30 million HIV infections and 7.8 million AIDS-related deaths during the same period. The number of pregnant HIV-positive women with access to ARV therapy has increased to 73 percent and new HIV infections among children have fallen by 58 percent.

Most of the progress has been made possible by an increase in the amount of money available to pay for HIV drugs and other interventions. In 1996, total global spending on HIV/AIDS was $300 million; by 2014, it had risen to roughly $22 billion. More than half (57 percent) of the $187 billion that has been invested in AIDS relief since 2000 has come from the domestic budgets of the countries most affected. But international aid programs have also been pivotal in building the momentum and infrastructure necessary to sustain effective HIV interventions. Since 2002, the Global Fund to Fight AIDS, Tuberculosis, and Malaria, for example, has disbursed more than $15.7 billion in more than 151 countries.

The U.S. government began funding the struggle against the AIDS pandemic in 1986, and in 2003 President George W. Bush initiated the President's Emergency Plan for AIDS Relief (PEPFAR). Since its founding, the PEPFAR program has committed over $65 billion to HIV/AIDS programs, to the Global Fund, and to bilateral tuberculosis initiatives. President Obama has since expanded PEPFAR, so that its funding, combined with private donors such as the Bill and Melinda Gates Foundation, make the United States the largest international contributor to HIV/AIDS relief in the world.

Money Well Spent?

Not all experts in global health policy believe that spending so much money on HIV/AIDS is the best strategy. It is not that they think the money is wasted or that millions of individuals don't benefit from treatment and prevention programs. Rather, they believe that the focus on HIV/AIDS has led the international community to ignore a range of health-related problems in developing countries that, if addressed, would produce greater reductions in mortality and morbidity.

Daniel Halperin, for example, points out that millions of African children and adults die of malnutrition, pneumonia, diarrhea, and many other common poverty-related conditions. Diarrhea, for example, is simple to treat, and most of the deaths it causes could be easily prevented. One-fifth of all global deaths from diarrhea occur in Congo, Ethiopia, and Nigeria. These three countries have relatively low levels of HIV infection, yet they receive hundreds of millions of dollars a year to support AIDS programs—money that might be better spent combatting conditions that cause childhood diarrhea.

Halperin's observations can be extended to other countries and circumstances. The majority of those in the underdeveloped world lack clean drinking water, yet contaminated water is the source of many infections. Money spent on sewage systems and clean water might make a bigger impact on the health of the population than a similar amount spent on HIV/AIDS. Further, the money from the Global Fund and other agencies sometimes exceeds the amount that can be put to use in a given country. Yet the funds cannot be used for any other purpose.

Given these circumstances, Halperin asks, "With 10 million children and a half million mothers in developing countries dying annually of largely preventable conditions, should we multiply AIDS spending, while giving out a pittance for initiatives like safe-water projects?"

The Promise of Prevention

An ounce of prevention is worth more than a pound of cure, and the most effective and definitive solutions to the HIV/AIDS pandemic still lie in efforts to stop the virus from spreading. Although traditional education measures that encourage condom use, abstinence, and monogamy have helped create double-digit decreases in new HIV infections, only slightly more than half (54 percent) of HIV-positive people know they have the virus. They help generate the unacceptably high number of new HIV infections each year—roughly 2 million in 2014.

The failure to bring HIV/AIDS under control has convinced many that the most effective means of curtailing the spread of the virus is likely to involve preventing the virus from causing infection after exposure. The ideal method for achieving this would be a safe, cheap, and effective vaccine, like those that brought smallpox, mumps, measles, whooping cough, and tetanus under control. Although promising steps have been taken toward developing a vaccine in recent years, none has yet emerged that is reliable and safe.

Another hope still unrealized is the development of a microbicide gel or cream that would kill HIV on contact. While limited as a preventive measure, it might at least reduce HIV infection in women, who suffer the majority of new HIV infections in the developing world. This approach would have the additional advantage of preventing the vertical spread of HIV: fewer babies would be born infected because their mothers are HIV-positive.

Perhaps the most promising option for stopping the spread of HIV is a worldwide program of intensive testing and treatment. In 2008, epidemiological researchers developed a mathematical model to predict what would happen if most adults and adolescents were tested every year for HIV and if those who tested positive were immediately treated with anti-retroviral drugs. The model predicted that the rate of HIV transmission could be reduced to such a low level that AIDS would virtually disappear within a decade. The reason is that those infected would, due to early treatment, carry such a low level of the virus that the chance of their passing it on to others would be significantly reduced. This phenomenon has likely already contributed to the significant drop in new HIV infections since 2000, but the large number (46 percent) of HIV-positive individuals who don't know their status suggests it may be difficult (and expensive) to achieve greater gains.

Regardless of which strategies are employed in the coming decades, two basic conclusions are accepted by most experts in public health. The first is that HIV/AIDS is a global catastrophe that continues to produce mass suffering and lost human potential around the world, especially in low- and medium-income nations. The second is that the richer nations, out of self-interest if not out of altruism or a sense of justice, must continue to invest the money and human capital that might finally bring the pandemic under control.

CASE PRESENTATION

Cutting and Culture

Like many other American teens, Leyla grew up connected to two radically different worlds. There was the Midwestern United States where she was born and raised, with its shopping malls, video games, and fast food. But there was also rural Somalia, with its camel herds and mud-walled houses, the place where her parents had grown up and where her extended family still lived.

Leyla was thrilled when her parents told her she would be travelling with them back to Somalia over her school vacation, and would finally meet her grandparents.

"Even though I'm one of four girls, I was the one they picked to go. I felt like I really was the lucky one," she recalled.

But soon after Leyla arrived in Somalia, she discovered another purpose of the visit. Her mother took her to a remote village and told her she was going to be "cut"—that her clitoris would be removed as part of a traditional practice opponents call *female genital mutilation* (FGM) and others call *female genital cutting* (FGC). Although she was terrified of what was about to happen to her, Leyla felt she had no choice but to cooperate when she was taken to the cutter's house.

"They had to hold me down," Leyla recalled. "There was no anaesthetic, no gloves, no pain medication after, no nurse to take care of you. It was the most painful thing I have experienced. They cut you like they are cutting paper. It's like you die."

Left bleeding in a side room with a wad of cotton in her underwear, Leyla struggled to reconcile what had just happened to her with the world she had grown up in.

"I felt like I was dreaming and that at some point I'd wake up. I felt so violated. Why didn't they just let me decide when I got older instead of ambushing me in the middle of nowhere?"

A Global Practice

Nearly a decade later, Leyla is still struggling with the physical and psychological effects of her ordeal. Although her experience is atypical for an American woman, it is strikingly common for women who live in some parts of Africa, the Middle East, and Asia.

The World Health Organization (WHO) estimates that as many as 30 million girls and women undergo FGM/C every year, most between infancy and age fifteen, and some 100 to 140 million are already living with its effects. In seven countries, including Egypt, Mali, and Guinea, more than 85 percent of women have undergone the procedure, and in Somalia, 98 percent have. The practice has also followed patterns of immigration, with the result that as many as five hundred thousand women in the United States have had or are at risk of having FGM/C, many through so-called "vacation cuttings" like Leyla's. Nevertheless, in the countries where it is practiced, as well as in the West, FGM/C has rarely been discussed in public, in part because understanding what it entails requires a frank discussion of female sexual anatomy.

The WHO has identified four types of FGM/C, although there are local variations among these categories. *Type I*, also called *clitoridectomy*, involves the total or partial removal of the clitoris, along with the clitoral hood. *Type II*, also called *excision*, typically involves clitoridectomy along with removal of the labia minora and/or the labia majora. Types I and II likely represent the majority of FGM/C, although there are strong regional differences in prevalence.

