. Nasser was puzzled. "Rich, how do you mean?

"Look, the whole family is saying to ease his suffering, right? If he were on a vent, we'd withdraw it, and he would die. If he needed a feeding tube, we wouldn't offer one, and he'd die. But now we're stuck. Amal, it just doesn't make any sense that we are allowed to cause his death by *not* giving him a G-I tube or by taking him off the vent, but we can't relieve his suffering in the quickest and most effective way by just. . ."

"By just *killing* him—isn't that what you mean? Of course we can't kill him; we're physicians, and doctors don't kill their patients!" Dr. Nasser was upset. She didn't expect one of her close colleagues to sanction the killing of patients.

"But don't you see? Amal, we do it all the time. We just use different words like 'withdrawing treatment.' But the result is the same."

"The *result* may be the same," Dr. Nasser responded, "but the *intent* is not, and that makes a world of difference."

Does it?

Dr. Uberti has a point. It *is* very common for health care professionals to remove treatment, knowing that the *result* will be a patient's death, the same result achieved by the kind of more active measures advocated by Dr. Kevorkian or proponents of euthanasia.

But Dr. Nasser's position seems reasonable too. Not continuing life-sustaining treatment may have the same result as euthanasia, but the results of actions may not be the only thing to look at when assessing their morality.

Drs. Uberti and Nasser are discussing one of the most controversial questions in bioethics today. At the same time, they are disagreeing about a fundamental issue of ethical *theory*, whether the morality[1] of an action is determined purely by its results.

Whether an action's morality is purely based on its consequences has been a subject of debate among philosophers for over a century. It is one example—and we will see many—where thoughtful deliberation about medical decisions leads inevitably to reflection on larger issues in ethics.

A working assumption of this manual is that our thinking about ethical issues in medicine will be more productive if we are familiar with basic concepts, principles, and theories of ethics in general. Providing that familiarity is the purpose of this text.

Euthanasia and physician-assisted suicide (or, more neutrally, "aid in dying") are dramatic and real issues faced by medical professionals. But less dramatic questions come up everyday, and

[1] There is no general agreement on any distinction between the terms "ethics" and "morality." Following most other writers in bioethics, I use them interchangeably.

many bioethicists believe that *every* medical question includes an ethical component. The following are some "bedside" ethical questions that patients, physicians, and other health care professionals confront regularly:

- How much information must a physician give a patient to meet the obligation to act only after securing "free and informed consent"? And how much does a clinician need to do in order to ascertain that the patient *understands* the information?
- If a patient requests medical care that health professionals consider harmful, which is the stronger ethical obligation: to respect a patient's free choice or to act on the best medical judgment?
- Must a health professional respect the confidentiality of medical information even when a patient has a contagious disease and seems intent on acting in ways that could endanger others?
- Nurses, interns, and medical students are expected to follow the orders of the attending physician, but what should they do if they think those orders are unethical and the attending won't listen to any challenges?
- In deciding on a treatment plan, should physicians think only of what will provide the greatest possible benefit to their *individual patients* or should they also consider whether they am making a wise use of scarce resources that could benefit other patients?

Bioethics is not only for clinicians. In addition to these "bedside" ethical issues, and sometimes as a counterpart of them, there exist equally pressing *policy* questions for administrators, health care planners, and legislators:

- What should the *law* say about physician assisted suicide and euthanasia?
- What should be the *legal* standard of informed consent? What should be the standard of evidence for holding a physician legally liable for "battery" or "unauthorized touching" as a result of acting *without* securing informed consent?
- What policy (or law) should dictate when physicians *must* disclose medical information to third parties and when physicians absolutely must respect confidentiality?
- Should people have a legal right to a hospital's or a country's resources to be given medical treatment regardless of their ability to pay for it? What *level* of medical care: just minimally adequate, the best available, or something in between?
- What restrictions, if any, should the government impose on new reproductive technologies like surrogate motherhood and "test tube babies"?
- What legal standard should determine when a fetus becomes a person with rights?
- What should be the legal definition of "death" and what should hospital policy be with respect to implementing it?
- What laws, if any, should be adopted to regulate cloning and the use of genetic manipulation to enhance human abilities (e.g., improve intelligence)?

Bioethics, as a field of study, is devoted to reflection about moral dilemmas such as these. There is enormous debate in the field about what this kind of "reflection" involves and how it

should proceed, but there is broad agreement that the kind of *reasoning* engaged in by philosophers should have a prominent (even if not exclusive) role. At the most basic level, philosophical reasoning is familiar to all of us; for example, we accept the need to *justify our decisions* and to be *consistent* in doing so. Like philosophical ethics in general, bioethics is committed to seeking the truth using a method of *rational deliberation*. In this very general respect, bioethical inquiry is like scientific or medical investigation. There are other ways in which ethical inquiry is different from scientific inquiry. Before we can look at particular bioethical issues, we should first get clear on what it means to *deliberate rationally* about bioethics and how this is, and is not, similar to the kind of rational method that medical science uses.

Bioethical and Scientific Inquiry: Similarities

Though bioethics is not a *science* like medicine, it shares with medicine the assumption that some beliefs are better or *more worth holding* than others. There may be *good reasons* for some beliefs and only the flimsiest basis for some others, and the belief we should accept (and the one most likely to be *true*) is the one with the *strongest reasons* in its favor. What counts as a strong reason is almost always "debatable," but that does not mean it is arbitrary. A trained medical researcher with expertise in evaluating medical experiments will be better equipped than a first-year medical student to judge whether a study's results lead to conclusions that are worthy of acceptance. Similarly, bioethical training should equip you with a greater ability to assess the "evidence" in favor of different courses of action and help you decide which is more likely to be the right one. "Expertise" in bioethics is different from expertise in medical science: the "evidence" is not empirical, and there is disagreement about what it means to have expertise in evaluating it. Nonetheless, there is increasing agreement on the value of bioethical *training* for medical caregivers, makers of health policy, and even for citizens and patients who are, of course, directly affected by medical decisions. This suggests that ethical judgments are not merely matters of subjective feeling or intuition and that learning the principles and methods of bioethics is an important part of decision-making in medicine.

The rejection of feelings

Like medical science, bioethics rejects subjective *feelings* as a basis for judgment. A physician might have a very strong feeling that a treatment will have a good outcome, but that feeling alone cannot count as *evidence* for a particular conclusion. No doubt an exhaustive account of how good judgments are made would include a role for trained "intuitions" and for leaps of thought that are difficult to specify and articulate, but mere strength of feeling is a notoriously poor guide to decisions in either science or ethics. How we feel is molded in large part by the particular circumstances of our lives—our parents, our friends, and our culture in general. These feelings have the power to cause us to believe certain things, but our being *caused* to believe something by powerful feelings is no guarantee of the worthiness or truth of a belief. Both sides in the debate on euthanasia are probably motivated by sincere feelings, perhaps strengthened by particular personal experiences. In both science and in ethics, emotion can lead us astray; we need to evaluate with reasoning whether the belief we have been led to by our feelings is actually well grounded; that is, whether it is supported with evidence.

There is another problem with using our feelings as a basis for ethics: we ourselves may have feelings that conflict with one another. We then need a method for judging which of our feelings is most *worthy* of being acted upon, and our feelings alone cannot determine this.

The rejection of custom

Bioethics shares with medicine the rejection of custom as a reliable basis of truth. Medical progress and the accumulation of new evidence often requires rejecting what had until that time been customary. Similarly, the fact that a practice is customary does not make it *ethically* sound. It was once customary to experiment on people without the protections we now believe there to be good reason to require. Patients were once customarily shielded from bad news whenever the *physician* judged that to be the better course; for example, forty years ago, 88% of physicians would not disclose a cancer diagnosis to a patient.[2] There may be a sound basis for many customs, but something's being customary does not *make* it right; the customary practice must be scrutinized to see whether there are *better reasons* for retaining it or for modifying or rejecting it.

A classic example where custom was rejected was in the 1972 Circuit Court decision, *Canterbury* v. *Spence*. The issue at stake was how much information the law should require physicians to give to patients about the risks and benefits of a proposed treatment. The decision notes that up to that point "the majority of courts dealing with the problem have made the duty depend on whether it was the custom of physicians practicing in the community to make the particular disclosure to the patient." But the Court rejects that standard and claims "In our view, the *patient's right of self-decision* shapes the boundaries of the duty to reveal [my emphasis]."[3]

The appeal to custom may seem especially powerful when it is made by a patient who chooses to adhere to customs, and this appeal seems even stronger if the patient is rooted in a culture very different from ours. We may regard as arrogant any attempt to impose *our* idea of a "rationally sound ethical judgment" on people whose customs and practices are radically different from our own. We may even begin to doubt that there is anything like universally valid *reasoning* about bioethical issues when someone from a very different culture rejects a practice that seems sound to us and insists that "this is not our way."

Before rejecting the method of rational deliberation in ethics, a few points should be kept in mind. First, the mere fact that cultures *differ* in their ethical beliefs based on different customs does not imply that each of those beliefs is equally true and worthy of acceptance. Cultures also differ in their medical beliefs, and we do not hesitate to call some beliefs well grounded and others simply false. We trust that the empirical method of science is a universally valid basis for determining which beliefs are true. This points to a crucial philosophical question: is there a method of reasoning in ethics, parallel to the method of empirical science, that can resolve differences of opinion and determine which customary beliefs to accept and which to reject? Many philosophers claim there is: that *reasoned argument* can accomplish in ethics what the empirical method has done for science. Without resolving that question—which we will take up again shortly—we should at least be open to the possibility that there may be *good reasons* to reject

---

[2] Donald Oken, "What to Tell Cancer Patients," *Journal of the American Medical Association 175 (April 1, 1961): 1120-1128.*

[3] U.S. Court of Appeals, 464 Federal Reporter, 2nd Series, 772, in Ronald Munson, *Intervention and Reflection*, Fifth Edition, 1996.

customs of other cultures just as we have rejected past customs in our own culture. For example, we might judge there to be good reasons not to grant a parent's request for genital cutting of a young girl even though the practice is customary in that family's culture.

A second point to keep in mind is that the degree to which we should respect diverse customs is, like the more general issue of respecting diverse patient values, itself a question that needs to be assessed rationally. We may decide that there are good reasons to respect some customary practices even when we regard the practice itself as morally problematic. We then need to deliberate, rationally, about the relative weight to give two competing values, the value of respecting the autonomous choice of our patients versus the value of doing what our own sense of reasoning dictates to be the right thing. Weighing conflicting values is a frequent feature of bioethical deliberation, as the debates about both abortion and euthanasia make clear.

The rejection of authority

Bioethics, like science, rejects *appeals to authority* as a foundation for truth. As with custom, there may be good reason to *respect* what others take to be sources of authority, such as a religious leader or a sacred book, and in medicine we may also have a strong obligation to do so. But this is not an absolute obligation; we may be obligated *not* to accede to a patient's request because we judge ourselves to have a stronger obligation not to be implicated in an immoral practice. Again, weighing conflicting obligations requires rational deliberation.

