
THE FOUR PRINCIPLES APPROACH TO HEALTH CARE ETHICS

My objective in this essay is to explain the so-called “four principles” approach and to explain the philosophical and practical roles these principles play. I start with a brief history and then turn to the four-principles framework, its practical uses, and philosophical problems of making the abstract framework specific.

THE ORIGINS OF PRINCIPLES IN HEALTH CARE ETHICS

Prior to the early 1970s, there was no firm ground in which a commitment to principles or even ethical theory could take root in biomedical ethics. This is not to say that physicians and researchers had no principled commitments to patients and research subjects. They did, but moral principles, practices, and virtues were rarely discussed. The health care ethics outlook in Europe and America was largely that of maximizing medical benefits and minimizing risks of harm and disease. The Hippocratic tradition had neglected many problems of truthfulness, privacy, justice, communal responsibility, the vulnerability of research subjects, and the like.¹ Views about ethics had been confined largely to the perspectives of those in the professions of medicine, public health, and nursing. No sustained work combined concerns in ethical theory and the health care fields.

Principles that could be understood with relative ease by the members of various disciplines figured prominently in the development of biomedical ethics during the 1970s and early 1980s. Principles were used to present frameworks of evaluative assumptions so that they could be used by, and readily understood by, people with many different forms of professional training. The distilled morality found in principles gave people a shared and serviceable group of general norms for analyzing many types of moral problems. In some respects, it could even be claimed that principles gave the embryonic field of bioethics a shared method for attacking its problems, and this gave some minimal coherence and uniformity to bioethics.

There were two primary sources of the early interest in principles in biomedical ethics. The first was the *Belmont Report* (and related documents) of the National Commission for the Protection of Human Subjects,² and the second was the book entitled *Principles of Biomedical Ethics*, which I coauthored with James F. Childress. I here confine discussion to the latter.

Childress and I began our search for the principles of biomedical ethics in 1975. In early 1976, we drafted the main ideas for the book, although only later would the title *Principles of Biomedical Ethics* be placed on it.³ Our goal was to develop a set of principles suitable for biomedical ethics. One of our proposals was that health care's traditional preoccupation with a beneficence-based model of health care ethics be augmented by an autonomy model, while also incorporating a wider set of social concerns, particularly those focused on social justice. The principles are understood as standards of conduct on which many other moral claims and judgments depend. A principle, then, is an essential norm in a system of moral thought and one that is basic for moral reasoning. More specific rules for health care ethics can be formulated *by reference to* these four principles, but neither rules nor practical judgments can be straightforwardly *deduced* from the principles.

THE FRAMEWORK OF PRINCIPLES

The principles in our framework have always been grouped under four general categories: (1) respect for autonomy (a principle requiring respect for the decision-making capacities of autonomous persons), (2) nonmaleficence (a principle requiring not causing harm to others), (3) beneficence (a group of principles requiring that we prevent harm, provide benefits, and balance benefits against risks and costs), and (4) justice (a group of principles requiring appropriate distribution of benefits, risks, and costs fairly). I will concentrate on an explication of each of the principles and how they are to be understood collectively as a framework.

Respect for Autonomy

Respect for autonomy is rooted in liberal moral and political traditions of the importance of individual freedom and choice. In moral philosophy, personal autonomy refers to personal self-governance: personal rule of the self by adequate understanding while remaining free from controlling interferences by others and from personal limitations that prevent choice. *Autonomy* means freedom from external constraint and the presence of critical mental capacities such as understanding, intending, and voluntary decision making.⁴ The autonomous individual acts freely in accordance with a self-chosen plan, analogous to the way an independent government manages its territories and sets its policies. A person of diminished autonomy, by contrast, is in some respect controlled by others or incapable of deliberating or acting on the basis of his or her desires and plans.

To respect an autonomous agent is to recognize with due appreciation that person's capacities and perspectives, including his or her right to hold certain views, to make certain choices, and to take certain actions based on personal values and beliefs. The moral demand that we respect the autonomy of persons can be expressed as a *principle* of respect for autonomy that states both a negative obligation and a positive obligation. As a negative obligation, autonomous actions should not be subjected to controlling constraints by others. As a positive obligation, this principle requires respectful treatment in informational exchanges and in other actions that foster autonomous decision making.