In Somalia as in several other North African countries, over 90 percent of FGM/C is *Type III*, also known as *infibulation*. Type III FGM/C involves the narrowing of the vaginal orifice by cutting and stitching together the labia, along with the removal of part or all of the other external genitalia. Practitioners often refer to this procedure as "sealing" girls, and indeed the catgut or wild-thorn stiches they put in generally leave only a pencil-sized hole to allow for urination and menstruation. Following infibulation, girls' legs are often bound together for ten days until the remaining vulvar tissue fuses together in a smooth surface, which is taken to be more attractively virginal than natural female anatomy. Upon marriage, these women must be "reopened" for sex, and men thus sometimes take a knife to their wedding nights.

Medical Risks

For most Westerners, not to mention many men and women from the societies where the practice is most prevalent, the details of FGM/C are deeply disturbing. They point out that female cutting almost always involves the infliction of excruciating pain on children. (This is true even in the small but growing number of cases where anesthesia is used in the initial procedure.) For this reason, FGM/C is often defined as a form of torture by international organizations as well as by several national constitutions, including Somalia's—although that war-torn country has yet to pass any specific laws against the practice.

The health risks of FGM/C have been well established by researchers around the world. It is not uncommon for girls to bleed to death during the procedure or to die of sepsis or tetanus shortly thereafter. Infibulation comes with particularly severe risks for women, including recurrent infections from backed-up urine and menses. All forms of FGM/C increase the risk of urinary incontinence, abscesses, HIV-transmission, and pain with sex, with many of these problems caused or aggravated by scar tissue. Reduced or deadened sexual sensation is a primary side effect—and often the primary objective—of FGM/C.

The practice also significantly increases childbirth complications. Studies have shown that women who have undergone some forms of FGM/C are 70 percent more likely to hemorrhage after giving birth and twice as likely to die in childbirth. Type III FGM/C has been shown to increase infant mortality by as much as 55 percent, often due to a neonate's head being crushed by scar tissue in the vaginal canal. No study has shown a medical benefit to FGM/C.

Cultural Imperialism?

Given this body of research, many Westerners view FGM/C as a backward and barbaric remnant of a tribal past, unrelated to modern medicine. But such assumptions ignore both the history of similar practices in the West and the increasing numbers of health care professionals who perform FGM/C around the world. According to estimates by the United Nations Population Fund, one out of five girls who undergo FGM/C are now cut by a trained medical professional, a proportion that appears to be on the rise. In Egypt, where roughly 90 percent of women have undergone some form of FGM/C, three out of four procedures now involve medical professionals.

Many of these practitioners claim that the health risks of FGM/C are exaggerated by critics and point to the continued popularity of male circumcision in the West and beyond as proof that cosmetic genital surgery can be legitimately performed on nonconsenting children. They argue that "female circumcision" is an equally acceptable expression of culture and religion and that Westerners who criticize it are thus engaging in hypocrisy and cultural imperialism. Proponents of the practice also typically make their own biomedical claims to justify it.

"One, it will stabilize her libido," said Lukman Hakim, chairman of a social service organization in Indonesia that performs hundreds of Type I cuttings every year. "Two, it will make a woman look more beautiful in the eyes of her husband. And three, it will balance her psychology." This mixed emphasis on female chastity, cultural standards of beauty or hygiene, and mental health benefits is common among FGM/C proponents and practitioners.

"This tradition is for keeping our girls chaste, for lowering the sex drive of our daughters," a Somali cutter named Maryan Hirsi Ibrahim told *New York Times* columnist Nicholas Kristof. "This is our culture."

Although such arguments may ring hollow to contemporary Western ears, they bear striking similarity to those that have been put forward for related practices in Western societies. As recently as the 1950s, clitoridectomies were still being performed in Europe and the United States to treat perceived pathologies such as hysteria, nymphomania, melancholia, homosexuality, and masturbation, as well as for diagnoses that are still part of clinical practice, such as epilepsy and schizophrenia. (In the West, however, cutting was not typically deemed necessary for "healthy" women.) Many defenders of FGM/C also point to the increasing popularity of breast augmentation and labial reduction surgeries as evidence that the West still believes that women's bodies must be "fixed" to make them socially acceptable. Although most of these nontherapeutic procedures are sought out by consenting adults, some involve younger patients, and defenders often argue that children in both sets of societies are eager to have such procedures performed. (See Case Presentation: Intersex Care in Chapter 10 for more on childhood cosmetic surgeries in the West.)

It is clear that nontherapeutic genital cutting is almost always bound up with cultural conceptions of gender, sexuality, and physical beauty. But with the exception of Judaism, there is little evidence that either male or female circumcision has a religious origin. Although FGM/C is more common among Muslim societies, there is no Islamic scriptural source for it and the practice predates both Christianity and Islam almost everywhere it is prevalent.

Finally, although both male circumcision and FGM/C appear to be primarily cultural in origin, comparison of the two practices often seems metaphorical and misleading. As many critics have pointed out, Types I and II FGM/C are generally more comparable to the removal of the entire penis and/or testes rather than just the foreskin. Although most male circumcisions are not performed for therapeutic reasons, the practice has been shown to have therapeutic benefits, with some studies showing a 60 percent reduction in HIV infection among circumcised men. No such benefits and significant harm has been shown to result from FGM/C. Although both practices raise questions about cosmetic surgeries on nonconsenting minors, FGM/C typically involves older children, less or no anesthesia, and more extensive damage to tissue with concentrated nerve endings. As described by Hibo

Wardere, an anti-FGM/C activist who was cut at age six in Somalia, the procedure involves "being engulfed in pain from head to toe—like fireworks going off everywhere and you don't know how to stop them. I prayed to God to just take me then and there."

Gender and Power

Unlike male circumcision, FGM/C is almost always defended in terms of the need for society to restrain the sexuality of those who are cut. (Clitoridectomies in the West were often framed in similar terms.)

"From a young age you were told girls who weren't cut were promiscuous," said Ayshah, a woman from Somalia who underwent FGM/C when she was five years old. "They were dirty, nobody wanted to marry them, they were ugly." Indeed, the most common defenses of the practice claim that it will ensure marital fidelity, preserve female virginity before marriage, and eliminate "masculine" traits in girls. These views are often shared by both men and women (who make up the majority of cutters), and some women, such as Ayshah, recall being eager to be cut—in part because of the gifts, sweets, and praise that typically follow the procedure.

It is clear, however, that FGM/C is almost always accompanied by deeply rooted structural inequalities between men and women. In societies where women are not allowed to work outside the home and cannot own or inherit land, girls are frequently viewed as an economic burden until they can elicit a dowry through marriage. In many societies in Africa and the Middle East, FGM/C is a rite of passage that marks the end of a girl's formal schooling and her eligibility for marriage—often to an older, financially secure man. Indeed, many critics suggest that FGM/C can only be discussed in the context of the seven hundred million women alive today who were forced into marriage as children. Furthermore, they suggest it reflects the same global devaluation of girls and women that produces lopsided sex-selection practices, neglect of girls' health, and more than 100 million females who are "missing" (relative to statistical norms) in Asia and Africa, because they failed to survive gestation, birth, or childhood. (See the Chapter 10 Briefing Session for more on this topic.) FGM/C is also strongly associated with patterns of sexual violence against women, "honor killings" of rape victims, and the global sex trafficking and forced prostitution that currently involves over four million women and girls.

Of course, these phenomena are hardly limited to societies or cultures that practice FGM/C, and many industrialized nations, including the United States, continue to struggle with disturbingly high levels of sexual violence, gender disparity, and human trafficking. Indeed, the persistence and global scale of gender-based inequality is so daunting that practices such as FGM/C can sometimes seem inevitable or unchangeable.

But the history of the twentieth century suggests that deeply entrenched cultural practices regarding gender may change with surprising speed. For ten centuries, generations of Chinese women had their feet systematically bound and broken to conform to an ideal of feminine beauty. But around the turn of the twentieth century, the practice fell out of favor and had almost completely disappeared in the span of a few decades. In the West, demands that women wear corsets or bonnets and not wear pants—cultural practices that were in place for generations—gave way as women gained greater social and financial autonomy. Medically unnecessary clitoridectomies and hysterectomies were also curtailed as women gained access to the medical profession and became active in patients' rights movements.