What is clear is that the mere fact that an authority says something does not make it *true* or worthy of our acceptance. One problem with appealing to authority is that authorities disagree, and one needs some basis, independent of any one of the authorities, to decide which authority is most worthy of obedience. If one were to try to decide to choose the authority with the "best" values, one would first need some way of learning which values are best *before* accepting any particular authority. In that case the authority would not actually function as the source of values, and one might as well simply try to learn which values are best directly, without relying on any appeal to authority.

If one tries to make *legal* authority the basis for ethics, then one is in the absurd position of saying that any change in the current laws must, by definition, be immoral. If the very standard of what is moral were current law, then any change in the law would be a deviation from morality. There may, of course, be an independent (though presumably not absolute) obligation for health professionals or patients to obey a law even when they think the law is wrong. But this suggests the need for bioethics also to address policy questions and to ask whether or not a change in the law is desirable. Thus, no bioethicist would say that the *morality* of physician-assisted suicide can be resolved simply by looking to what the present law actually is, but in our role as citizens and policy-makers we need to address a separate moral question: what *should the law be*? Citizens of many states, including Michigan, confront this very question.

A practical difficulty with appealing to authority, particularly religious authority, is that if each person loyally follows the ethical dictates of her chosen authority, there will be little chance of reaching agreement. (We might *first* need to discuss the truth of our respective religious authorities.) At the bedside, in ethics committees, and in State legislatures—particularly in a multicultural society—the only "authority" likely to carry weight with people from diverse backgrounds is the authority of reason.

Bioethical and Scientific Inquiry: Different Meanings of "Evidence"

The standard bioethical approach, then, is to base judgments not on feelings, customs, or authority but on *reasoned argument*. What a repeatable experiment is to a medical claim, reasoned argument is to an ethical claim. And just as a researcher must learn not only to do experiments but to construct the right experiments and draw the right conclusions from them, so one engaged in ethical inquiry must not only offer arguments but learn to construct good, well-reasoned, appropriate arguments. When scientists oppose each other, each has *some* evidence to support his position, and it is even more usual for bioethicists to disagree with one another. Each will offer arguments as "evidence," and then our task is critically to analyze which arguments are the *strongest*.

The parallel between ethical and scientific inquiry must not be carried too far, however. Though both insist that claims be backed up with "evidence," the form of evidence appropriate to ethics is different from scientific evidence because ethical questions are different in their very nature from scientific questions. A scientific problem is generally a dispute about what the *facts* are in a given situation, so the "evidence" for a scientific claim usually takes the form of providing new data that are gained through the *empirical* method; that is, through observation and experience (e.g., experimentation). An ethical problem, in contrast, may exist even after all the facts are known, so the evidence that is needed to resolve the problem will not take the form of new data. An ethical problem focuses not on factual judgments but on judgments about *values*.

Think, for example, of two physicians who disagree on whether abortion is morally appropriate in a particular case. Most likely they already agree on the relevant facts: for example, the structure of the fetus at six weeks, the degree of risk of pregnancy to the mother, and the likelihood that the fetus will develop normally. There may be no new medical or psychological fact to be discovered that would alter either physician's opinion about the moral permissibility of an abortion. Yet they still disagree about *something*, perhaps about what an entity needs to have in order to be considered a *person with rights*. The kind of "evidence" that they need to decide on the ethics of abortion is a *reasoned argument about values*, which is something quite different from simply amassing new data.

The relation between factual (or empirical) judgments and ethical (or normative) judgments is also clear with respect physician-assisted suicide. Factual data, such as the chance of relieving a patient's pain or the odds that the patient will ever be able to communicate again, are crucially relevant. But such factual information is not *sufficient* for making a final moral decision.

An important starting point for discussing any bioethical problem is to separate the relevant empirical (factual, scientific) judgments from the ethical judgments. Physicians and other health care professionals have training and expertise in answering the empirical questions. And in some cases getting clear on the medical facts makes a decision relatively straightforward. But there are other cases—the ones of greatest interest to bioethics—where even after all the facts are known, the moral issue is not resolved and requires further exploration.

Bioethics thus uses a method of inquiry that is both different from most other areas of study, especially the sciences, and also different from the way many people think about moral judgments. Its project is an ambitious one: to recognize that factual judgments alone cannot provide a foundation for moral judgments, yet not to base moral judgments merely on feelings, customs, or appeals to authority. Many people think that if a question cannot be settled by an appeal to *facts*, then one must either depend on some authority or else rely on subjective feelings. Bioethics offers

an alternative approach: it attempts to *justify* ethical claims, to show that some ethical claims are more worthy of belief than others, by asking which claims have *better reasons* in their favor and are therefore *more rationally defensible*.

How can this method of *reasoned argument* show that one claim is more worth believing than another? The best way to answer this question is by actually examining arguments for particular ethical positions, such as those found in articles and books on bioethical issues. But you will be better able to evaluate ethical arguments if you consider more closely how reasoned argumentation works.

Bioethics and Logic

A reasoned argument is trying to reach a *conclusion* (C). It does this by setting out some claims called *premises* (P) or assumptions which, if true, *logically imply* (lead to) the desired conclusion. Another way of saying that certain premises logically imply a conclusion is to say that the conclusion *logically follows* from those premises. The most famous example is

P: All men are mortal.

P: Socrates is a man.

C: Therefore, Socrates is mortal.

Here the two premises clearly logically lead to the conclusion; the conclusion logically follows from them. For the conclusion to be *true*, not only must (a) the conclusion *logically follow* from the premises but (b) all the necessary premises leading to that conclusion must be *true.* Notice that these are two independent requirements. The requirement that a conclusion logically follow from its premises is a purely formal one. A conclusion's "logically following" from premises only means that *if* those premises are true, the conclusion must be true. If the above argument started with the claim that all men are *immortal* and concluded that Socrates is *immortal*, that conclusion would logically follow just as much as the first example. But since the first premise would then be false, the argument does not meet condition (b) above.

Typically a bioethical argument, reaching a conclusion that is an ethical claim, will involve both factual and ethical premises. As discussed in the previous section, factual claims alone, while important, will not lead to an ethical conclusion, so a bioethical argument will always require at least one ethical premise. It is often helpful to set out the premises in an argument both in order to see whether the argument is sound and to distinguish clearly between the moral dimensions of a problem and its empirical or scientific components.

Imagine that I am arguing for the conclusion (C) that *it is morally permissible to withdraw nutrition and hydration from Anthony Lucento*. I might argue as follows, using both ethical premises (EP) and factual premises (FP).

EP: It is morally permissible to withdraw nutrition and hydration from a patient whenever the patient is suffering from severe pain that cannot be relieved and there is almost no chance the patient will live more than a few weeks with continued treatment.

FP: Anthony Lucento is suffering from severe pain that cannot be relieved and there is almost no chance that he will live more than a few weeks with continued treatment.

C: Therefore, It is morally permissible to withdraw nutrition and hydration from Anthony Lucento.

Both of these premises might be controversial. Physicians might well disagree about whether the second, factual premise is true, so even if they agree with the first premise, they will still need to debate the medical facts with respect to this patient. But the *ethically* more interesting issue arises where there is a consensus about the second premise (which actually contains two different factual claims). Then we would need to debate the more general ethical question of whether just these conditions make withdrawal of food and water *morally permissible*. One might argue that withdrawal of food and water is never permissible, perhaps because every moment of life is of infinite value or because food and water is a basic human need and removing it is morally different from withdrawing medical treatments such as ventilators or antibiotics.[4]

Of course this example, meant only to introduce the logical structure of a rational argument, is greatly oversimplified. For example, it makes no reference to Mr. Lucento's own wishes. This might be a good reason for rejecting the ethical premise (EP); perhaps it needs to be amended to include a requirement that the patient or his representative *consent* to the removal of nutrition and hydration. Frequently judgments in medical ethics involve asking two questions: *what is the right thing to do?* and *who should decide?* A complex bioethical dilemma may require assessing many elements involving both of those questions.

Here is another example of a simplified two-step argument that focuses on a patient's claim to decide on her own care.

EP: A patient is competent if she can understand the risks and benefits of treatment.

FP: Maria Ricardo can understand the risks and benefits of treatment.

C: Therefore, Maria Ricardo is competent.

EP: The refusal of treatment by a competent patient should be respected.

FP: Maria Ricardo refuses treatment.

C: Therefore, Maria Ricardo's refusal of treatment should be respected.

The structure of the argument is fine—each of the two conclusions logically follows from the stated premises—but the truth of each premise may be controversial. The first premise is ethically problematic. It might not be obvious why the first claim is ethical rather than factual, but it is important to distinguish the ethical judgment about the criteria needed in order to consider any person competent from the factual judgment about whether a specific patient meets the criteria. (and from the factual judgment of how the law defines competence). Deciding the criteria for competence is a complex moral issue. Competing *values* are at stake: we are obligated to respect patients' choices about their treatment, but we also have a duty to protect vulnerable patients who are unable to exercise a free and informed choice. If we set the standard for competence too high—e.g., requiring (at one extreme) that a patient have medical knowledge equivalent to a physician's—we will deny most patients' free choice, but if we set it too low—e.g., being able to utter decipherable sounds—we will fail to protect many patients who need protection. The *law* must also deal with the question of competence, and "being able to understand the risks and benefits of treatment" is one familiar legal standard. But this is another example of how an appeal to the law cannot resolve the ethical issue. Stating what the law establishes as a standard for

[4] See Joanne Lynn and James F. Childress, "Must Patients Always Be Given Food and Water" (available in the University of Michigan Health Systems Infant and Child Care Ethics Committee Resource Library, and in Munson, pp. 188-195).

competence would be a *factual* claim, but what our *moral* obligation is with respect to a patient's expressed wishes (which may go beyond our legal obligation) is a normative claim.[5]

If one wants to challenge an argument, one will either try to show that one or more of the premises is false or that the conclusion does not follow from the given premises. In all the above examples, our focus was on the truth of the premises because the indicated conclusion did logically follow from them. But you will often confront attempted arguments where the conclusion does *not* follow. These are the *non sequiturs* made famous by advertisers and politicians, but they may even find their way into college classrooms and ethics committee meetings. The best way to show that a conclusion does not follow from the *stated premises* is to supply what is missing, the *hidden premise* that would be needed to make the argument work. And if the truth of that hidden premise is controversial, then the truth of the conclusion may well depend on whether or not the hidden premise is supportable.

A familiar example of an argument with a hidden premise, stated in ordinary language, is: "Killing an innocent person is always immoral; therefore, abortion is always immoral." Notice how in ordinary speech we sometimes state the conclusion before the premise; e.g., "Abortion is always immoral *because* killing an innocent person is always immoral." The structure of the argument is simple:

EP: Killing an innocent person is always immoral.

C: Therefore, abortion is always immoral.

This conclusion would follow from the single premise only with the addition of at least one more (previously "hidden") premise:

EP: Killing an innocent person is always immoral.

EP: Abortion is the killing of an innocent person.

C: Therefore, abortion is always immoral.