Many autonomous actions could not occur without others' material cooperation in making options available. Respect for autonomy obligates professionals in health care and research involving human subjects to disclose information, to probe for and ensure understanding and voluntariness, and to foster adequate decision making. True respect requires more than mere noninterference in others' personal affairs. It includes, at least in some contexts, building up or maintaining others' capacities for autonomous choice while helping to allay fears and other conditions that destroy or disrupt their autonomous actions. Respect, on this account, involves acknowledging the value and decision-making rights of persons and enabling them to act autonomously, whereas disrespect for autonomy involves attitudes and actions that ignore, insult, demean, or are inattentive to others' rights of autonomy.

Many issues in professional ethics concern failures to respect a person's autonomy, ranging from manipulative underdisclosure of pertinent information to nonrecognition of a refusal of medical interventions. For example, in the debate over whether autonomous, informed patients have the right to refuse medical interventions, the principle of respect for autonomy suggests that an autonomous decision to refuse interventions must be respected. Although it was not until the late 1970s that serious attention was given to the rights of patients to refuse, this is

no reason for thinking that respect for autonomy as now understood is a newly added principle in our moral perspective. It simply means that the implications of this principle were not widely appreciated until recently.⁵

Controversial problems with the principle of respect for autonomy, as with all moral principles, arise when we must interpret its significance for particular contexts, determine precise limits on its application, and decide how to handle situations in which it conflicts with other moral principles. Many controversies involve questions about the conditions under which a person's right to autonomous expression demands actions by others, as well as questions about the restrictions society may rightfully place on choices by patients or subjects when these choices conflict with other values. If restriction of the patient's autonomy is in order, the justification will always rest on some competing moral principle such as beneficence or justice.

Nonmaleficence

Physicians have long avowed that they are obligated to avoid doing harm to their patients. Among the most quoted principles in the history of codes of health care ethics is the maxim *primum non nocere*: "Above all, do no harm." British physician Thomas Percival furnished the first developed modern account of health care ethics. He maintained that a principle of nonmaleficence fixes the physician's primary obligations and triumphs even over respect for the patient's autonomy in a circumstance of potential harm to patients:

To a patient . . . who makes inquiries which, if faithfully answered, might prove fatal to him, it would be a gross and unfeeling wrong to reveal the truth. His right to it is suspended, and even annihilated; because . . . it would be deeply injurious to himself, to his family, and to the public. And he has the strongest claim, from the trust reposed in his physician, as well as from the common principles of humanity, to be guarded against whatever would be detrimental to him.⁶

Some basic rules in the common morality are requirements to avoid causing a harm. They include rules such as do not kill; do not cause pain; do not disable; do not deprive of pleasure; do not cheat; and do not break promises.⁷ Similar but more specific prohibitions are found across the literature of biomedical ethics, each grounded in the principle that intentionally or negligently caused harm is a fundamental moral wrong.

Numerous problems of nonmaleficence are found in health care ethics today, some involving blatant abuses of persons and others involving subtle and

unresolved questions. Blatant examples of failures to act nonmaleficently are found in the use of physicians to classify political dissidents as mentally ill, thereafter treating them with harmful drugs and incarcerating them with insane and violent persons.⁸ More subtle examples are found in the use of medications for the treatment of aggressive and destructive patients. These common treatment modalities are helpful to many patients, but they can be harmful to others.

A provocative question about nonmaleficence and physician ethics has been raised by Paul S. Appelbaum in an investigation of “the problem of doing harm” through testimony in criminal contexts and civil litigation—for example, by omitting information in the context of a trial, after which a more severe punishment is delivered to the person than likely would have been delivered. Appelbaum presents the generic problem as one of nonmaleficence:

If physicians are committed to doing good and avoiding harm, how can they participate in legal proceedings from which harm may result? If, on the other hand, physicians in court abandon medicine’s traditional ethical principles, how do they justify that deviation? And if the obligations to do good and avoid harm no longer govern physicians in the legal setting, what alternative principles come into play? . . . Are physicians in general bound by the principles of beneficence and nonmaleficence?⁹

Beneficence

The physician who professes to “do no harm” is not usually interpreted as pledging never to cause harm, but rather to strive to create a positive balance of goods over inflicted harms. Those engaged in medical practice, research, and public health know that risks of harm presented by interventions must often be weighed against possible benefits for patients, subjects, and the public. Here we see the importance of beneficence as a principle beyond the scope of nonmaleficence.