It is worth noting that none of these changes was accomplished through legal means, and some critics argue that prohibitions on FGM/C are counterproductive, driving the practice underground in a spirit of resistance to foreign norms. As long as the perceived social advantages of FGM/C remain as a necessary precondition to marriages and dowries, it will continue, they contend. Others argue that the lasting physical and psychological harm of FGM/C preclude any toleration of one of the world's most pervasive abuses of human rights. In line with such arguments, Nigeria banned FGM/C in May 2015, making it the twenty-third African nation to do so. (The practice has been illegal in the United States since 1996, but "vacation cuttings" such as Leyla's were not included in the law until 2012.)

One of the most successful campaigns against FGM/C has been led by the human rights group Tostan, which has focused on medical education campaigns in African villages. As a result, thousands of such villages have taken pledges to replace FGM/C with rituals involving "circumcision with words." Such rituals are designed to ensure that girls who have not been cut will still be marriageable and will not face ostracism or abuse. Still, groups like Tostan are competing with the "medicalization" of FGM/C by health professionals, who can earn extra money by performing the procedures and making unsubstantiated claims about their medical benefit. Thus, in Africa as in the rest of the world, medicine often serves as the stage for fierce debates over gender, power, and identity.

Leyla, for her part, emphasizes the potential benefits of a broad-based educational campaign in the United States and beyond, to educate physicians about how to treat FGM/C and to teach families about its risks.

"There needs to be more warning, more information. There needs to be a place where people can get help. My mother thought she was fulfilling her motherly duties," Leyla told a reporter from *The Guardian*. "That's why I'm telling my story. If I can stop this happening to just one more girl then it will have been worth it."

BRIEFING SESSION

In June of 2015, a commission of more than forty leading physicians, epidemiologists, and climatologists, organized by the *Lancet* medical journal, warned that climate change constituted a "medical emergency" that threatened to erase fifty years of progress in global public health. Based on sharp increases in heat stress, drought, flooding, violent storms, vector-borne diseases, population displacement, and food insecurity, "[c]limate change is the biggest global health threat of the 21st century," the commission concluded. In a separate report, the World Health Organization estimated that by 2030, at least 250,000 deaths a year will be directly attributable to climate change.

The serious health effects of our planet's changing climate have been on the horizon for nearly a century. In 1917, the inventor Alexander Graham Bell noted that the industrial revolution was contributing to "a sort of greenhouse effect," in which the earth's radiant heat was trapped by pollutants in the atmosphere. The continued burning of coal and oil, he claimed, was likely to produce a net result in which "the greenhouse becomes a sort of hot-house," raising average global temperatures and wreaking havoc on human societies.

Bell's warning has been followed by almost a century of dramatic temperature increase, violent shifts in regional climates, and the melting of the polar ice caps. Since comprehensive record keeping began in 1880, the Earth's average global temperature has risen 0.8 degrees Celsius, and NASA predicts that at current emissions rates, it will have risen between 2°C and 6°C by the end of the twenty-first century. Even a 2°C increase has not been seen on Earth for 125,000 years, and the speed of the current temperature increase is over twenty times faster than any past warming pattern after ice ages. Thirteen of the fifteen hottest years on record have occurred since 2000, with 2014 eclipsing both 2005 and 2010 for the warmest global average.

Climate change is already having a significant impact on human health, with unprecedented heat, drought, and extreme weather patterns killing thousands every year, and rising sea levels threatening the coastal living conditions of nearly half the world's population. Regional climate shifts are already increasing the range of diseases borne by vectors such as mosquitos and snails, and studies suggest that climate change will expose at least an additional two billion people to malaria and dengue fever before the end of the century. Many other changes, such as the impact on viral and bacterial mutations, are only beginning to be understood.

Destructive dynamics such as climate change do not respect national boundaries. As the 2014 Ebola outbreak and the 2010 H1N1 swine flu pandemic made clear, effective action for any single nation's public health is often intertwined with global health. The interdependence of national economies and cultures is paralleled by increasing interdependence of their populations' health status, a pattern that is mirrored by the global scope of our environmental problems and the moral

Defenders of this status quo argue that patients are, in effect, voting with their dollars, and that the higher salaries for urologists, for example, reflect their greater value to society. On this view, suffering and sickness in other societies might be an occasion for optional private charity, but it does not engender professional or moral obligations. Critics argue that biomedical markets are a particularly poor guide to social value. Prescribing antibiotics to people with dysentery will never be a revenue-generating enterprise, nor will treating destitute HIV patients with antiretroviral medications. But both interventions have enormous potential to alleviate suffering.

Others argue that we have a moral obligation to address global health disparities because we are materially involved in the medical conditions of foreigners. They contend that international financial institutions such as the World Bank and International Monetary Fund are structured to favor wealthy countries. Poor nations are, in part, poor because of the colonial and postcolonial theft of their natural resources and the long-standing suppression of human rights and human potential. On this view, the obligation to alleviate poverty-driven illness abroad is less like rescuing a drowning child than assisting a child one has hit with one's car.

Globalized Clinical Trials?

Issues of equity also arise with regard to the design and implementation of clinical trials. Health care disparities are so steep between poor and rich countries that many desperately ill people in the former effectively receive no treatment at all while most people in the latter receive high-quality (but expensive) treatment. Some researchers have argued that because many expensive treatments are not widely available in poor countries, trials in those countries need not provide standard-of-care treatments to control groups—as is generally required by the Declaration of Helsinki and other codes of medical research ethics. In a number of controversial studies in the 1990s, for example, researchers gave placebos to poor HIV-positive pregnant women, because their impoverished societies' "standard" treatment for avoiding maternal–fetal HIV transmission was typically no treatment at all. (See Social Context: Clinical Trials in the Developing World.)

Defenders of such trials argue that they made some subjects better off and no subjects worse off than they would otherwise have been. They also claim that if the results are used to benefit the broader population, then it is clearly better that such trials occur rather than not occur. But critics suggest that the real question is not whether such studies should occur or not, but why offering First-World standards of care in Third-World research contexts is deemed "unrealistic" or "unfeasible" in the first place. If such current best treatments are, in fact, medically and scientifically inappropriate for destitute patients in poor countries, then lower-standard studies may be justified. (This position is similar to that adopted in the most recent version of the Declaration of Helsinki, which is reprinted in Chapter 2.) But if the decision has been influenced by financial considerations (either in denying poor people expensive treatments or in selecting a poor country as a cheaper place to conduct research), then subjects are being exploited. Critics argue that ethical research must not blindly accept and perpetuate immoral health care disparities between rich and poor nations. Defenders argue that such disparities are the responsibility of government and aid organizations, and that biomedical research is ill suited to address such problems.

Relativism and Pluralism

A persistent challenge for any global approach to bioethics is that different societies and cultures appear to disagree over specific biomedical practices and policies, and their members often cite different or even incompatible values to justify these positions. Some societies prohibit abortion, homosexuality, and birth control, while others legalize euthanasia, recreational drug use, and prostitution. Some societies tolerate female genital cutting, while others restrict both male and female circumcision to varying degrees. Approaches to informed consent and patient autonomy vary widely across societies, with some cultures in Asia and Africa, for example, emphasizing a more collective approach to medical decision-making that involves family or community groups.

In the Anglo-American tradition, bioethics developed with an emphasis on the application of universally binding ethical principles such as nonmaleficence and informed consent, with much of the debate revolving around the scope of their application. While this analytic approach, sometimes known as *principalism*, is both efficient and powerful, it has often failed to engage directly with biomedical concepts from non-Western cultures. When such issues do arise, they have often been framed as a conflict between cultural relativism and objectivism. *Cultural relativism* is the view that there are no universally objective principles or values beyond the beliefs and customs of specific cultures, and thus there can be no basis for cross-cultural criticism or debate. By contrast, ethical *objectivism* holds that there are a set of universal and objective values and principles that have application everywhere.