Now, of course, the second premise is extremely controversial, and much of the abortion debate revolves around whether the fetus is truly a "person" (by which bioethicists usually mean not just a member of the human species but an entity with full moral status, often including a right to life). One can well imagine the debate continuing with a focus on the second premise, where it could function as the conclusion of a separate argument:

FP: Abortion is the killing of a entity with human genes.

EP: Whatever has human genes is also a person.

C: Abortion is the killing of a person.

Of course the second premise here is also problematic, as is the whole issue of defining *personhood*. (Is an anencephalic infant with no chance of gaining consciousness a person? Is a comatose elderly patient who has no chance of regaining forebrain activity still a person?) But leaving that substantive issue aside for now, we can see that reasoned argument demands of us that we set forth specific claims as our premises, demonstrate their truth (unless truth can be easily assumed) and then draw only those conclusions that logically follow. Often after we identify

[5] In clinical practice, one might approach the ethical question by stating the law's requirements (a factual claim) and then arguing for the controversial ethical claim that clinicians have no moral duty beyond obeying the law. Another dimension of this issue is the policy question: what *should* the law require?

exactly what premises our beliefs rest on, we will see the need to rethink claims whose truth we had previously taken for granted.

How far can we keep going back further and further like this, continually examining more and more basic premises? Well, as long as we want, but in bioethical controversy there is generally no need to claim an absolute "proof" but only to present good reasons that would show a *reasonable* person that a given claim is worthy of belief. For instance, if someone were to argue that "Policy X causes needless suffering; therefore Policy X is wrong," he probably only needs to demonstrate that the stated premise, "Policy X causes needless suffering," is true. Of course, someone could still say, "Okay, you've shown that Policy X causes needless suffering, but there is a hidden premise; namely, "whatever causes needless suffering is wrong." One could *say* this, but it might be fair to ignore such an objection because the hidden premise is just not very controversial or worth arguing about. If someone doesn't see the avoidance of *needless* suffering as a *good reason* to reject Policy X, then maybe there is nothing more to be said.

Generally, there *is* more to be said, and the best way to test the strength of an argument is to see whether it can *withstand criticism*. Thoughtful people disagree on any of the bioethical issues that are worth sustained attention. The best way to test our own ethical positions is to consider the strongest criticisms that would be presented by those who disagree with them and see if we can respond to these criticisms. Most books and articles that develop bioethical positions not only present arguments to support them but anticipate and respond to the kind of critical objections that a thoughtful opponent would offer. Ethics committee meetings often include this kind of dialogue between contrasting positions. Where there is no party who actually disagrees with a line of thought, it is frequently helpful for someone to play "devil's advocate" in the constructive spirit of testing whether a line of argument is as strong as it may first appear to be.

Bioethics can usefully employ this method of logical reasoning because reasoning is universal—unlike feelings, customs, and authorities which differ from person to person, from place to place, and from one historical period to another. This only means that the *laws of logic* are universal: there is not an "American logic" and a separate "Pakistani logic." Of course there will still be disagreement when the reasoning toward a conclusion is more complex than the examples we have looked at and where it is, initially at least, unclear even what the exact premises are or what the structure of the argument looks like. But commitment to *rational reflection* on bioethical issues gives us a shared framework for deliberation and offers us the best hope of reaching a consensus or a reasonable compromise. And even where agreement is impossible, our principled commitment to reasoned reflection allows us to get clear about, and ideally to respect, the sources of our disagreement.

Chapter Two: Bioethics and Ethical Theory

It should now be clear that medical ethics is a field of inquiry committed to a method of rational reflection and that this not just a matter of appealing to our subjective feelings or to current laws and customs or to what our own tradition considers an authoritative source. But all rational reflection must begin with some foundational assumptions, and there is considerable disagreement among ethical theorists about these starting points. This chapter presents an outline of some of the main candidates.

Three levels of judgment: particular cases, principles, and theories

We can think of ethical judgments about actions on three levels, from the most concrete to the most general. The most concrete judgment would be one dealing with a *particular case*; for example, judging that a particular instance of germ-line genetic therapy for a particular patient is not morally permissible. But it is hard to imagine even the least theoretically oriented person remaining completely on this level. To do so would be to take each instance of genetic therapy and to make an ethical judgment only on it, making no attempt to relate the different cases to one another. A person who did this could judge one case of genetic therapy morally permissible and then another case morally wrong, even though it is similar in all relevant respects to the first case.

Most people demand a degree of consistency in the different judgments they make about particular cases of genetic therapy (or anything else). Another way to make this point is to say that we expect people to be able to answer the question, "why?" when they assert, for example, that a certain procedure would be morally wrong. And we expect an answer in the form of a moral *principle*, such as "It is morally wrong to experiment on people without their informed consent" or "Genetic manipulation that can affect unborn persons (as germ-line therapy can) is always immoral." Principles, then, function to explain our judgments in particular cases and to provide consistency among them.

Principles are *more general* than judgments in particular cases, and different principles can also be more or less general. The principle that "genetic therapy without consent is immoral" is not as general as "all genetic therapy is immoral." A person who appeals to a general principle is expected to employ it consistently in all cases where it applies. Someone who held a position against genetic therapy and appealed to the (dubious) principle "All *unnatural* medical interventions are morally wrong" would, if consistent, need to oppose just about *any* medical intervention—or else more clearly specify a meaning of the concept "unnatural."

Formulating a principle that one can hold *consistently* does not guarantee that one's principle is sound, however. One might consistently maintain the principle that "no weight should be given to parents' wishes when deciding on the medical care of children" or "the best treatment plan for each patient is the one which most aids the financial health of the hospital" or any of hundreds of other unacceptable principles. The real challenge is to know which of all the possible ethical principles are the ones most worthy of belief. According to many bioethicists, this requires moving to the level of *theories*, a level even more general than that of principles.

Just as we expect people to be consistent in their judgments on particular cases, we also expect people to be consistent in their principles. A person may hold one principle that covers informed consent for treatment, another principle addressing the use of scarce resources, and still

another dealing with surrogate motherhood. But there must be a reason for holding each of these principles rather than others, and there must be some connection between the different principles we hold. We cannot, for example, claim that "scarce medical resources should go to the patients who can derive the most medical benefit from them" and then decide to offer a liver transplant to a distinguished scientist over a homeless person on the sole grounds that the scientist is more valuable to society. Just as we can ask people "why?" when they assert that a particular action is right or wrong, we can also demand a reason for believing one principle rather than another. Typically the answer will take the form of a more general principle; for example, one might defend the principle that "scarce resources should be distributed solely on the basis of medical need" by appealing to the more general principle that "every person, regardless of wealth or social contribution, has an equal right to medical care." But we can press the demand for rational justification still further and ask for a reason to accept even that more general principle. We can demand a *first* principle from which more particular principles are derived. When we demand this *most general* kind of standard of ethical action, we are at the level of ethical *theory*.

An ethical theory attempts to answer the question, "what makes a right act right?" If sound, a theory tells us what it is that all morally right actions have in common that *make* them right actions. A theory provides a definition of right action.

Theories of right action are theories about *moral obligation*, and it is important to be clear about what bioethicists do and do not mean by this phrase. A basic assumption in ethics is that certain actions are morally required, and it is the job of an ethical theory to tell us what *kind* of actions these are. To say that an action is morally required is also to say that performing the act is a *moral obligation*, a *moral duty*, something that one *should* do, and something one *ought* to do. Each of these is a different way of saying the same thing. To say that a person is "morally obligated" to do something generally means that if that person chooses freely, then she is *morally praiseworthy* if she does the action and *morally blameworthy* if she fails to do it. A person who fails to do an action that she is "morally obligated" to do is "morally blameworthy" (if she chooses freely) because she chose not to do what she ought to have done. But it is important to realize that saying someone is morally obligated to do something is not to say that the person should be *forced* to do it or that any particular *punishment* should be administered if the person does not do it. Those are separate matters, questions of political and legal philosophy. What an ethical theory attempts to do is to define a standard for moral obligation, something that can instruct people who wish to be moral what choices they *should* freely choose to make.

Two Approaches to Ethical Theory: Overview

Moral philosophers have debated ethical theory for centuries. For the sake of simplicity, we can distinguish two broad approaches to ethical theory. One approach, known as utilitarianism or consequentialism[6], judges the rightness or wrongness of an action purely on the basis of the action's likely consequences. The second approach, non-consequentialism (also referred to as formalism, Kantianism, rule-based ethics, or deontological ethics), considers the consequences of an act either morally irrelevant or *only one of several* factors that must be considered in

[6]Moral philosophers differ in their use of the terms "consequentialism" and "utilitarianism." Some philosophers define utilitarianism so that it is one *form* of consequentialism. Here I am using the two terms interchangeably.

evaluating whether an action is moral. Thus, the most fundamental controversy in standard ethical theory is whether an action's consequences are the sole basis for determining its morality.

Consider again the dispute about euthanasia between Drs. Uberti and Nasser in chapter one. For Dr. Uberti, in a situation where the results of withdrawing life-sustaining treatment and the results of actively killing are the same, the morality of both actions is the same. Of course this does not mean that Dr. Uberti or other utilitarians approve of killing patients. Where a patient wants to live or has a reasonable prospect for recovery, *of course* killing patients is wrong—but then so is withdrawing life-sustaining treatment. The utilitarian point is that there is no inherent difference between the two actions. An opposite approach would look not at the results of each action (or not only at the results) but instead focus on the *kind* of act that it is. They might support a moral rule claiming that any action that aims at killing an innocent person is immoral. They might even concede that in *some* cases withdrawing treatment would have worse *results* than active euthanasia—a more drawn out death, pain that cannot be relieved, anguish of loved ones, use of medical resources that could go to other patients—but still claim that the inherent character of the act of killing makes the action wrong.

Many (but by no means all) disagreements in medical ethics are, at their core, based upon a more fundamental disagreement on the level of ethical theory; in particular, whether the morality of an action is determined by its likely consequences.

Utilitarianism

The utilitarian approach to ethics claims that the *only* basis for judging the morality of an act is by its consequences for all those who are affected. If an action is likely to produce, for all affected, a better balance of good over bad *results* than any other action, then, according to utilitarianism, it is the morally right thing to do. One way the utilitarian defends this approach is to claim that there can be no point to moral theory other than to advance *general human welfare*. The utilitarian must be impartial, not favoring himself or the people closest to him. But if an act promotes general human welfare more than any other act, it must be the morally required action.

Though utilitarianism appears straightforward enough, it is hardly uncontroversial. It is possible to imagine medical experiments that could be done without consent, even using deception and coercion, even harming many people, that would yield important medical insights that could save thousands of lives. A consequentialism *might* say such experimentation is morally permissible, even obligatory, if there were no other course of action that could yield the same benefits. Of course the consequentialist needs to take into account the people harmed, whether there is good evidence that the experiments will indeed save many lives, whether there are any long-term, subtle effects that would be destructive, and whether there is any other less harmful way of gaining the same life-saving information. But if, having taken all of this into account (using a method to be discussed further below), coercive experimentation creates more benefits and reduces more human suffering than any other action, the consequentialist would advocate doing the experiments.