In ordinary English, the term *beneficence* connotes acts of mercy, kindness, charity, love, and humanity. In its most general meaning, it includes all forms of action intended to benefit other persons. In health care ethics beneficence commonly refers to an action done to benefit others, whereas benevolence refers to the character trait or virtue of being disposed to act for the benefit of others. The principle of beneficence refers to a moral obligation to act for the benefit of others. No demand is more important in taking care of patients: The welfare of patients is medicine’s context and justification. Beneficence has long been treated as a

foundational value—and sometimes as *the* foundational value¹⁰—in health care ethics.

The principle of beneficence requires us to help others further their important and legitimate interests, often by preventing or removing possible harms. This principle includes rules such as “maximize possible benefits and minimize possible harms” and “balance benefits against risks.” Many duties in medicine, nursing, public health, and research are expressed in terms of a positive obligation to come to the assistance of those in need of treatment or in danger of injury. The harms to be prevented, removed, or minimized are the pain, suffering, and disability of injury and disease. The range of benefits that might be considered relevant is broad. It could even include helping patients find appropriate forms of financial assistance and helping them gain access to health care or research protocols. Sometimes the benefit is for the patient, at other times for society.

Some writers in health care ethics suggest that certain duties not to injure others are more compelling than duties to benefit them. They point out that we do not consider it justifiable to kill a dying patient in order to use the patient’s organs to save two others, even though benefits would be maximized, all things considered. The obligation not to injure a patient by abandonment has been said to be stronger than the obligation to prevent injury to a patient who has been abandoned by another (under the assumption that both are moral duties). Despite the attractiveness of these notions that there is a hierarchical ordering rule, Childress and I reject such hierarchies on grounds that obligations of beneficence do, under many circumstances, outweigh those of nonmaleficence. A harm inflicted by not avoiding causing it may be negligible or trivial, whereas the harm that beneficence requires we prevent may be substantial. For example, saving a person’s life by a blood transfusion clearly justifies the inflicted harm of venipuncture on the blood donor. One of the motivations for separating nonmaleficence from beneficence is that these principles themselves come into conflict. Since the weights of the two principles can vary, there can be no mechanical decision rule asserting that one obligation must always outweigh the other.

Perhaps the major theoretical problem about beneficence is whether the principle generates general moral duties that are incumbent on everyone—not because of a professional role, but because morality itself makes a general demand of beneficence. Many analyses of beneficence in ethical theory (most notably utilitarianism¹¹) seem to demand severe sacrifice and extreme generosity in the moral life—for example, giving a kidney for transplantation or donating bone marrow to a stranger. However, many moral philosophers have argued that such beneficent action is virtuous and a moral ideal, but not an obligation; therefore, there is no principle of beneficence of the sort proclaimed in the four-principles approach.

I agree that the line between what is required and what is not required by the principle is difficult to draw, and that drawing a precise line independent of context is impossible. I do not agree, however, with the radical view that there are no obligations of beneficence—both general and specific obligations. I return to this problem of weighing, judging, and specifying later in this essay in a discussion of the notion of *prima facie* duties.

Justice

Every civilized society is a cooperative venture structured by moral, legal, and cultural principles of justice that define the terms of cooperation. A person in any such society has been treated justly if treated according to what is fair, due, or owed. For example, if equal political rights are due to all citizens, then justice is done when those rights are accorded. The more restricted notion of *distributive justice* refers to fair, equitable, and appropriate distribution in society. Usually this term refers to the distribution of primary social goods, such as economic goods and fundamental political rights, but burdens are also within its scope. Paying for forms of national health insurance is a distributed burden; medical-welfare checks and grants to do research are distributed benefits.

There is no single principle of justice in the four-principles approach. Somewhat like principles under the heading of beneficence, there are several principles, each requiring specification in particular contexts. But common to almost all theories of justice—and accepted in the four-principles approach—is the minimal (formal) principle that like cases should be treated alike, or, to use the language of equality, equals ought to be treated equally and unequals unequally. This elementary principle, or formal principle of justice, states no particular respects in which people ought to be treated. It merely asserts that whatever respects are relevant, if persons are equal in those respects, they should be treated alike. Thus, the formal principle of justice does not tell us how to determine equality or proportion in these matters, and it lacks substance as a specific guide to conduct.