As many critics have argued, however, the choice between these two positions can be artificial and misleading. Cultural relativism in ethics, critics argue, is a notoriously ambiguous and possibly self-contradictory view that few people actually hold. The cultural relativist's most basic claim—that cultures have no right to judge and should therefore tolerate each other—is itself the type of cross-cultural evaluation that relativism forbids. Indeed, cultural relativism appears to have no ethical grounds for criticizing aggressive intolerance between or within cultures. So long as a culture's dominant or majority view endorses slavery or involuntary euthanasia, for example, the relativist must consider such practices "right for them" in that culture. Also, critics say, there are notorious problems assessing the boundaries of any given culture and formulating its views, and most cultures (and individuals) are in fact *multicultural*, shaped by many different subcultural and foreign influences.

These long-recognized problems with relativism do not, however, require that we endorse a view of ethics as consisting of a set of universally objective "moral facts" or principles that can be applied everywhere without regard to cultural difference. As many ethical theorists have pointed out, our everyday use of moral language mostly involves the practice of persuasion and dialogue, rather than the application of universal concepts and principles. We use the language of values and principles primarily to try to persuade others and to shape their thoughts, emotions, and actions—not to make objective statements about the world. Based on this perspective, other approaches to the problem of moral diversity have been developed that avoid the extreme versions of both objectivism and relativism.

Some of these approaches seek to identify a rough set of overlapping human norms and values that are shared across cultures but expressed differently in different societies. Finding such overlap in global bioethics involves careful attention to the specific cultural and religious values that shape disputed practices or policies. Take the expectation in some Asian and African cultures that competent patients may delegate their informed consent to family members or community leaders. Rather than simply imposing 'Western'

principles of truth-telling and patient autonomy on these situations, an alternative approach would take seriously the collectivist values that may underlie such delegations of responsibility. Recognizing the importance of familial and communal affiliation to patients' well-being may challenge us to expand traditional conceptions of patient autonomy, while also being mindful of the risks of cultural coercion and undue influence. (Michelle Gold's article in Chapter 1 provides an example of this type of moral reflection.) The premise of such cross-cultural exchange is that there is enough overlap in the moral vocabularies of different societies to engage in meaningful bioethical dialogue. Although this process is unlikely to resolve all sources of disagreement, the goal is to identify some basic values that are central to human flourishing and human dignity, even as societies have come to express them in substantially different ways.

A key feature of many newer approaches to developing a global bioethics is that they are explicitly *pluralistic*. They assume that different people and cultures will have different conceptions of the good and that their biomedical practices will be shaped by different views on health, morality, and religion. For example, unlike utilitarian or most Kantian approaches to moral judgment, which reduce complex situations to a single metric of value (e.g., the promotion of utility or application of the categorical imperative), pluralism avoids commitment to a single comprehensive view. Instead, pluralistic approaches seek to find overlapping consensus on some basic principles and policies that different groups can endorse even though they hold radically different comprehensive views. According to John Rawls, such approaches are "political" in that they seek moral agreement for pragmatic purposes, rather than to abolish religious and philosophical disagreements.

A pluralistic approach to evaluating the moral legitimacy of female genital cutting, for example, could not simply rely on the notion, associated with contemporary Western ethics, that individuals are autonomous (self-determining) and so it is wrong to perform surgery on them without their consent. Relying on such a notion would run counter to the idea in many traditional societies that individuals are not, in themselves, fully autonomous, but gain status as a person only as part of a society or in relationship to divinity. The pluralistic approach might instead endorse a broader human capability, such as the ability to secure one's bodily integrity, to challenge practices such as genital cutting.

As such, this view might be endorsed both by those who value bodily integrity as a way of honoring God's creation and those who value it based on Kant's conception of autonomous control over one's own body. Such overlapping consensus may not incorporate all viewpoints or convince all supporters of FGM/C, but pluralists argue that it usefully allows room for local variations on broadly shared conceptions of human dignity and human rights.

It is worth noting that some of the most successful campaigns against FGM/C appear to have employed strongly pluralistic approaches to the issue. The Senegal-based organization Tostan has successfully brokered hundreds of intervillage agreements in Africa to replace FGM/C with a "circumcision by words"—a ceremony that preserves marriageability and a cultural rite of passage for girls who receive it. Drawing on local community values, such educational campaigns seek to align diverse religious and cultural perspectives around a pragmatic shift in biomedical practice. (See Case Presentation: Cutting and Culture.)

Pluralistic approaches to bioethics are rooted in the view that moral community—both within and among societies—admits of multiple conceptions of the good. In this, they follow the work of W. D. Ross and other theorists going back to Aristotle who see social goods as varied and incommensurable. The difficulty for all such views is that they provide little guidance when we face conflicts among recognized values and the duties they engender. Critics argue that they provide insufficient moral clarity to resolve the tragic biomedical choices posed by such phenomena as HIV, global health disparities, and climate change. Defenders of a pluralistic approach to bioethics argue that it provides the only appropriate and effective framework through which the world's diverse peoples can take collective action to address the urgent problems we face.

READINGS

Section 1: Global Health and Distributive Justice

The Distribution of Biomedical Research Resources and International Justice

David B. Resnik

Resnik cites research showing that only 10 percent of the world's biomedical research funding is devoted to the problems that cause 90 percent of the world's disease burden, most of it shouldered by the developing world. Resnik attributes this disparity to two main causes: (1) research on diseases that afflict the developing world (such as malaria and TB) is not viewed as financially rewarding by corporations; (2) such research is not viewed as politically rewarding by governments, in part because it is not promoted by advocacy groups as cancer and heart disease research is.

Nevertheless, the 90/10 divide is profoundly unjust, Resnik claims. He argues for a cosmopolitan version of distributive justice, in which we are not excused from obligations to the least advantaged just because they happen to live in a different nation. In order to address this injustice, Resnik calls for (1) advocacy groups to lobby corporations and governments for developing-world diseases as they do for developed-world ones; and (2) a UN-administered trust to help fund and patent biomedical R&D for the 90 percent.

The distribution of resources for research and development (R&D) in biomedicine has a direct impact on the progress of the health sciences and the distribution of health. Money spent on biomedical R&D can have a positive impact on disease and disability, longevity, infant mortality, and other measures of the health of a population. Differences in research funding can also be a factor in racial and ethnic health disparities. For many years there have been racial and ethnic disparities in health in the United States (US). While many different factors, such as poverty, discrimination, and educational attainment contribute to these disparities, inadequate

From David B. Resnik, "The Distribution of Biomedical Research Resources and International Justice," *Developing World Bioethics* 4, no. 1 (2004). ISSN 1471-8731 (print); 1471–8847 (online). (Most references omitted.)

Chapter 12

Ethical Theories, Moral Principles, and Medical Decisions

CHAPTER CONTENTS

BASIC ETHICAL THEORIES 902
- Utilitarianism 903
  - The Principle of Utility 903
  - Act and Rule Utilitarianism 905
  - Preference Utilitarianism 909
  - Difficulties with Utilitarianism 910
- Kant's Ethics 911
  - The Categorical Imperative 911
  - Another Formulation 912
  - Duty 912
  - Kant's Ethics in the Medical Context 913
  - Difficulties with Kantian Ethics 915
- Ross's Ethics 916
  - Moral Properties and Rules 916
  - Actual Duties and Prima Facie Duties 917
  - Ross's Ethics in the Medical Context 919
  - Difficulties with Ross's Moral Rules 920
- Rawls's Theory of Justice 921
  - The Original Position and the Principles of Justice 921
  - Rawls's Theory of Justice in the Medical Context 923
  - Difficulties with Rawls's Theory 925
- Natural Law Ethics and Moral Theology 926
  - Purposes, Reason, and the Moral Law as Interpreted by Roman Catholicism 926
  - Applications of Roman Catholic Moral–Theological Viewpoints in the Medical Context 929
  - Difficulties with Natural Law Ethics and Moral Theology 931

MAJOR MORAL PRINCIPLES 932
- The Principle of Nonmaleficence 933
- The Principle of Beneficence 935
- The Principle of Utility 936
- Principles of Distributive Justice 938
  - The Principle of Equality 939
  - The Principle of Need 939
  - The Principle of Contribution 940
  - The Principle of Effort 940
- The Principle of Autonomy 941
  - Autonomy and Actions 942
  - Autonomy and Options 943
  - Autonomy and Decision-Making 943
  - Restrictions on Autonomy 944