It would be a mistake to think that the core problem here is that "we never know for sure that an experiment will save thousands of lives." Our difficulty knowing about the future is a problem of the world we live in, not a problem unique to utilitarian theory. Trying to figure out the likely future consequences makes ethical decision-making complex, but recognizing complexity is what we should *expect* of a good ethical theory. The general question to ask is whether any action is

morally okay *if we do have compelling evidence that it will yield the best results*. There are many cases—perhaps deception with an individual patient to gain consent for what a physician deems to be beneficial surgery or pronouncing death a bit prematurely to gain use of body parts that will help many other people—where the evidence about the likely results is quite strong, yet one may have moral objections to these practices. That objection to the consequentialist approach would rest on something more fundamental than doubt about knowing what the consequences will be; it would be a claim that the morality of actions depends on something *other than* their likely consequences.

By definition a utilitarian insists that there are *no* morally relevant considerations in assessing an action's morality other than the likely consequences. A utilitarian might even point out that in this situation a *failure* to do coercive experiments would, in effect, be causing the death of thousands of people who could be saved. This approach seems at odds with the way many people think about morality, but that alone does not show it to be mistaken. The utilitarian would claim, in fact, that one of the virtues of utilitarian theory is that it challenges the unreflective and largely intuitive morality that many people hold and replaces it with a more rational method.

At the core of the utilitarian approach is the conviction that no action is inherently good or bad just because of the *kind* of action that it is, and this is a feature of the theory that separates it from some other approaches. There are certain *kinds* of actions—doing dangerous experiments without consent, killing innocent people, rape, torture—that generally provoke moral revulsion, but the utilitarian claim is that we should oppose such actions when, and only when, they lead to worse consequences than another available act (or doing nothing). Usually this *is* the case, so generally utilitarians oppose actions like these. But if there are cases—and perhaps some instances of active euthanasia qualify—where killing innocent people produces better results than refraining from killing, then acting with an intent to kill is the *right* action in that situation.

Utilitarians prefer to characterize actions as "right" or "wrong" rather than as good or bad. Whether an action is the right one at a particular time depends on whether the *consequences* that follow from it are good or bad. According to the consequentialist, the only things in the world that can be characterized as "good" or "bad" are the results of actions, not the actions themselves.

If we apply Dr. Uberti's utilitarian approach to active euthanasia, we have an argument where the first premise is the utilitarian standard itself:

(P1) An action is morally required if it is, of all available acts, the one most likely to promote the best possible balance of good over bad consequences for all those affected.

(P2) Active euthanasia would relieve suffering to the greatest degree and thereby promote the best possible balance of good over bad consequences for all those affected.

(C) Therefore, active euthanasia is morally required.

What Consequences Are "Good"?

Though all utilitarians agree that actions are right or wrong depending only on whether their consequences are good or bad, they disagree about what constitutes a "good" consequence. Though this disagreement among utilitarians does not go to the core of the debate between the two different approaches to ethics, we need to say something about what the utilitarian means by "ultimate good" and to examine some of the leading candidates for this title.

There are, of course, many things that we consider good in the world, but when utilitarians refer to the "ultimate good," they mean not those things that are good only because they *lead to* something else that is good but rather those things that are intrinsically good, good in and of themselves. For example, people often consider money a "good," but only because it may lead to other things that people value. Money, therefore, has no intrinsic value; its value is *derived* from something else that is valued for its own sake. In some situations, it may not be good at all. When utilitarians say that an act's rightness depends on its producing "good," they are referring to that which is good in itself, and therefore always good. If we can ask about something, "what is it good *for*?" and can give an answer, then it probably is not an intrinsic good.

What things *are* intrinsically good? We might think of things like friendship, beauty, and kindness, but we need to consider whether these would be good if they did not lead to human happiness or satisfaction. If they are only good because they lead to greater happiness, then they are not intrinsically good after all. But if we ask "what is happiness itself good for?" we may find it impossible to answer; happiness seems to be worthwhile in and of itself. For this reason many utilitarians have proposed that *happiness* is in fact the only intrinsically good thing. Historically, in fact, utilitarianism has often been identified as a theory that approves of actions based on whether they produce happiness (rather than "good consequences" generally). "The greatest happiness for the greatest number" is sometimes thought to capture the essence of the utilitarian approach. The view that our obligation is to maximize happiness has been called *hedonistic* or *eudamonistic utilitarianism*.

Two views of "happiness"

Even utilitarians who agree that "happiness" is the ultimate good may disagree about what happiness is. There is an historic debate between two utilitarians, Jeremy Bentham and John Stuart Mill. Jeremy Bentham, claimed that happiness was nothing but a quantity of pleasure: the more pleasure, the more happiness. John Stuart Mill agreed that happiness is pleasure but argued that *more* pleasure alone does not necessarily create happiness. According to Mill, the amount of happiness an act produces depends not only on the quantity of pleasure that results from it but on the *quality* of the pleasure. Mill thought that some very intense pleasures might be of low quality and not produce as much happiness as smaller amounts of higher quality pleasures.

Mill thought that identifying happiness, the ultimate good, with amounts of pleasure was mistaken and could lead to putting things like wrestling or roller derby (or the equivalent in his day) ahead of Shakespeare and Plato. True, Mill might grant, some people get *more* pleasure out of watching wrestling on television than seeing Shakespeare in Stratford, but they can't be said to be happier because the pleasure is of such low quality. Mill is famous for saying, "better Socrates dissatisfied than a pig satisfied," better the somewhat frustrated experience of "higher pleasures" than the carefree satisfaction of lower ones. For Mill the ultimate aim of our actions should be promoting genuine happiness, and this cannot be equated with mere contentment.

Mill's conception of happiness as the ultimate good has implications for medical care. Presumably, if Mill were in a position to choose between saving the life of a retarded patient who is capable of taking great pleasure only from simple experiences and extending the life of a poet or scientist, he would choose—other things being equal—to promote greater "happiness" by saving a life capable of "higher" pleasures. In deciding the treatment of an individual patient, Mill might consider preserving a person's ability to engage in complex "higher" pleasures to be morally more

important than maintaining a patient's life for a somewhat longer period in which the patient lacked that ability, even if she were *content* during that time. Of course both Mill and Bentham would agree that *reducing suffering* is a good consequence and therefore an important goal of medical care.

Hedonistic versus Preference Utilitarianism

There are, however, forms of utilitarianism that reject happiness itself as the ultimate intrinsic good, partly because of the difficulty of determining what "happiness" means for people with different temperaments and life goals. Unlike hedonistic utilitarianism, which needs to judge, perhaps based on some view of human nature, what outcomes make people truly happy, "preference utilitarianism" sees the ultimate good as the *satisfaction of desires or preferences*, whatever these happen to be.

Implications for paternalism

Whether one regards happiness or the satisfaction of preferences as the ultimate good may dictate different actions with respect to patients who seem to be making the wrong choices. A "happiness (or hedonistic) utilitarian" may be more inclined to favor *paternalistic* behavior, actions that restrict a patient's freedom for that person's own good.[7] To defend paternalistic action, one must claim to know better than another adult what is good for that person, just as a parent could reasonably claim to know better than a child whether it would be good for the child to play in the street. Since the preference utilitarian counts as good whatever satisfies a person's preference, promoting a person's "genuine happiness" by frustrating the person's own desires would not be regarded as an appropriate action.

Paternalistic action is often a part of medical care because physicians may believe that in many cases they *do* know better than even their adult patients what is good for them. A patient may desire to avoid a treatment because of fear or out of an inability to understand the complexities of medicine, and a physician, operating as a "happiness utilitarian," may have good grounds to believe that her obligation is to promote her patients' long-term happiness when what they desire conflicts with their own happiness. The preference utilitarian will have a more difficult time doing this since frustrating a person's preferences is itself a bad consequence.

However, this difference between the two forms of utilitarianism should not be exaggerated. The preference utilitarian does not regard the act of violating a patient's desires as *inherently* bad—no actions are inherently good or bad for any utilitarian—and there are still many things that could outweigh whatever negative value is attached to it. These include not only the desires of other people but the patient's own future desires. Presumably the preference utilitarian will have a harder time justifying paternalism because patients' current desires are generally a better predictor of their own future desires than they are of their "genuine happiness."

Paternalism is an issue for medical policy as well as for clinical practice. Current laws that require people to get a doctor's prescription for certain medications are paternalistic. People would be freer if they could choose to buy whatever drugs they wanted over the counter. But restricting people's freedom to act on their desires may promote overall happiness. The same

[7] The word "paternalistic" comes from a root meaning "father," and to act paternalistically involves treating an adult as a father would appropriately treat his children. A more gender-neutral term would be "parentalism."

might apply to a government's decision not to use Medicare or Medicaid funds for some alternative treatments that are rejected by scientific medicine and yet strongly desired by many people. For the preference utilitarian satisfying those desires will be an intrinsic good. Even though this good could get outweighed if enough other desires were frustrated by the same policy, the preference utilitarian will find it more difficult to override popular wishes on the grounds that "medical authorities know that this treatment will not make you *truly* happy."

The origin of preferences

For many critics of preference utilitarianism, examples such as these suggest one of its weaknesses, that it takes no account of whether a person's current preference is the result of lack of information or the manipulative effects of advertising or propaganda. Some men may have a strong desire to smoke cigarettes (which is still an "over-the-counter drug") because advertising has made them believe that smoking will make them strong and virile. If we make satisfying the actual desires of people the ultimate moral standard, then we cannot take into account *how* people got the desires that they have. We cannot distinguish between "genuine" desires and those that are the result of ignorance or manipulation.

Preference utilitarians respond to this challenge in one of two ways. Some preference utilitarians attempt to devise formulas to exclude consideration of certain kinds of desires, such as those based on ignorance, misinformation, sadism, or even jealousy. (What if a son's desire for his father's death is greater than the father's desire to continue living?) They then need to develop a way to distinguish between preferences that count and those that don't, and it's not clear whether the method of preference utilitarianism itself can do this. Other preference utilitarians simply "bite the bullet" and claim that even if a preference is based on misinformation or suspicious motives, *the preference still exists*, and the ethical action is truly the one that will improve the balance of actual satisfaction over dissatisfaction in the world.

Utilitarianism "in Practice"

Whatever conception of the ultimate good utilitarians adopt, they generally try to quantify the good and bad consequences of each available action and recommend the one likely to produce the best possible balance of "plus points" over "minus points." In assigning these numbers, the utilitarian must consider (a) the number of people who will be affected, (b) the amount or "intensity" of the effect, and (c) the likelihood of the effect actually occurring. Let's imagine that only two actions are possible, A and B. If A produces two units of good for two people and B produces 2 units of good for 4 people, then B would appear to produce twice as many "plus points" (twice as much good) as A. They each have the same "intensity" for each person affected, but B affects more people. But the amount of good produced by an action must be weighed against the bad it creates, so if A produces no harm while B produces 6 units of harm, then A's *net* value (+4 points) will be higher than B's (8 minus 6, or 2).