Many controversies about justice arise over what should be considered the relevant characteristics for equal treatment. Principles that specify these relevant characteristics are often said to be “material” because they identify relevant properties for distribution. Childress and I take account of the fact that philosophers have developed diverse theories of justice that provide sometimes conflicting material principles. We try to show that there are some merits in egalitarian theories, libertarian theories, and utilitarian theories; and we defend a mixed use of principles in these theories. We regard these three theories of justice as appropriately capturing some of our traditional convictions about justice,

and we think that they can all be tapped as resources that will help to produce a coherent conception of justice.

However, many issues of justice in health care ethics are not easily framed in the context of traditional principles and abstract moral theories.¹² For example, some basic issues in health care ethics in the last three decades center on special levels of protection and aid for vulnerable and disadvantaged parties in health care systems. These issues cut across clinical ethics, public health ethics, and research ethics. The four-principles approach tries to deal with several of these issues, without producing a grand theory for resolving all issues of justice. For example, we address issues in research ethics about whether research is permissible with groups who have been repeatedly used as research subjects, though the advantages of research are calculated to benefit all in society. We argue that since medical research is a social enterprise for the public good, it must be accomplished in a broadly inclusive and participatory way, and we try to specify the commitments of such generalizations. In this way, we incorporate principles of justice but do not produce a general theory of justice.

THE CHOICE OF FOUR PRINCIPLES AND THE EVOLUTION OF THE THEORY

The choice of our four types of moral principle as the framework for moral decision making in bioethics derives in part from professional roles and traditions. As I noted earlier, the obligations and virtues of health professionals have for centuries, as found in codes and learned writings on ethics, been framed by professional commitments to provide medical care and to protect patients from disease, injury, and system failure. Our principles build on this tradition, but they also significantly depart from it by including parts of morality that traditionally have been neglected in health care ethics, attending especially to the traditional neglect of principles of respect for autonomy and justice. All four types of principles are needed to provide a comprehensive framework for biomedical ethics, but this general framework is abstract and spare until it has been further specified—that is, interpreted and adapted for particular circumstances.

Principles of Biomedical Ethics has evolved appreciably since the first edition in its understanding of abstractness and the demands of particular circumstances. This is not because the principles have changed, but because over the years Childress and I have altered some of our views about the grounding of the principles and about their practical significance. Two major changes deserve special attention. The first is our development of the idea that the four principles are already embedded in public morality—a universal common morality—and are presupposed in the formulation of public and institutional policies.

The second is our adoption of Henry Richardson's account of the specification of moral norms. These supplements to our theory and their significance will be discussed in the next two sections.

The Centrality of the Common Morality

The source of the four principles is what Childress and I call *the common morality* (a view only incorporated at the point of the third edition of *Principles*, following the language of Alan Donagan). The common morality is applicable to all persons in all places, and all human conduct is rightly judged by its standards. The following are examples of *standards of action* (here rules of obligation) in the common morality: (1) don't kill, (2) don't cause pain or suffering to others, (3) prevent evil or harm from occurring, (4) rescue persons in danger, (5) tell the truth, (6) nurture the young and dependent, (7) keep your promises, (8) don't steal, (9) don't punish the innocent, and (10) treat all persons with equal moral consideration.

Why have such norms become parts of a common morality, whereas other norms have not? To answer this question, I start with an assumption about the primary goal—that is, objective—of the social institution of morality. This objective is to promote human flourishing by counteracting conditions that cause the quality of people's lives to worsen. The goal is to prevent or limit problems of indifference, conflict, suffering, hostility, scarce resources, limited information, and the like. Centuries of experience have demonstrated that the human condition tends to deteriorate into misery, confusion, violence, and distrust unless norms of the sort just listed—the norms of the common morality—are observed. When complied with, these norms lessen human misery and preventable death. It is an overstatement to maintain that all of these norms are necessary for the survival of a society (as some philosophers and social scientists have maintained¹³), but it is not too much to claim that these norms are necessary to ameliorate or counteract the tendency for the quality of people's lives to worsen or for social relationships to disintegrate.¹⁴

These norms are what they are, and not some other set of norms, because they have proven over time that their observance is essential to realize the objectives of morality. What justifies them is that they achieve the objectives of morality, not the fact that they are universally shared across cultures. It is conceivable, of course, that the set of norms that is shared universally is not the same set of norms as the set pragmatically justified by their conformity to the objectives of morality. I agree that if another set of norms would better serve the objectives of morality, then that set of norms ought to displace the norms currently in place. However, there are no clear candidates as alternatives to these norms.