BEYOND PRINCIPLISM 946
- The Capabilities Approach 947
  - Central Capabilities 948
  - The Capabilities Approach in the Medical Context 949
  - Difficulties with the Capabilities Approach 950
- Virtue Ethics 950
  - The Virtues 951
  - Virtue Ethics in the Medical Context 952
  - Difficulties with Virtue Ethics 952
- Care Ethics 953
  - Values, Not Principles 954
  - Care Ethics in the Medical Context 955
  - Difficulties with Care Ethics 956
- Feminist Ethics and Critical Theories of Identity 957
  - Critical Theories of Identity in a Medical Context 958
  - Difficulties with Identity-Based Ethical Theories 959

RETROSPECT 960

"He's stopped breathing, Doctor," the nurse said. She sounded calm and deliberate. By the time Dr. Sarah Cunningham had reached Matteo Sabatini's bedside, the nurse was already providing cardiopulmonary resuscitation. But Sabatini still had the purplish blue color of cyanosis, caused by a lack of oxygen in his blood.

Sarah Cunningham knew that if Sabatini was to survive, he would have to be given oxygen fast and placed on a respirator. But should she order this done?

Matteo Sabatini was an old man, almost ninety. So far as anyone knew, he was alone in the world. His health was poor. He had congestive heart disease and was dying slowly and painfully from intestinal cancer. Wouldn't it be a kindness to Sabatini to allow him this quick and painless death? Why condemn him to lingering on for a few extra hours or weeks?

The decision that Sarah Cunningham faces is a moral one. She has to decide whether she should order the medical procedures that might prolong Sabatini's life or not order them and accept the consequence that he will almost surely die within minutes.

This kind of case rivets our attention because of its immediacy and drama. But there are many other situations that arise in the context of medical practice and research that require moral decisions. Some are equal in drama to the problem facing Dr. Cunningham, while others are not so dramatic but are of at least equal seriousness. There are far too many to catalogue, but consider this sample: Should physicians ever lie to their patients? Should parents be able to allow their children to be used as experimental subjects? Should people suffering from genetic diseases get access to assisted reproduction? Should women's access to abortion be restricted in certain cases? Should children with serious birth defects be allowed to die? Do terminally ill people have a right to physician-assisted death? Does everyone have a right to medical care?

Such questions are likely to strike some of us as overly abstract or academic. This attitude often changes, however, when we find ourselves in a position to make or influence biomedical decisions that impact specific individuals. It changes, too, when we find ourselves on the receiving end of such decisions.

But whether we view these problems abstractly or concretely, they generate the same question: Are there any rules, standards, or principles that we can use as guides when we are faced with moral decisions? If there are, then Dr. Cunningham need not be wholly unprepared to decide whether she should order steps taken to save Matteo Sabatini's life. Nor need we be unprepared to decide issues like those discussed above.

The branch of philosophy concerned with principles that allow us to make decisions about what is right and wrong is called *ethics* or *moral philosophy*. *Bioethics* is specifically concerned with moral principles and decisions in the context of medical practice, policy, and research. Bioethics has a specialized focus but remains a part of the discipline of ethics. Thus, if we are to answer our question as to whether there are any rules or principles to use when making moral decisions in the medical context, we must turn to general ethical theories and to a consideration

of moral principles that have been proposed to hold in all contexts of human action.

In the first section, we will discuss five major ethical theories that have been put forward by philosophers. Each of these theories represents an attempt to supply basic principles we can rely on in making moral decisions. We'll consider these theories and examine how they might be applied to moral issues in the medical context. We will discuss the reasons that have been offered to persuade us to accept each theory, but we will also point out some of the difficulties each theory presents.

In the second section, we will examine and illustrate several moral principles that are of special relevance to biomedical research and practice. These principles are frequently appealed to in discussions of practical ethical problems and are sufficiently uncontroversial to be endorsed in a general way by any of the ethical theories mentioned in the first section. (The application of these principles independent of traditional moral theories is sometimes called *principlism*.)

In the third and final section, we will consider approaches to ethical theory that are not identified with a specific set of foundational principles. These include the capabilities approach, virtue ethics, care ethics, and various identity-based theories. We will consider how these theories may be used in making moral decisions, but we will also call attention to some of the criticisms urged against them.

The three sections are not dependent on one another, and it is possible to profit from one without reading the others. (The price for this independence is a small amount of repetition.) Nevertheless, reading all three sections is recommended. Readers will find it easier to follow and evaluate many of the book's Briefing Sessions, as well as Social Context and Case Presentations, if they have at least some familiarity with these various approaches to ethical theory and moral principles.

BASIC ETHICAL THEORIES

Ethical theories attempt to articulate and justify principles that can be employed as guides for moral decision-making and as standards for the evaluation of actions and policies. In effect, such theories attempt to define what it means to act morally, and in doing so, they stipulate in a general fashion our duties or obligations.

Ethical theories also offer a means to explain and justify actions. If our actions are guided by a particular theory, then we can explain them by demonstrating that the principles of the theory required us to act as we did. In such cases, the explanation also constitutes a justification. (In some cases, we may justify our actions by showing that the theory *permitted* our actions—that is, didn't require them, but didn't rule them out as wrong.)

Advocates of a particular ethical theory present what they consider to be good reasons and relevant evidence in its support. Frequently, their aim is to show that the theory is one that any reasonable individual would find persuasive or would endorse as correct. Accordingly, appeals to faith, ideology, or other comprehensive views of the world are not considered to be either necessary or legitimate to justify the theory. Rational persuasion alone is regarded as the basis of justification.

In this section, we will briefly consider four general ethical theories and one theory of justice that has an essential ethical component. In each case, we will begin by examining the basic principles of the theory and the grounds offered for its acceptance. We will then explore some possible applications of the theory that arise within the medical context. Finally, we will mention some of the practical consequences and conceptual difficulties that raise questions about the theory's adequacy or correctness.

Utilitarianism

The ethical theory known as utilitarianism was given its most influential formulation in the nineteenth century by the British philosophers Jeremy Bentham (1748–1832) and John Stuart Mill (1806–1873). Bentham and Mill did not produce identical theories, but both of their versions have become known as "classical utilitarianism." Subsequent elaborations and qualifications of utilitarianism usually draw on the formulations of Bentham and Mill, so their theories are worth careful examination.

The Principle of Utility

The foundation of utilitarianism is a single apparently simple principle. Mill calls it the "principle of utility" and states it this way: *Actions are right in proportion as they tend to promote happiness, wrong as they tend to produce the reverse of happiness.*

The principle focuses attention on the *consequences* of actions, rather than on some feature of the actions themselves. The *utility* or "usefulness" of an action is determined by the extent to which it produces happiness. Thus, no action is *in itself* right or wrong. Nor is an action right or wrong by virtue of the actor's hopes, intentions, or past actions. Consequences alone are important. Breaking a promise, lying, causing pain, or even killing a person may, under certain circumstances, be the right action to take. Under other circumstances, the action might be wrong.

We need not think of the principle as applying to just one action that we are considering. It supplies the basis for a kind of cost–benefit analysis to employ in a situation in which several lines of action are possible. Using the principle, we are supposed to consider the possible results of each action. Then we are to choose the one that produces the most benefit (happiness) at the least cost (unhappiness). The action we take may produce some unhappiness, but it is a balance of happiness over unhappiness that the principle tells us to seek.

Suppose, for example, that a patient in a large hospital is near death: she is in a coma, an EEG shows only minimal brain function, and a respirator is required to keep her breathing. Another patient has just been brought to the hospital from the scene of an automobile accident. His kidneys have been severely damaged, and he is in need of an immediate transplant. There is a good tissue match with the first patient's kidneys. Is it right to hasten her death by removing one of them for transplant?