All this talk about plus and minus points of goodness (or happiness or satisfaction of desires) may seem artificial, and this is a criticism we will need to examine. But there is a common objection to the utilitarian approach that is actually quite weak, and we should dispose of it first. The objection is that utilitarianism is a poor theory because we can't *know* what the consequences will be in the future.

The utilitarian response to this objection is incorporated into the very formula for calculating good and bad consequences; namely, that the *likelihood* of a result occurring is part of what determines its numerical value. This is actually a point we all take for granted in decision-making and is not unique to the utilitarian. If a surgeon judges that a successful operation will produce very good consequences, she must nonetheless consider the likelihood ("the percentage chance") that the surgery will be successful and also the likelihood of bad consequences, such as death of the patient. If surgery to cure blindness has a 97% chance of being successful and only a 0.1% chance of causing the patient's death, obviously there are stronger grounds for going ahead with the operation than if the same chance of success needs to be balanced against a 30% chance of death. The utilitarian is only recommending in the area of ethics what most of us take for granted as sound practice in many other areas. The fact that we cannot know *for sure* what will happen is a feature of the world we live in and of our limited capacities to predict the future. But the question we need to ask, in assessing utilitarianism, is whether *making the attempt* to predict consequences, and basing our decision on that attempt, is a defensible method for making moral decisions.

Imagine deciding between surgery that has a 97% chance of success and a 5% mortality rate and chemotherapy with a 70% success rate where 10% of the patients die from the treatment. Clearly, other things being equal, the surgery is preferable. Someone might object that we can't predict the future and proclaim, "who is to say that in this case chemotherapy isn't better?" But there would be no good reason to heed such a person's objection. If we didn't think it rational to "play the odds" in this way, we might never drive home from our jobs because "who is to say that something wonderful won't happen if you stay overnight at work?" Playing the odds, with the best information we can get (which will always be imperfect), is clearly rational.

The objection to utilitarianism that is based on our not knowing for sure what all the future consequences will be seems to be based on a false hidden assumption; namely, that the best ethical approach is the one that is simplest to follow, the one that tells you without a doubt what is ethically right to do.

The "argument" would be as follows:

(P1) A good ethical theory must be simple and provide definite answers to guide action.

(P2) Utilitarianism is not simple and does not provide definite answer to guide action.

(C) Therefore, utilitarianism is not a good ethical theory.

The conclusion logically follows from the premises, but there is no good reason to believe that the first premise (P1) is true. Acting ethically may be a complex matter, and a theory that tried to simplify it would not accurately reflect the complexity of moral life.

Criticisms of Utilitarianism

One of the central questions of ethics is whether consequentialism is a good ethical theory. Put another way, assume that we are committed to rational principles as the foundation for ethics, assume that we know what the intrinsic good is, assume that we are able to predict future consequences, and assume that we can somehow quantify the good and bad consequences perfectly. Would basing our moral judgment of an action on whether it is likely to produce the best balance of good over bad results be a fully adequate approach to ethics? The most important ethical criticisms of utilitarianism are directed toward showing that this *whole approach* is flawed

because *there are good reasons to consider other things about an action besides what is likely to result from it*. We will examine three criticisms of utilitarian ethical theory, each of which also provides some insight into the spirit of *non*-consequentialist ethical theories.

Utilitarianism does not take into account rights and rules.

The concept of human *rights* is central both to many popular discussions of morality and to the philosophical approach of non-consequentialist ethical theories. To say that someone has a *right* to something means, at the least, that he has a strong claim that does not depend entirely on whether respecting the claim will yield good consequences. For example, if we accept the claim that *all patients have a moral right to full medical information about their condition*, then it may be wrong to withhold bad news from a patient even if doing so will reduce his suffering. A more controversial assertion of moral rights is the claim that *all people*—say, in any country of great wealth like the United States—*have moral right to a certain level of health care* regardless of their ability to pay for it. If people have a moral right to health care, then denying health care would be immoral even if the money were used for another purpose that created a greater amount of happiness or satisfaction.

Those who argue for moral rights disagree both about what particular rights actually exist and whether rights are absolutely binding in all cases. What rights-oriented theorists share is the conviction that there are some claims that have moral force *independent* of considerations of consequences. The most extreme kind of rights-oriented position would give no weight at all to the results of an action and might claim that certain rights—for example, the right not to be killed—are absolute, meaning that nothing could ever justify violating them. A non-absolute rights theorist might hold that certain rights must be respected unless outweighed by some other moral consideration that is even more important. Among these considerations could be other rights or an overwhelming balance of consequences. But even this kind of non-absolute rights position is not the same as consequentialism because it does not judge the rightness of an act *solely* on the basis of consequences. Any rights-oriented theories is committed to the claim that there are, at least in principle, *some* circumstances in which the action which produces the best results may not be the morally correct one. Respecting rights is a moral obligation with independent moral force.

Those who employ the concept of moral rights as an argument against utilitarianism cannot argue that we should respect rights *because* doing so, in the long run, produces better results for all. Anything that produces better consequences is exactly what the utilitarian favors, so this claim is an appeal to the utilitarian standard rather than a criticism of it. Moreover, it would make the need to respect rights depend, in each case, on whether, in fact, better results would be likely to occur. Certainly good consequences may flow from being truthful with patients—it increases trust, for example—but if patients have a *right* to truthful information, the obligation to provide it does not *depend* on these consequences since the whole point of moral rights is to assert a claim that is not dependent on the consequences. Whether moral rights, properly understood in this way, are a necessary part of ethics is a central point of dispute between consequentialists and non-consequentialists.

The utilitarian response to a rights-oriented criticism is to deny that any such rights exist. Jeremy Bentham called the notion of rights “nonsense on stilts,” or high-minded nonsense. If the only purpose of morality is to promote human welfare, then “rights” must be mythical entities because they require one to respect a claim even when doing so *sacrifices* general human good.

Support for the idea of rights, and criticism of utilitarianism for its exclusion of rights, is closely related to an appeal to the need for ethical *rules*. Rules and rights are actually two sides of a coin: a right guarantees something for the person acted upon, whereas a rule spells out an obligation for the person acting. Any rights-oriented principle could be expressed as a moral rule. For example, if we claim that "innocent people have a moral right not to be killed," that implies a moral rule, such as "it is morally wrong to kill innocent people." Anytime we assert that someone has a right, we are suggesting that there is a moral rule obligating other people to respect that right.

Notice that the spirit behind both rights and rules is to characterize certain kinds of actions as inherently moral or immoral, and this is precisely what the utilitarian claims is impossible to do. If the entire basis of an action's moral worthiness is based on its *results*, as the utilitarian claims, then there can be no way to characterize kinds of actions as inherently good or bad irrespective of those results. One of the most fundamental differences between the two approaches to ethics is on this exact issue. The utilitarian claims there are no actions that are inherently moral or immoral just because of the *kind* of acts they are, whereas those who believe in absolute rules or absolute rights may judge certain actions (e.g., acting with an intent to bring about a patient's death or breaking a promise of confidentiality) to be inherently immoral, regardless of the consequences.

It is important to stress that this entire criticism refers to *moral* rights and *moral* rules rather than to legal rights and rules. Though a utilitarian denies any independent status to moral rights and rules, she need not disapprove of legal rules and legal rights. For a utilitarian whether or not to enact a particular *law* is simply an action that must be judged by its consequences, like any other action. A utilitarian would claim that whether or not a particular legal right or rule should be adopted is something that must be decided on the basis of comparing the results that would be likely to follow from adopting or not adopting the particular right or rule in question. But classical utilitarianism, as we have discussed it, cannot endorse a *moral* rule, such as "lying is always wrong," because the justification of such a moral rule would need to be based on something about the quality of lying itself, not—as a utilitarian would insist—on what the consequences would be if people lie or do not lie.

Utilitarianism does not take into account justice.

Justice is the element of morality that focuses on the distribution of goods, privileges, burdens, and punishments. Frequently questions of justice are formulated in terms of what is owed to people or what people *deserve*, and it is therefore closely related to the spirit of rights-oriented ethics. Utilitarianism is criticized for totally ignoring these considerations.

Consider the question of how health care good should be allocated in society. The utilitarian answer must, by definition, be that the morally best distribution scheme is the one that maximizes the benefits of health and extended life to the greatest degree. The utilitarian approach is entirely future-oriented; it looks only at results. The critic of utilitarian argues that there may be other morally important things to consider.

Image that we had to choose between giving a scarce resource, such as a liver available for transplantation, either to an alcoholic with liver disease or to a person who developed liver failure after taking medications prescribed by his physician. Even if the alcoholic were younger and likely to derive more "quality adjusted life years" from the transplant than the older non-alcoholic patient, the critic of utilitarianism might argue that the older patient *deserves* the treatment more

than the alcoholic. The notion of *desert* is non-consequentialist because it refers not to some future result but to something that occurred in the past. Utilitarianism is criticized because it fails to take into account this morally important consideration.

Health policy might be greatly influenced by whether or not these kinds of non-utilitarian considerations are included. If we had an effective treatment for lung cancer caused by cigarette smoking, some might argue that money should instead be used to treat another disease even if using it for lung cancer treatment would save more lives or produce more quality-adjusted life years. Again, the argument might be that those who cause their disease do not *deserve* treatment.

Of course not all critics of utilitarian theory who invoke the idea of desert necessarily agree about what it entails. Leaving aside the question of whether a person who causes her own disease is less deserving of treatment for it, another critic of utilitarianism might argue that when determining who deserves life-saving treatment, it is *not* relevant to consider the quality of life she is likely to enjoy or, especially, how much her continued life means to others. But for a utilitarian those factors would seem to be crucial. If one had to choose between saving the life of a research scientist who is the mother of two dependent children, an unemployed bachelor, or a schizophrenic patient in a mental hospital, utilitarian considerations would appear to dictate choosing the scientist. The critic of utilitarianism might claim that it is *unjust* to consider a person's social worth and perhaps unjust as well to distinguish the quality of life of a "normal" person compared to someone mentally ill, retarded, or in other ways "disabled."

Another argument about the failure of utilitarianism to take into account justice is that if we only look at the consequences, we might unjustly benefit some people *at the expense of others*. Imagine we could enormously improve the effectiveness of medical research if instead of trying to extrapolate from experiments on animals and having elaborate protections for human subjects, we established a "draft" to require some people to submit themselves to medical experimentation. Even if this procedure actually reduced more suffering and produced more extended life and health than the present system, it might still be judged to be *unjust* because some people are benefiting at the expense of others. (The rights-oriented critic could make a similar point by claiming that such experimentation violates the basic rights of people not to be subjected to experimentation without their free and informed consent.)

The utilitarian is likely to argue that no utilitarian would actually sanction a draft like this because it would create so much fear and insecurity in society that the consequences would actually be worse. This is the kind of long-term, subtle consequence that must always be taken into account when applying utilitarianism, and the defender of this theory often makes the point that critics of utilitarianism oversimplify it and fail to consider these less obvious results of our actions.