What Childress and I call “principles” are the most general and basic norms of the common morality. In *Principles of Biomedical Ethics*, we devote an entire chapter to each principle in the attempt to explain its nature, content, specification, and the like. The assumption behind the argument in each chapter is that our framework of four principles should incorporate and articulate the most general values of the common morality.

Our framework encompasses several types of moral norms, including not only principles but also rules, rights, and moral ideals. We treat principles as the most general and comprehensive norms, but we make only a loose distinction between rules and principles. Rules, we argue, are more precise and practical guides to action that depend on the more general principles for their justification. We defend several types of rules, all of which should be viewed as specifications of principles. These include substantive rules (e.g., truth telling, confidentiality, and privacy rules), authority rules (e.g., rules of surrogate authority and rules of professional authority), and procedural rules (e.g., rules for determining eligibility for organ transplantation and rules for reporting grievances to higher authorities).

The Prima Facie Character of Principles and Rules

These principles and rules (or other norms in the common morality) can in some circumstances be justifiably overridden by other moral norms with which they conflict. For example, we might justifiably not tell the truth in order to prevent someone from killing another person, and we might justifiably disclose confidential information about a person in order to protect the rights of another person. Principles, duties, and rights are not absolute (or unconditional) merely because they are universal. There are exceptions to all principles, each of which is merely presumptive in force.

Oxford philosopher W. D. Ross developed a theory that has been influential on *Principles* since the first edition. Ross’s theory is intended to assist in resolving problems of conflict between principles. His views are based on an account of prima facie duties, which he contrasts with actual duties. A prima facie duty is one that is always to be acted upon unless it conflicts on a particular occasion with another duty. One’s *actual* duty, by contrast, is determined by an examination of the respective weights of competing prima facie duties in particular situations. When principles contingently conflict, no supreme principle is available—in the four-principles approach—to determine an overriding obligation. Therefore, discretionary judgment becomes an inescapable part of moral thinking in our approach.

Here is an example. A physician has confidential information about a patient who is also an employee in the hospital where the physician practices. The employee is seeking advancement in a stress-filled position, but the physician has good reason to believe this advancement would be devastating for both the employee and the hospital. The physician has duties of confidentiality, nonmaleficence, and beneficence in these circumstances: Should the physician break confidence? Could the matter be handled by making thin disclosures only to the hospital administrator and not to the personnel office? Can such disclosures be made consistent with one's general commitments to confidentiality? Addressing these questions through a process of moral justification is required to establish one's actual duty in the face of these conflicts of *prima facie* duties. Once we acknowledge that all general principles have exceptions, we are free to view every moral conclusion that is supported by a principle and every principle itself as subject to modification or reformulation. Change of this sort is to be accomplished through *specification*, the means by which principles come to have their practical value.

The Specification of Norms and Moral Coherence

To say that principles have their origins in and find support in the common morality and in traditions of health care is not to say that their appearance in a developed system of biomedical ethics is identical to their appearance in the traditions from which they spring. *Prima facie* principles underdetermine moral judgments because there is too little content in abstract principles to determine concrete outcomes. Every norm and theory contains regions of indeterminacy that need reduction through further development of their commitments in the system, augmenting them with a more specific moral content. I turn, then, to these questions: How does the *prima facie* conception of principles work in practical bioethics? How are general principles to reach down to concrete policies? How does one fill the gap between abstract principles and concrete judgments?

The Method of Specification

The answer is that principles must be specified to suit the needs and demands of particular contexts, thus enabling principles to overcome their lack of content and to handle moral conflict. Specification is a process of reducing the indeterminateness of abstract norms and providing them with specific action-guiding content.¹⁵ For example, without further specification, "do no harm" is too abstract to help in

thinking through problems such as whether physicians may justifiably hasten the death of patients. The general norm has to be specified for this particular context.

Specification is not a process of either producing or defending general norms such as those in the common morality; it assumes that they are available. Specifying the norms with which one starts (whether those in the common morality or norms that were previously specified) is accomplished by *narrowing the scope* of the norms, not by explaining what the general norms *mean*. The scope is narrowed, as Henry Richardson puts it, by “spelling out where, when, why, how, by what means, to whom, or by whom the action is to be done or avoided.”¹⁶ For example, the norm that we are obligated to “respect the autonomy of persons” cannot, unless specified, handle complicated problems of what to disclose or demand in clinical medicine and research involving human subjects. A definition of “respect for autonomy” (as, say, “allowing competent persons to exercise their liberty rights”) might clarify one’s meaning in using the norm, but would not narrow the general norm or render it more specific.