If we view this scenario in isolation, the principle of utility would likely prompt us to view the removal of the kidney as justified. The woman is on the brink of death, with little-to-no chance of recovery, while the man has a good chance of surviving if he receives the kidney. It is true that the woman's life is threatened even more by the surgery. It may in fact kill her. But, on balance, the kidney transplant seems likely to produce more happiness than unhappiness. In fact, it seems better than the alternative of doing nothing. For in that case, both patients are likely to die.

The principle of utility is also called the "greatest happiness principle" by Bentham and Mill. The reason for this name is clear when the principle is stated as follows: *Those actions are right which produce the greatest happiness for the greatest number of people.* This alternative formulation makes it obvious that, in deciding how to act, it is not just my happiness or the happiness of a particular person or group that must be considered. According to utilitarianism, every person is to count just as much as any other person. That is, when we are considering how we should act, everyone's interest must be considered. The right action, then, will be the one that produces the most happiness for the largest number of people.

Mill is particularly anxious that utilitarianism not be construed as a sophisticated justification for crude self-interest. He stresses that, in making a moral decision, we must look at the situation objectively. We must, he says, be a "benevolent spectator" and then act in a way that will bring about the best results for all concerned. This view is summarized in a famous passage:

> *The happiness which forms the utilitarian standard of what is right in conduct is not the agent's own happiness, but that of all concerned. As between his own happiness and that of others, utilitarianism requires him to be as strictly impartial as a disinterested and benevolent spectator. In the golden rule of Jesus of Nazareth, we read the complete spirit of the ethics of utility. To do as you would be done by, and to love your neighbor as yourself, constitute the ideal perfection of utilitarian morality.*

The key concept in both formulations of the principle of utility is "happiness." Bentham simply identifies happiness with pleasure—pleasure of any kind. The aim of ethics, then, is to increase the amount of pleasure in the world to the greatest possible extent. In furtherance of this aim, Bentham recommends the use of a "calculus of pleasure and pain," in which characteristics of pleasure such as intensity, duration, and number of people affected are measured and assigned numerical values. To determine which of several possible actions is the right one, we need only determine which one receives the highest numerical score. Unfortunately, Bentham does not tell us what units to use or how to make the measurements.

Mill also identifies happiness with pleasure, but he differs from Bentham in a major respect. Unlike Bentham, he insists that some pleasures are "higher" than others. Thus, pleasures of the intellect are superior to, say, purely sensual pleasures. This difference in the concept of pleasure can become significant in a medical context. For example, Mill's view might assign reduced value to the life of a dementia patient with severely diminished cognition, but who still gets intense pleasure from eating ice cream. The loss of "higher pleasures," such as reading or meaningful conversation with loved ones, cannot be compensated by any quantity of pleasure derived from ice cream. (For Bentham, by contrast, pleasure is pleasure, no matter what the source.)

Both Mill and Bentham regard happiness as an intrinsic good. That is, it is something good in itself or for its own sake. Actions, by contrast, are good only to the extent to which they tend to promote happiness. Therefore, they are only *instrumentally* good. Since utilitarianism determines the rightness of actions in terms of their tendency to promote the greatest happiness for the greatest number, it is considered to be a *teleological* ethical theory. (*Teleological* comes from the Greek word *telos*, which means "end" or "goal.") A teleological ethical theory judges the rightness of an action in terms of an external goal or purpose—"general happiness" or utility for utilitarianism. However, utilitarianism is also a *consequentialist* theory, for the outcomes or consequences of actions are the only considerations relevant to determining their moral rightness. Not all teleological theories are consequentialist.

Some more recent formulations of utilitarianism have rejected the notion that happiness, no matter how defined, is the sole intrinsic good that actions or policies must promote. Critics of the classical view have argued that the list of things we recognize as valuable in themselves should be increased to include ones such as knowledge, beauty, love, friendship, liberty, and health. According to this *pluralistic* view, in applying the principle of utility we must consider the entire range of intrinsic goods that an action is likely to promote. Thus, the right action is the one that can be expected to produce the greatest sum of intrinsic goods. In most of the discussion that follows, we will speak of the greatest happiness or benefit, but it is easy enough to see how the same points can be made from a pluralistic perspective.

Act and Rule Utilitarianism

Utilitarians generally accept the principle of utility as the standard for determining the rightness of actions. But they divide into two major groups over how to apply the principle. *Act utilitarianism* holds that the principle should be applied to particular acts in particular circumstances. *Rule utilitarianism* maintains that the principle should be used to test rules, which can in turn be used to decide the rightness of particular acts. Let us consider each of these views and see how it works in practice.

Act utilitarianism holds that an act is right if, and only if, no other act could have been performed that would produce a higher utility. Suppose an infant is born with severe impairments. The child has an open spine, severe brain damage, and dysfunctional kidneys. What should be done? (We will leave open the question of who should decide.)

The act utilitarian holds that we must attempt to determine the consequences of the various actions that are open to us. We should consider, for example, these possibilities: (1) Give the child only the ordinary treatment that would be given to a healthy infant; (2) give the child special treatment for his problems; (3) give the child no treatment—allow him to die; (4) actively end the child's life in a painless way.

According to act utilitarianism, we must explore the potential results of each possibility. We must realize, for example, that when such a child is given only ordinary treatment, he will be worse off, if he survives, than if he had received special treatment. Also, an impaired child allowed to die may suffer more pain for a longer period of time than one killed by a lethal injection. Furthermore, a child treated aggressively will have to undergo numerous surgical procedures of limited effectiveness. We must also consider the child's family and assess the emotional and financial effects that each of the possible actions will have on them. Then, too, we must take into account such matters as the "quality of life" of a child with severe brain damage and multiple impairments, the effect on physicians and nurses in killing or allowing the child to die, and the financial costs to society in providing long-term care.

After these considerations, we should then choose the action that has the greatest utility. We should act in the way that will produce the most benefit for all concerned. Which of the possibilities we select will depend on the precise features of the situation: how impaired the child is, how good his chances are for living a tolerable life, the values and financial status of the family, and so on. The great strength of act utilitarianism is that it invites us to approach each case as unique. In another case with different circumstances, we might, without inconsistency, choose another of the possible actions.

Act utilitarianism shows a sensitivity to specific cases, but it is not free from difficulties. Some philosophers have pointed out that there is no way that we can be sure that we have chosen the right action. We are sure to be ignorant of much relevant information. Besides, we can't predict with much confidence what the results of our actions will be. There is no way to be sure, for example, that even a severely impaired infant will not recover sufficiently to live a better life than we predict.

The act utilitarian can reply that acting morally doesn't require omniscience. We must only make a reasonable effort to obtain relevant information, which we can use to determine the likely consequences of our actions. Acting morally doesn't require anything more than this.

Another objection to act utilitarianism is more serious. According to the doctrine, I am obligated to keep a promise only if keeping it will produce more utility than some other

action. If some other action will produce the same utility, then keeping the promise is permissible but not obligatory. Suppose a surgeon promises a patient that only she will perform an operation, then allows a well-qualified resident to perform part of it. Suppose all goes well and the patient never discovers that the promise was not kept. The outcome for the patient is exactly the same as if the surgeon had kept the promise. From the point of view of act utilitarianism, there is nothing wrong with the surgeon's failure to keep it. Yet critics charge that there is something wrong—that, in making the promise, the surgeon took on an obligation. Act utilitarianism is unable to account for obligations incurred by such actions as promising and pledging, critics say, for such actions involve something other than consequences.

A third objection to act utilitarianism arises in situations in which nearly everyone must follow the same rules in order to achieve a high level of utility, but even greater utility can be achieved if a few people disregard the rules. Consider the relationship between physicians and the Medicaid insurance program. The program pays physicians for services provided to patients with very low incomes. The program would collapse if most physicians were not honest in billing Medicaid for their services. Not only would many poor people suffer, but many physicians would lose a source of income.