But critics have an effective response, too. They would insist that what makes a policy of drafting some people for medical research unjust is inherent in the way it exploits some people and is not dependent on whether it actually brings about the *result* of fear and insecurity. What if the policy were to draft only small children in orphanages who would never be told that they might be drafted? Or only the mentally retarded who could not understand the policy and develop a fear of it? The critic of utilitarianism claims that such a draft of vulnerable people would be every bit as immoral and that utilitarianism just shows how unworthy a theory it is when it appeals to "subtle" long-term results in a "desperate" attempt to show that the theory does not lead to such horrible things as it first appears. What makes something like a draft for medical

experimentation "horrible," the critic insists, is not the results but the inherently unjust nature of the practice itself.

Consequentialism does not take into account the morally relevant difference between acts and omissions

Critics of utilitarianism claim that there is a morally relevant difference between acting and failing to act ("omissions"), even if the results in both cases are the same. But for the consequentialist there is no distinction between the consequences that flow from an action and those that follow an "omission" or failure to act. The objection is that this utilitarian position disregards the special responsibility moral agents have for what they actually *do*, in contrast to things they let happen.

Determining the morality of euthanasia (and physician-assisted suicide) rests in large part on precisely this question: is there a morally relevant difference between acts and omissions? A major argument for euthanasia is that it reduces suffering. Imagine a terminally ill patient for whom continued life means unrelenting suffering. The utilitarian argues that if we would be willing to withhold life-sustaining treatment such as a ventilator or feeding tube from this patient, knowing that the *result* will be the patient's death and we judge that result to be better than the alternatives, then we should have no objection to actively bringing about that patient's death, the same beneficial result. (In fact, the utilitarian may add, the result will be better because the patient will be spared a somewhat more prolonged process of dying.) Of course the utilitarian would not approve euthanasia for a patient whose continued life would be a benefit, but in that case a failure to treat—an "omission"—would be equally blameworthy. As we would expect, for a utilitarian it is only the results that count.

According to the critic of utilitarianism, what is central to morality is the idea of a *moral agent*, a person who *chooses to act* in a certain way with a certain *intention*. A moral agent's *action* is not the same as a physical *event* occurring in the world. A natural event such as a disease can have tragic *consequences*, but we do not judge it in moral terms. Malignant cells are not the kind of thing that can be *morally blameworthy*. They are not moral agents that make choices and can be held morally accountable for those choices. The problem with utilitarianism is that it totally disregards the special features of moral agency which make actions different from natural events.

Following from this analysis, the critic claims that it is a certain *kind of action*, acting with an intent to kill a patient, that is immoral. When a physician withholds treatment from a patient, she may know that it is likely the patient will die, but it is the disease that causes the patient's death, not the physician's action. The withholding of treatment may be intended to relieve pain and suffering, but the intent is not to kill, and the morality of an action depends not on the results of the action (or not only on the results) but on the intent of the moral agent. The difference is shown by the fact that if the patient lives even without the treatment, the physician would not then do something to bring about the patient's death, whereas in euthanasia a dose of potassium chloride might be increased if it failed to achieve the action's *intent* of killing the patient.

The nonconsequentialist might argue that there are other areas, particularly areas of health *policy*, where failing to distinguish between acts and omissions leads to absurdity. For example, Congress might know that not appropriating a certain amount of money for a particular disease will result in the loss of lives that would have been saved. But in using the money for some other

purpose, the *intent* of Congress was not to kill those people. Surely, the critic might argue, the failure to save lives is morally different from acting with an intent to kill people.

The utilitarian might respond that there is only a *psychological but not a moral* difference between acting and failing to act where the results are the same. It is true that when we act to kill a person, we know the identity of the person we are killing, whereas when we fail to develop or fund a life-saving program (e.g., to use genetic manipulation to cure muscular dystrophy), we are dealing with anonymous lives. The emotional impact on us is greater when we know the identities of the people we are helping or harming, but according to the utilitarian that just obscures the moral issue. If we could have saved the victims of muscular dystrophy and we fail to do so, we are morally in the same position as if we had killed them.

For the critic of utilitarianism, the failure to distinguish acts and omissions leads to the absurdity that we are responsible for every bad thing that happens that we choose not to prevent, even though our intent might be to do good work in other areas. If Congress does *not* act to fund genetic therapy but instead promotes childhood nutrition, knowing that with limited funds it cannot act in all areas, its *intention* is not to kill those who might have been saved through genetics. People as moral agents can only be held responsible for actions they aim at, not for all the bad things in the world they do not work to change.

Kantian Ethics

Any theory which denies that actions can be morally judged *purely* on the basis of their consequences is a non-consequentialist theory, but non-consequentialists differ about whether consequences are *one* of the factors relevant to judging actions. Some non-consequentialists claim that several factors, including consequences, are morally relevant to assessing actions; others hold that consequences are not relevant at all. The theory most sharply opposed to utilitarianism is that of Immanuel Kant (1724-1804). Whereas utilitarians judge actions by their results and ignore the intentions of the persons acting, Kant denies that consequences have any moral significance and begins his consideration of the moral worth of actions by focusing on the intention of the person acting.

According to Kant actions can be judged morally only because a person, something with rational nature, freely chooses whether to do the right thing. Good and bad consequences can flow from many things besides the actions of moral agents, but those events are not subjects of ethics. Imagine, for example, that on a given day two children are killed: one child is murdered by a person scheming to gain the child's inheritance; the other child is struck by lightning. The death of the two children might be equally bad results, but Kant would insist that we not call the lightning striking the child a "bad action," the kind of thing subject to moral evaluation. What distinguishes the *action* of a moral agent from a mere *event* in the world is that moral agents have intentions; specifically, they can formulate the idea of a moral principle and decide whether or not to follow it.

The Categorical Imperative

Kantian ethics stresses adhering to correct moral principles and is appropriately considered a rule-based ethics. But to say that the right act is one that follows correct moral rules does not give any indication of which rules are the correct ones. A person can formulate and act on all kinds of *wrong* rules, too: "always exploit other people," "lie whenever you can get away with it," and so

forth. So Kant needs to tell us how to know which rules are the right ones, and he does this through his idea of the *categorical imperative*.

The "categorical imperative" can be considered a kind of "super-rule" that instructs us which of all the possible rules are the morally good ones. An "imperative," of course, is a command, something that tells us we should do something. Kant thinks that most of the imperatives we face are hypothetical rather than categorical. A hypothetical imperative is a command that tells us to do something *if* some particular condition holds. It is only binding if that condition exists. The following are some typical examples of hypothetical imperatives:

- If you are an American citizen, you should vote in Presidential elections.
- If you want to earn lots of money, you should acquire a marketable skill.
- If you have diabetes, you should receive insulin treatments.

Hypothetical imperatives, like those above, apply only for people in certain roles or with certain desires or in certain kinds of conditions. Kant claimed that the very idea of morality was of a categorical imperative, a command that applies to all persons at all times without any ifs, ands, or buts. The categorical imperative is the supreme rule of morality, that which we are commanded to do not because we happen to be in a particular situation but simply because we are persons; that is, beings with rational nature.

Kant formulated the categorical imperative as a rule that would test other rules; for a rule to be a *moral* rule, it must meet the test of the categorical imperative. Kant stated this test in three ways but claimed that each of them is simply a different way of saying the same thing (though this may not be immediately obvious). The clearest, most easily understood formulation of the categorical imperative is:

> So act as to treat humanity, whether in your own person or in that of any other, always as an end and never merely as a means.

An underlying assumption of Kant's ethics is that all persons have intrinsic and equal moral worth. For a rule to be a moral rule, it must be one that does not treat persons merely as means but as ends. Persons, according to Kant, must be thought of as having worth in and of themselves, not because they can be useful in achieving some other purpose, however noble it might be. For Kant it is never moral to "use" people, even in order to help a greater number of people.

Kant's requirement that we treat all persons as "ends" takes the equal worth of persons farther than utilitarianism. A utilitarian must be impartial in judging the effects of actions, and each person must count equally when judging good and bad results. But utilitarian theory allows, indeed requires, that we balance some people's good or harm against that of others. This means that even under the impartial standard followed by consequentialism persons may be used as means.

Imagine that doing a painful experiment against a person's will is the *only* way to save fifty other human lives. From an impartial utilitarian perspective, it seems likely that doing the experiment is the morally correct action. You are not *favoring* each of the people whose lives can be saved; it's just that there are more of those people, so considering each person equally results in "trading off" one life in order to save fifty. Kant opposes this kind of "trade off" because it treats a person as a means and fails to recognize the intrinsic worth of that person.

Kantians criticize utilitarians for lacking "respect for persons"; that is, for disregarding the inherent dignity and literal "integrity" (unity) of each person. From a Kantian point of view, what

utilitarians do is equivalent to merging all human consciousnesses and considering them to be one vast container of good and bad experiences. For a utilitarian, an action is right if the container (the world of conscious beings) is kept at the highest net value. But this ignores the fact that the world's supply of conscious experience is divided into separate persons, *each* of which has inherent worth.

Kantians see utilitarians doing with a group of persons what is only reasonable when restricted to the experience of a single person. An individual person may trade off good and bad experiences. For example, you may decide to endure thirty minutes of pain at the dentist in order to gain years of freedom from dental problems. You may choose to suffer through many minus points of pain in a residency program in order to achieve the benefits of a medical degree. It is reasonable for a single person to regard moments of experience, near and far, as of equivalent value and to decide to sacrifice the quality of some moments to gain much greater benefits at other moments. But what the utilitarian does, and what Kantian theory forbids, is to trade off one person's conscious moments for that of other people. It merges all the conscious moments together and claims that it is reasonable to make some people suffer if enough other people can benefit. Kant's imperative that people never be treated as "mere means" deems all such actions as morally unacceptable.

Kant's vision of the perfectly moral society is one that is a "kingdom of ends," one in which there are no superiors or inferiors, one in which each person recognizes the equal and inherent worth of every other person. This does not rule out all relationships where one person serves another. A person may choose to volunteer for a medical experiment, but this is morally acceptable if and *only* if the person freely consents. The core ethical imperative in medical experimentation that a person must give *free and informed consent* follows from the Kantian idea that persons have an intrinsic worth and dignity and must be treated as ends.

Difficulties of Kantian Ethics

Though the "spirit" of Kantian ethics has been enormously influential, Kantian ethical theory has also been subject to severe criticism. There is much debate about exactly what rules would or would not be acceptable under Kantian ethics. Since Kant's own discussion remains very abstract, various parties can quite reasonably make the claim that he would or would not approve of the morality of certain practices. Another problem is that Kant offers us no way to determine how to resolve conflicts between different rules. Kant himself thought that no two rules that met the test of the categorical imperative could ever conflict, but many interpreters of Kant disagree.

One of the most frequent criticisms of Kant is that his rules are too rigid. Kant specifically claims that lying is inherently immoral, and of course for Kant no reference to the disastrous *consequences* of telling the truth are morally relevant. But imagine that you have been treating both a mother and child after an auto accident. The child has just died, and you have good grounds for believing that lying to the mother about the child's death for the next 24-48 hours is necessary to avoid a trauma that would almost certainly take her life. Many critics, including many who would not go to the other extreme of accepting a purely utilitarian standard, consider an absolute prohibition against lying unjustifiable. It may be that lying is not justified whenever better consequences would be created by lying and yet not be the case that lying is *always* immoral. Finding this "middle ground," possibly a synthesis of consequentialist and Kantian ethics, has been the aim of many ethical theorists.