Suppose a particular physician believes that the requirements to qualify for Medicaid are too restrictive and that many who urgently need medical care are denied coverage. As an act utilitarian, she reasons that it is right for her to apply for funding under Medicaid to open a free clinic. She intends to bill for services she does not provide, then use that money to treat those not covered by Medicaid. Her claims will be small compared to the entire Medicaid budget, so it is unlikely that anyone who qualifies for Medicaid will go without treatment. Since she will tell no one what she is doing, others are not likely to be influenced by her example and make false claims for similar or less altruistic purposes. The money she is paid will bring substantial benefit to those in need of health care. Thus, she concludes, by violating the rules of the program, her actions will produce greater utility than would be produced by following the rules.

The physician's action would be morally right, according to many act utilitarians. Yet, critics say, we expect an action that is morally right to be one that is right for everyone in similar circumstances. If every physician in the Medicaid program acted in this way, however, the program would be destroyed and thus produce no utility at all. Furthermore, according to critics, the physician's action produces unfairness. Although it is true that the patients she treats at her free clinic gain a benefit they would not otherwise have, similar patients elsewhere must go without treatment. The Medicaid policy, whatever its flaws, is at least prima facie fair in providing benefits to all who meet its requirements. Once again, according to critics, more seems to be involved in judging the moral worth of an action than can be accounted for by act utilitarianism.

In connection with such objections, some critics have gone so far as to claim that it is impossible to see how a society in which everyone was an act utilitarian could function. We could not count on promises being kept nor take for granted that people were telling us the truth. Social policies would be no more than general guides to action, and we could never be sure that people would regard themselves as obligated to adhere to the provisions of those policies. The critics are not necessarily right, of course, and defenders of act utilitarianism have made substantial efforts to answer such criticisms. Some have denied that the theory has those implications; others have argued that some of our common moral perceptions should give way to more rational

standards. For example, concerning euthanasia, Carl Wellman writes,

> *Try as I may, I honestly cannot discover great hidden disutilities in the act of killing an elderly person suffering greatly from an incurable illness, provided that certain safeguards like a written medical opinion by at least two doctors and a request by the patient are preserved. In this case I cannot find any way to reconcile my theory with my moral judgment. What I do in this case is to hold fast to act-utilitarianism and distrust my moral sense. I claim that my condemnation of such acts is an irrational disapproval, a condemnation that will change upon further reasoning about the act.... That I feel wrongness is clear, but I cannot state to myself any rational justification for my feeling. Hence, I discount this particular judgment as irrational.*

Rule utilitarianism maintains that an action is right if it conforms to a rule of conduct that, according to the principle of utility, will produce at least as much utility as any other rule applicable to the situation. A rule like "Provide only ordinary care for severely brain-damaged newborns with multiple impairments" would, if established, allow us to decide about the course of action to follow in situations like the one described above.

The rule utilitarian is concerned with assessing the utility not of individual actions but of particular rules. In practice, then, we do not have to go through the calculations involved in determining in each case whether a specific action will increase utility. All that we have to establish is that following a certain rule will, in general, result in a situation in which utility is maximized. Once rules are established, they can be relied on to determine whether a particular action is right.

The basic idea behind rule utilitarianism is that having a set of rules that are consistently observed produces the greatest social utility. Having everyone follow the same rule in each case of the same kind yields more utility for everyone in the long run. An act utilitarian can agree that having rules may produce more social utility than not having them. But the act utilitarian insists that the rules be regarded as no more than rough guides to action, as "rules of thumb." Thus, for act utilitarianism it is perfectly legitimate to violate a rule if doing so will maximize utility in that instance. By contrast, the rule utilitarian holds that rules must be followed, even though following them may produce less net utility (more unhappiness than happiness) in a particular case.

Rule utilitarianism can endorse rules such as "Keep your promises." Thus, unlike act utilitarianism, it can account for the general sense that, in making promises, we are placing ourselves under an obligation that cannot be set aside for the sake of increasing utility. If "Keep your promises" is accepted as a rule, then the surgeon who fails to perform all of an operation herself when she has promised her patient she would do so has not done the right thing, even if the patient never learns the truth.

Rule utilitarians recognize that circumstances can arise in which it would be disastrous to follow a general rule, even when it is true that, *in general,* greater happiness would result from following the rule all the time. Clearly, we should not keep a promise to meet someone for lunch when we have to choose between keeping the promise and rushing a heart-attack victim to the hospital. It is consistent with the theory to formulate rules that include appropriate escape clauses. For example, "Keep your promises, unless breaking them is required to save a life" and "Keep your promises, unless keeping them would lead to a disastrous result unforeseen at the time the promise was made" are rules that a rule utilitarian might regard as likely to produce greater utility than "Always keep your promises no matter what the consequences may be." What a rule utilitarian cannot endorse is a rule like "Keep your promises, except when breaking a promise would produce more utility."

This would cause the rule utilitarian's position to collapse back into act utilitarianism.

Of course, rule utilitarians are not committed to endorsing general rules only. It is compatible with the view to offer quite specific rules, and in fact there is no constraint on how specific a rule may be. A rule utilitarian might, for example, establish a rule such as, "If an infant is born with an open spine, severe brain damage, and dysfunctional kidneys, then the infant should receive no life-sustaining treatment."

The possibility of formulating a large number of rules and establishing them separately opens this basic version of rule utilitarianism to two objections. First, some rules are likely to conflict when they are applied to the same case and basic rule utilitarianism offers no way to resolve such conflicts. What should a physician do when faced both with a rule like that above and with another that directs him to "Provide life-sustaining care to all who require it"? Rules that pass the test of promoting utility when considered individually may express contradictory demands. A further objection to basic rule utilitarianism is that establishing rules to cover many different circumstances and situations results in such an abundance of rules that employing them to make moral decisions becomes difficult, if not impossible.

Partly due to such difficulties, some rule utilitarians have instead tried to establish the utility of a fixed set of rules or an entire moral code. The set can include rules for resolving potential conflicts, and an effort can be made to keep the rules few and simple to minimize the practical difficulty of employing them. Once again, as with basic rule utilitarianism, the principle of utility is employed to determine which set of rules, out of the various sets considered, ought to be accepted.

In this more sophisticated form, rule utilitarianism can be characterized as the theory that an action is right when it conforms to a set of rules that has been determined to produce at least as much overall utility as any other set. It is possible to accept as constraints certain social and economic institutions, such as private property and a market economy, and then argue for the set of rules that will yield the most utility under those conditions. However, it is also possible to argue for a different set of rules that would lead to the greatest possible utility, quite apart from those social forms. Indeed, such a set might be defended in an effort to bring about changes in present society that are needed to increase the overall level of utility. Utilitarianism, whether act or rule, need not be limited to individual moral obligation. It is also a social and political theory.

We have already seen that rule utilitarianism, unlike act utilitarianism, makes possible the sort of obligation we associate with making promises. But how might rule utilitarianism deal with the case of the physician who files false Medicaid claims to raise funds to operate a free clinic? An obvious answer, although certainly not the only one possible, is that any set of rules likely to be adopted by a rule utilitarian will contain at least one rule making fraud morally wrong. Without a rule forbidding fraud, no social program that requires the cooperation of its participants is likely to achieve its aim. Such a rule protects the program from miscalculations of utility that individuals may make for self-serving reasons, keeps the program focused on its goals, and prevents it from becoming fragmented. Even if a few individuals commit fraud, the rule against it is crucial in discouraging a large majority from committing it. Otherwise, as we pointed out earlier, such a program would collapse. By requiring that the program operate as it was designed, rule utilitarianism also preserves prima facie fairness, because only those who qualify receive benefits.

The most telling objection to rule utilitarianism, according to some critics, is that it is inconsistent. The justification for a set of moral rules is that the rules maximize utility. If rules are to maximize utility, then it seems obvious that they may require that an act produce more utility than any other possible act in a particular situation. Otherwise, the maximum amount of utility would not result. But if the rules satisfy this demand, then they will justify exactly the same actions as act utilitarianism. Thus, the rules will deem it permissible to sometimes break promises, make fraudulent claims, and so on. When rule utilitarianism moves to block these possibilities by requiring that rules produce only the most utility overall, it becomes inconsistent: the set of rules is said to maximize utility, but the rules will require actions that do not maximize utility. Thus, rule utilitarianism seems both to accept and reject the principle of utility as the ultimate moral standard.

Preference Utilitarianism

Some critics have questioned the wisdom of using happiness or any other intrinsic value (e.g., knowledge or health) as a criterion of the rightness of an action. The notion of an intrinsic value, they have argued, is too imprecise to be used as a practical guide. Furthermore, it is not at all clear that people share the same values, and even if they do, they are not committed to them to the same degree. One person may value knowledge more than health, whereas someone else may value physical pleasure over knowledge or health. As a result, there can be no clear-cut procedure for determining what action is likely to produce the best outcome for an individual or group.

The effort to develop clear-cut procedures for determining the best action or policy has led some thinkers to replace considerations of intrinsic value with considerations of actual preferences. What someone wants, desires, or prefers can be determined, in principle, in an objective way by consulting the person directly. In addition, people are often able to do more than merely express a preference. Sometimes they can rank their preferences from that which is most desired to that which is least desired.

Such a ranking is of special importance in situations involving risk, for individuals can be asked to decide how much risk they are willing to take to try to realize a given preference. A young woman with a hip injury who is otherwise in good health may be willing to accept the risk of surgery to increase her chances of being restored to many years of active life. By contrast, an elderly woman in frail health may prefer to avoid surgery and accept the limitations that the injury imposes on her physical activities. For the elderly patient, not only are the risks of surgery greater because of her poor health, but even if the surgery is successful, she will have fewer years to benefit from it.

Alternatively, the older woman may place such a premium on physical activity that she is willing to take the risk of surgery to improve her chances of securing even a few more years of it. Only she can say what is important to her and how willing she is to take the risk required to secure it.

These considerations about personal preferences can also be raised with regard to social preferences. Statistical information about what people desire and what they are willing to forego to see their desires satisfied becomes relevant to institutional and legislative deliberations about what policies to adopt. For example, a crucial question facing our own society is whether we are willing to provide everyone with at least a basic minimum of health care, even if this requires harming the private insurance industry or reducing spending on other social goods.

Employing the satisfaction of preferences as a criterion for the moral value of an action or policy makes it possible to incorporate determinate data into our ethical deliberations. The life expectancy of infants with specific impairments at birth can be estimated on a statistical basis; surgical

procedures can be associated with a certain success rate and a certain mortality rate. Similarly, a particular social policy has a certain financial cost, and if implemented, the policy is likely to mean the loss of other possible benefits and opportunities.

Ideally, information of this kind should allow a preference utilitarian to calculate the best course of action for an individual or group. The best action will be the one that best combines the satisfaction of preferences with other conditions (e.g., financial costs and risks) that are at least minimally acceptable. To use the jargon of the theorists, the best action is the one that maximizes the preference utilities of the person or group.

A utilitarianism that employs preferences has the advantage of suggesting more explicit methods of analysis and decision-making than the classical formulation. It also has the potential for being more sensitive to the expressed desires of individuals. However, preference utilitarianism is not free from specific difficulties.

Perhaps most obvious is the problem posed by preferences that we would generally regard as unacceptable. What are we to say about those who prefer mass murder, child abuse, or torturing animals? A second problem involves *adaptive preferences*, or the way people's preferences are often shaped (and limited) by the social conditions currently available to them. These problems suggest that subjective preferences cannot be treated equally or taken at face value, and we must have a way to distinguish acceptable from unacceptable ones. Whether this can be done by relying on the principle of utility alone is doubtful. In the view of some commentators, some other moral principle (or principles) is needed.

Difficulties with Utilitarianism

Classical utilitarianism is open to a variety of objections. Here we will concentrate on only one, however, for it seems to reveal a major flaw in the structure of the entire theory. This serious objection is that the principle of utility appears to justify the imposition of great suffering on a few people for the benefit of many people.

Certain kinds of human experimentation forcefully illustrate this possibility. Suppose an investigator is concerned with acquiring a better understanding of human brain functions. She could learn a great deal by systematically destroying the brain of one person and carefully noting the results. Such a study would offer many more opportunities for increasing medical knowledge than traditional studies that focus on animals or human subjects with accidental brain damage. We may suppose that the investigator chooses a subject without education or training, without family or friends, who cannot be regarded as making a significant contribution to society. The subject will die from the experiment, but it is not unreasonable to suppose that the knowledge of the human brain gained from the experiment will improve the lives of many more people.

The principle of utility seems to make such experiments legitimate because the outcome is a greater quantity of good than harm. One or a few have suffered immensely, but the many will profit to an extent that far outweighs that suffering.

Clearly, what is missing from such scenarios is the concept of *justice*. It cannot be right to increase the general happiness at the expense of one person or group. There must be some way of distributing happiness and unhappiness and avoiding exploitation. Mill acknowledged that utilitarianism needed a principle of justice, but many contemporary philosophers do not believe that such a principle can be derived from the principle of utility. In their opinion, utilitarianism as an ethical theory suffers severely from this defect. Yet some philosophers, while acknowledging the defect, have still held that utilitarianism is the best substantive moral theory available.

Kant's Ethics

For utilitarianism, the rightness of an action depends upon its consequences. In stark contrast to this view is the ethical theory formulated by the German philosopher Immanuel Kant (1724–1804). For Kant, the consequences of an action are morally irrelevant. Rather, an action is right when it accords with a rule satisfying a broader principle that he calls the *categorical imperative*. Since this is the basic principle of Kant's ethics, we will begin our discussion with it.

The Categorical Imperative

If I decide to have an abortion and act on my decision, it is possible to view my action as involving a rule. I can be thought of as endorsing a rule to the effect that "Whenever I am in circumstances like these, then I will have an abortion." Kant calls such a rule a *maxim*. In his view, all reasoned and considered actions can be regarded as involving maxims.

The maxims in such cases are personal or subjective, but they can be thought of as candidates for moral rules. If they pass the test imposed by the categorical imperative, then we can say that the actions they approve are right. Furthermore, in passing the test, these maxims cease to be merely personal and subjective. They gain the status of objective rules of morality that hold for everyone.

Kant formulates the categorical imperative in this way: *Act only on that maxim which you can will to be a universal law.* Kant calls the principle *categorical* to distinguish it from *hypothetical* imperatives. The latter tell us what to do if we want to bring about certain consequences—such as happiness. A categorical imperative prescribes what we ought to do without reference to any consequences. The principle is an "imperative" because it is a command.

The test imposed on maxims by the categorical imperative is one of generalization or *universalizability*. The central idea of the test is that a moral maxim is one that can be generalized to apply to all cases of the same kind. That is, it could be adopted as a maxim by everyone who is in a similar situation, and thus willed as a universal law. For a maxim to satisfy the categorical imperative, it is not necessary that we be willing in some psychological sense to see it universally adopted. Rather, the test is one that requires us to avoid inconsistency or logical conflict in what we will as a universal law.

Suppose, for example, that I am a physician and I tell a patient that he has a serious illness, although I know that he doesn't. This may be to my immediate advantage, for the treatment and the supposed cure will increase my income and reputation. The maxim of my action might be phrased as, "Whenever it is to my advantage, I may lie to healthy patients and tell them that they have serious illnesses."

Now suppose that I try to generalize my maxim. In doing so, I will discover that I am willing the existence of a universal practice with contradictory properties. If "Whenever it is to my advantage, I may lie to healthy patients and tell them that they have serious illnesses" is made a universal law, then trust in the diagnostic pronouncements of physicians will soon be destroyed. But my scheme depends on my patients' trusting me and accepting the truth of my lying diagnosis.

It is as if I were saying, "Let there be a rule of truth telling such that people can assume that physicians are telling them the truth, but let there also be a rule that physicians may lie to their patients when it is in the interest of the physician to do so." In willing both rules, I am willing something contradictory. Thus, I can will my action in an exceptional case, but I can't will that my action be universal without generating a logical conflict.

Kant claims that such considerations show that it is always wrong to lie. Lying produces a contradiction in what we will. On one hand, we will that people believe what