Kantian rules may often be thought to correspond to rights. If people are obligated to follow a rule to keep their promises to other people, then we may also say that persons have a *right* that promises made to them be kept. The issue of how rigid or flexible a rule should be parallels the question of whether rights are absolute. It is consistent with Kantian ethics to suppose that certain actions, actions which treat people as mere means, are violations of rights. Kant's position (based on his position on rules) seems to be that it is never permissible to violate a person's rights, but this, too, may be overly rigid.

Imagine that the *only* way to save thousands of children's lives would be to do a dangerous and painful experiment on one person who refuses to give his consent. Experimenting on someone against his will is, as we've seen, a classic case of using a person as a "mere means." We might well suppose that Kant is right to think that there is something inherently bad about doing this and that it would be wrong to do such an experiment even when it would allow us to bring about slightly better *results* in the world than not doing the experiment. If so, we agree with Kant that morality is not exhausted by a consideration of consequences. But if the effects of not doing the experiment were *overwhelmingly* bad, we might think that there comes a point when it is morally appropriate to infringe the *right* of a person not to be "used" in an experiment. After all, each of the other people whose lives will be saved are equally innocent, equally valuable. But Kant's rule against using people as a means in this way is absolute, and so is the corresponding right of a person not to be used. To many of Kant's critics, this approach to morality is too rigid.

Kantian ethical theory is, along with utilitarianism, one of the two main approaches to contemporary ethics. However, most Kantians do not accept Kant's own formulations in their original form. The task then becomes to offer a theory in the spirit of Kant that captures the essence of his theory, the idea that each person has an intrinsic value that cannot be "traded off," while allowing for the flexibility that is needed to serve as a standard for concrete decision-making by flesh and blood human beings.

Other Rule-Based and Rights-Based Approaches

If we wish to remain open to the idea that an action may be moral or immoral for reasons other than its likely consequences but we cannot accept the idea of *absolute* rules (as in Kant) or *absolute* rights, then we need some ethical theory that combines the best elements of both approaches. These will, by our definition, be *nonconsequentialist* theories since they do not base the morality of an action *exclusively* on its consequences.

Prima Facie Rules or Rights

One approach to ethics often associated with the philosopher W.D. Ross is a theory of *prima facie* duties. In Ross's ethics there are a number of rules that define duties that people generally have, some of which are consequentialist in nature (e.g., promoting good for others, not harming others) and some of which are nonconsequentialist (e.g., telling the truth for its own sake, compensating others for past wrongs we may have done to them). In a situation where only one rule applies, then that becomes our actual duty. But if a rule conflicts with another rule—for example, if our consequentialist obligation to promote a patient's good conflicts with our duty to tell the patient the truth—then each rule should be regarded as only a *prima facie* duty. "*Prima facie*" means "on first view" or non-absolute. In other words, each duty has some strong moral

standing on its own but only on first view. On second or third view—"in the final analysis"—one may determine that one duty is outweighed by others.

Ross's idea of *prima facie* rules can, of course, be extended to rights as well. One might disagree with the utilitarian rejection of moral rights and yet hold that there exist rights that are not absolute. Just as one may be in a situation where two rules conflict, one might confront a conflict of rights. Abortion is a classic instance where two proposed rights, the right of a fetus to develop and the right of a mother to control her body, are often thought to conflict.

Ross develops the theory much further than this brief account suggests, but even in its full form, Ross does not tell us exactly how to decide when one duty outweighs another. He offers some guidelines but claims that in the end, we must use "our best moral intuitions." What is most important about his approach for bioethics is that it opens up the possibility of going beyond utilitarianism to develop a rule-based or rights-based ethical framework and yet not be committed to Kant's idea that such rules or rights admit of absolutely no exceptions.

Rule Utilitarianism

Since the original development of utilitarianism, philosophers have formulated a theory known as "rule utilitarianism," which *does* appeal to moral rules. What we have previously discussed as "utilitarianism" is then designated as "act-utilitarianism" to distinguish it from rule-utilitarianism. Rule utilitarianism claims that we should follow those moral *rules* which, if followed by *everyone*, would yield the best consequences..

Rule-utilitarianism" differs both from "regular" or "act" utilitarianism and from a fully rule-based ethical system like Kant's. In act utilitarianism the question is what the consequences would be of *this particular* action, and an action is morally right if the results are likely to be better than the results of any alternative action. In rule utilitarianism an action is right because it adheres to a good rule, and a good rule is one which, *if followed by everyone*, would produce the best results. Thus it could be that the act of treating a particular patient without securing informed consent would have good consequences, but the rule-utilitarian would ask whether there would be good consequences if everyone followed the rule (or policy), "treat patients without securing informed consent." It is likely that the consequences of that general policy would be worse than following a rule that required informed consent.

Rule-utilitarianism also differs from a Kantian approach or any rule-based approach that sees the goodness of a rule as based on some quality of the rule itself. For Kant a rule is a good rule because it respects the inherent dignity of each person and treats each person as an end. For the rule-utilitarian the only way of determining which rules are good rules is to see what the consequences would be if all people actually followed the rule.

One problem with rule-utilitarianism is that it is not always clear how specific to make a proposed rule. For example, one might try out the rule "never lie to a patient" but decide that a more specific rule, such as "never lie to a patient unless doing so will prevent that person's imminent death" is a better rule. What would *make* it better is that better consequences would result if everyone followed it than if everyone followed another rule. Some critics of rule-utilitarianism claim that when more and more conditions and specifications are added to the rule, it becomes no different from act utilitarianism. Though rule-utilitarians agree that no proper names can be used ("do not lie unless your patient's name is Jeremy K. Outofit"), enough conditions might be added to a rule so that it pretty much defines the exact situation one faces and

then the test of the rule becomes similar to the act-utilitarian approach of examining what the consequences will be of performing a particular act.

Which Ethical Theory Is the Right One?

You may be tempted to say that for some issues utilitarianism is the right theory but that something closer to Kant's ethics works best on other issues. However, to pick and choose a different ethical theory for each issue is to abandon the whole project of finding a consistent standard for moral action. The whole point of an ethical theory is to guide us to the right action in particular cases and to insure consistency in our judgments. If we first decide which action is right and then invoke the theory that accords with our particular judgment in a given case, we are using the theory more as an after-the-fact rationalization than as a guiding standard.

It is fair to point out that recently some bioethicists have expressed doubts about the possibility of finding consistent standards for moral action. One might say they have come to share a skepticism about "principlism" that has been more widespread among physicians. This skepticism may be based on the idea that the diversity of cases is so great that no principle, or set of principles, can ever capture what makes a course of action morally appropriate. Or it might rest on the notion that a special kind of "intuition" tells us what is morally right in a way that no purely rational-deductive method—reasoning from premises to conclusions—could ever adequately articulate. But even if there is no one clear ethical theory or "first principle" that can guide us, there seems no *rational* alternative to exploring principles and attempting to determine which ones *justify* action. If we abandon the project of justifying our behavior with rationally defensible principles, we seem to be left only with our subjective feelings. Even Ross, who is in fact criticized for relying too much on "intuition," does not base his ethics *entirely* on intuition. For him it is only at the final stage of decision, and only when carefully considered duties conflict, that intuition is called upon, and Ross does not consider these "moral intuitions" to be just subjective feelings.

Even if we judge that no one theory tells the whole story about what makes an action morally right, we do not need to abandon the search for a defensible ethical theory. We may try to formulate higher-level criteria to instruct us *to what extent* to use consequentialist thinking and *to what extent* to employ nonconsequentialist principles. Another way of putting this is that we need a new ethical theory that incorporates the best elements of the primary theories.

One way of developing principles and theories in bioethics is to examine a wide range of particular cases. The case method is the one most often used to teach bioethics in medical schools and residency programs and on hospital ethics committees. By dealing with the diversity and "thickness" of real-life cases, we may come to see that we appeal to certain principles in which we gain increasing confidence.

This appears paradoxical: theories are supposed to be guidelines for particular cases, yet in the case method we are developing our theory in light of its application to actual cases. But this procedure has analogies to the method of science. In science we develop general explanations or theories to explain phenomena, and the more phenomena a theory explains, the more confidence we have in the theory. If a theory in which we have confidence seems to conflict with what we observe in a particular case, we may decide that our initial observation is wrong, but we might also decide that our theory needs to be further refined. Similarly, an ethical principle that has worked in many cases may seem to imply something unsound in a new case. It is possible that the

theory is flawed. But it is also possible that our judgment in the particular case is for some reason distorted, perhaps by an emotional consideration that we should not allow to influence us. Having that principle in mind may help us correct for such distortions.

Developing a sound ethical theory may require continual revisions and refinements, and one of the ways of testing a theory is to see whether it accords with our judgments on particular cases in which we have the greatest confidence. For example, we are probably more confident that "it is wrong to transplant a vital organ against the wishes of a competent patient" than we are of any theoretical standard we might formulate. Therefore, we might safely conclude that any theory that implied that we are morally required to take organs from living, competent patients against their wishes is a theory in need of revision.

In practice we go back and forth between judgments in particular cases and judgments on larger principles and theories. One famous philosopher, John Rawls, speaks of our seeking a "reflective equilibrium," the point where our judgments about ethical theory are in balance or in "equilibrium" with our deepest intuitions about particular cases. This is an ongoing project and probably an accurate characterization of how we search for ethical clarity in bioethics.

Appendix: Risk-Benefit and Cost-Benefit Analysis as Applied Utilitarianism

Some may regard the utilitarian idea of assigning plus and minus points to good and bad consequences as artificial, but this is a method that is widely employed in health care. Risk-benefit and cost-benefit analysis are accepted—and sometimes criticized—methods of decision-making; they are basically "applied utilitarianism." Of course clinicians, hospitals, or governments rarely adopt the full utilitarian point of view that would require considering all persons in the world equally. Perhaps for non-utilitarian reasons, physicians consider the effects their actions have on their own patients more than on other people, and governments weigh the interests of their citizens more than the welfare of foreigners. But even though clinicians and policy-makers do not adopt the fully impartial point of view, they do, *within the range of their concern*, often employ a method of deciding policy based on quantifying consequences that parallels that of utilitarian theory.

Risk-Benefit Analysis in Clinical Practice

Risk-benefit analysis to decide among various treatment options seems, on the face of it, uncontroversial. It is a rational method that looks at the likely outcomes and helps to overcome distortions introduced by emotions such as fears and false hopes and the unwarranted comfort of what is customary and familiar. It may seem hard to imagine anything *else* to look at except the likely results of different treatments, once the decision is made to "do what is best" for a particular patient.

There will, of course, be problems *applying* this method. Though it is important to see that these are not themselves criticisms of the whole approach, the issues are serious and worth mentioning. One problem is determining what is to count as a "successful" outcome. "Percentage of cases in which life is extended for five years or more" is easily quantified, but if qualitative measures such as "return to normal life" are introduced, it will be much more difficult to determine what constitutes success. Similar problems exist in quantifying just how bad certain side effects might be. The hedonistic utilitarian would have to determine how much any given complication impairs happiness. Following preference utilitarianism, one could survey people to get some sense of how strong a preference people have to avoid particular risks, but that average would mean little in relation to particular patients, even if the average were broken down into different subgroups of patients. If one treatment poses a risk of a slight hearing loss and another treatment poses the same risk of a slight loss of sexual function, it is not obvious how to quantify the respective deficits for a particular patient. Your patient who is an elderly musician must surprise you—if you asked.

A criticism of risk-benefit analysis, which goes beyond just a problem of application and could be asserted even if all problems of application were overcome, is that it does not take into account any right a patient may have to decide to refuse the most "rational" treatment option, or any right a clinician may have not to be implicated in a procedure she finds morally objectionable. The lack of consideration for moral *rights* is part of a more general objection to the whole utilitarian approach and is discussed in chapter two under "Criticisms of Utilitarianism" (see page 19).

Cost-Benefit Analysis In Health Policy

Policy-makers frequently must decide, given limited resources, what programs *should* be pursued and which are either not desirable to pursue or of lower priority. An underlying assumption of cost-benefit analysis, like the utilitarian theory on which it is based, is that the morally best policies are those that maximize good results at the least expense or, to put it in policy-maker's terms, those policies that are economically most efficient.

Cost-benefit analysis may, for example, be used in evaluating public health policies. The government may need to decide whether to enact stricter air pollution regulations; in particular, more stringent controls on automobile emissions. The benefits or "plus points" resulting from such a policy might include improved health, a lower mortality rate from respiratory disease, and the aesthetic value of clearer air. The costs or "minus points" will be the dollars required for such a program: government inspections, higher automobile prices that everyone will have to pay, and the resulting inflation. The ethical question is whether the government *should* expend some of the society's limited wealth on this program, rather than some other program.

Cost-benefit analysis aims to answer this question by quantifying the good and bad results likely for a proposed policy. Ideally, decision-makers would then choose those polices that are most likely to produce the best ratio of positive over negative consequences. The benefits of a particular clean air program could be compared to alternative air pollution programs and also to proposals for spending the same amount of money both on other health programs and on programs that yield other benefits, such as national security, education, or housing.

One increasingly popular and controversial use of cost-benefit analysis is *outcomes research* to determine which procedures yield sufficient benefits to justify their costs. A health maintenance organization (HMO) might determine which procedures to cover based on studies of patient outcomes. The "benefits" of each procedure would need to be quantified; e.g., the percentage of patients who survive five years or longer after chemotherapy, the proportion who return to pain-free ambulation after surgery, the ratio of patients who return to a "reasonable quality of life" after dialysis. The goal is to eliminate the subjective elements of an individual physician's judgment and replace them with objective data based on the actual results achieved by a large number of physicians with a very large number of patients. If the goal of health policy is to produce the greatest amount of good health for the greatest number of people while expending only a finite amount of scarce resources, cost-benefit analysis promises to help determine which medical procedures are worth paying for.

Problems of Cost-Benefit Analysis: Quandaries of Quantifying Goodness

In theory cost-benefit analysis is the ideal utilitarian tool because it offers a mathematical way to compare the "good" coming out of different actions and thereby to resolve the ethical questions that clinicians and policy-makers confront. In practice cost-benefit analysis is more problematic. Its success depends on being able to quantify benefits. Though cost can usually be expressed quantitatively in terms of the amount of money a program requires, the benefits, such as "reasonable quality of life," are more difficult to quantify. This section offers an overview—admittedly quite simplified—of some of the ways this is attempted.

The problem of quantifying benefits exists both on the level of clinical practice and on the institutional or social level of health policy. Clinically we might want to know how much good a certain procedure, such as a bone-marrow transplant for a patient with metastatic breast cancer,

will *on average* bring about, based on past experience with other similar patients. On the level of policy, those who must decide which treatments Medicare, Medicaid, or an HMO will pay for need to determine how much good a treatment like bone marrow transplants is likely to produce. The "goods" that medicine aims for are generally the extension of life and the improvement of patients' quality of life, so policy-makers often employ the notion of "quality adjusted life years" (QALYs) to measure these goods. Two approaches are possible. One is to start with a specific amount of money that is available and then to compare how many QALYs would result from different procedures or policies. If spending a million dollars on bone-marrow transplants will produce more QALYs (longer and improved lives) than spending the same amount of money on heart transplants, then bone-marrow transplants are a more *cost-effective* use of a scarce resource (money). In this approach, sometimes called cost-effectiveness analysis, a measurable benefit is compared only with another measurable benefit, and no attempt is made to judge whether the number of extended and improved lives is actually worth a million dollars or what exact amount of money each benefit justifies spending. The even more ambitious approach of cost-benefit analysis tries to judge how much money it is *worth* to achieve a particular benefit, such as extending a life. To do this, one must have some way of determining the monetary value of a human life.

Assigning a dollar value to human life raises disturbing issues, but before deciding categorically that doing so is obnoxious and immoral, one must ask what alternatives are available to one who aims to make moral choices with limited resources. The cost-benefit analyst insists that a rational, mathematical procedure is likely to produce greater benefits than a decision influenced by emotional considerations that may be psychologically powerful but morally irrelevant. So, for example, the cost-benefit analyst may decide that, given the resources available for health care, it is "too expensive" to treat elderly cancer patients, but his justification will be that the same amount of money could *do more good* somewhere else; for example, vaccinating children. It is important to realize that he is not saying that treating older patients is morally right but impractical; rather, he is saying that funding these programs would be morally wrong because they would not create benefits (e.g., QALYs) as great as could be created if the resources were used differently. For cost-benefit analysis, and for utilitarianism generally, morality requires using all resources—not just money, but limited time, as well—in the most efficient way, the way most likely to produce the greatest benefits with the least expenditure. (Of course a separate and important ethical question is what proportion of our *total* budge should go to health care programs.)

Applied utilitarians have developed methods for assigning dollar values to human lives. The usual method is one based on preference utilitarianism, which judges the amount of "good" in a human life by how strongly a person actually prefer to live. Cost-benefit analysis adds that the strength of this preference is to be measured in the dollars they are willing to spend to avoid death. Generally they do not ask people how much they would spend to preserve their lives (because the amount might be unlimited) but how much they would be willing to pay to have their *risk* of death reduced. The average dollar value of each life can then be calculated, based on how much people are willing to pay to increase their chances of living.

For example, imagine that the government is considering a water treatment program that would reduce the death rate from cancer by 1%. To decide whether the benefits outweigh the costs, surveys would ask people how much they would be willing to pay to have their risk of dying reduced by 1%. Since 1% is 1/100th, whatever the average amount people are willing to

pay would be multiplied by 100 to determine, in numbers of dollars, how strong a preference people have to live and, for policy purposes, how much "good" preservation of a single life represents. Ideally, this program could then be evaluated against the amount of good expected from other life-saving programs (and the good of other programs that merely improve the quality of life), and the most "cost effective" programs would be the ones judged most worthy of adoption.

The intensity of people's preferences can also used to adjust the value of a year of life for its quality. If a survey of people shows that, on average, most people would choose to be deaf only if they could live twice as long as they would live with full hearing, that shows that people consider the quality of a life without hearing to be half of the quality of life with full hearing. If people said they would exchange a life with monthly migraine headaches for a life 10% shorter than the one they otherwise would expect, they are judging a year of life with migraines to have only 90% of the value of a life free of migraines. The aim is to use surveys of this kind of translate the intensity of people's preferences into a measure of how much "benefit" to assign to a life with or without a particular disability and thereby to develop a measure of the "quality-adjusted life-years" likely to be produced by a given expenditure of resources. One could also try to determine the dollar value of removing a certain kind of disability by asking people how much they would be willing to spend to have the disability removed or how much they would need to be paid to judge that their life with both the disability and the extra money is as desirable as a life without either.

Many people have strong reactions either for or against this approach, but it is important not to make a final judgment too quickly. In particular two extreme reactions are worth avoiding.

One extreme is to endorse cost-benefit analysis or applied utilitarianism as a scientific answer to solving moral dilemmas. On this view traditional ideas about moral wisdom, careful judgment, and the judicious weighing of arguments are expendable because all we need to do is to calculate accurately—or to find computers that can—and then churn out the plus and minus values. Indeed, the original aim of utilitarianism was to provide a scientific foundation for ethical judgments, and cost-benefit analysis might be thought to have accomplished this.

However, it should be clear that cost-benefit analysis itself requires making many ethical judgments that cannot be made scientifically. It is not clear that people can judge accurately either how much they are willing to spend to reduce a risk to their lives or what fraction of a year of life they would be willing to give up in order to be free of a disability. Sometimes it depends on how the question is asked. If people are asked how much they would pay to have the chance of *their own* death reduced by 5%, they will be much more generous than if asked about their willingness to pay to have the death rate in *their community* reduced by 5%, though of course the two questions are logically equivalent. And if one wants to know how strongly people prefer to avoid a certain *disability*, one will get different results depending on whether one surveys only those who have experienced the condition or the whole population.

These, and other technical issues, jeopardize the objectivity of any cost-benefit analysis. These problems do not challenge the desirability, in principle, of using the method in the best way one can. However, they should make people extremely skeptical about the actual use of cost-benefit analysis in public life. There are so many variables that can be played with that any consulting firm that uses a cost-benefit analysis and comes out with a result different from the one wanted by the organization or interest group that hired it can simply change one of the variables. But even if one is not so cynical about the way cost-benefit analysis is actually used, one still must realize that

there is no clear, scientific way to answer ethical questions such as how much "good" extending a human life represents.

Another extreme reaction is to condemn cost-benefit analysis either because of the difficulty of assigning dollar values to everything or simply on the grounds that it is a cold and heartless procedure.

The difficulty of assigning dollar values to life and health is important to recognize; however, it does not necessarily doom cost-benefit analysis. One must simply be aware of these limitations and try as well as one can to take them into account. The most important question about any ethical approach is whether it is desirable to *aspire* to. Just because a method is difficult to employ does not negate the value of the theory it is based on.

The objection to cost-benefit analysis that it is cold and heartless may come closer to focusing on a core question. Of course just to say that it would be "cold and heartless" (or something equivalent) is not sufficient. The critic may be saying one of two things. One possibility is that the critic is challenging the whole method of making bioethical decisions on the basis of rational principles and arguing that each decision is so distinct that other human faculties such as "intuition" must also play a role. We touched upon this critique of "principlism" earlier (see page 28). Another possibility is that the critic does not object to rational principles in general but only to the specific method of quantifying benefits and harms, a method which may yield a "false precision" that is so hopelessly inadequate that it is *not* even worth aspiring to. Or the critic may oppose the whole utilitarian approach to ethics and argue for rational principles that are not utilitarian.