Specification adds content to general norms. For example, one possible specification of “respect the autonomy of persons” is “respect the autonomy of competent patients when they become incompetent by following their advance directives.” This specification will work well in some medical contexts, but will not be adequate in others, thus necessitating additional specification. Progressive specification can continue indefinitely, gradually reducing the conflicts that abstract principles themselves cannot resolve. However, to qualify all along the way as a specification, a transparent connection must always be maintained to the initial norm that gives moral authority to the resulting string of norms.

Now we come to a critical matter about particular moralities, by contrast to the common morality. There is always the possibility of developing more than one line of specification of a norm when confronting practical problems and moral disagreements. It is simply part of the moral life that different persons and groups will offer different (sometimes conflicting) specifications, potentially creating multiple particular moralities. On any problematic issue (such as abortion, animal research, aid in disaster relief, health inequities, and euthanasia), competing specifications are likely to be offered by reasonable and fair-minded parties, all of whom are committed to the common morality. We cannot hold persons to a higher standard than to make judgments conscientiously in light of the relevant basic and specified norms while attending to the available factual evidence. Conscientious and reasonable moral agents will understandably disagree with equally conscientious persons over moral weights and priorities in circumstances of a contingent conflict of norms.

Nothing in the model of specification suggests that we can always eliminate circumstances of intractable conflicting judgments. However, we should always try to do so by justifying the specification we put forward. This suggests that

specification as a method needs to be connected to a model of justification that will support some specifications and not others. Only brief attention can be paid here to this difficult philosophical problem.

Justifying Specifications Using the Method of Coherence

A specification is justified, in the four-principles approach, if and only if it maximizes the coherence of the overall set of relevant, justified beliefs. These beliefs could include empirically justified beliefs, justified basic moral beliefs, and previously justified specifications. This is a version of so-called wide reflective equilibrium.¹⁷ No matter how wide the pool of beliefs, there is no reason to expect that the process of rendering norms coherent by specification will come to an end or be perfected. Particular moralities are, from this perspective, continuous works in progress—a process rather than a finished product. There is no reason to think that morality can be rendered coherent in only one way through the process of specification. Many particular moralities present coherent ways to specify the common morality. Normatively, we can demand no more than that agents faithfully specify the norms of the common morality with an attentive eye to overall coherence.

The following are some of the criteria for a coherent (and therefore, according to this model, justified) set of ethical beliefs: consistency (the avoidance of contradiction), argumentative support (explicit support for a position with reasons), intuitive plausibility (the feature of a norm or judgment being secure in its own right), compatibility or coherence with reasonable nonmoral beliefs (in particular, coherence with available empirical evidence), comprehensiveness (the feature of covering the entire moral domain, or as much of it as possible), and simplicity (reducing the number of moral considerations to the minimum possible without sacrifice in terms of the other criteria).¹⁸

CONCLUSION

I have explained, and argued in defense of, what has often been called the four-principles approach to biomedical ethics and is now increasingly called *principlism*.¹⁹ The four clusters of principles derive from both considered judgments in the common morality and enduring and valuable parts of traditions of health care. Health care ethics has often been said to be an “applied ethics,” but this metaphor may be more misleading than helpful. It is rarely the case that we can directly apply a principle to resolve a tough problem. We will almost always be engaged in collecting evidence, reasoning, and specifying general principles. This

is how problems should be treated and how progress can be made in health care ethics. From this perspective, the four principles form only a starting point—the point where the practical work begins.

Notes

1. Albert R. Jonsen, *The Birth of Bioethics* (New York: Oxford University Press, 1998), 3ff; Edmund D. Pellegrino and David C. Thomasma, *The Virtues in Medical Practice* (New York: Oxford University Press, 1993), 184–189.
2. National Commission for the Protection of Human Subjects of Biomedical and Behavioral Research, *The Belmont Report: Ethical Principles and Guidelines for the Protection of Human Subjects of Research* (Washington, DC: DHEW Publication OS 78-0012 1978); James F. Childress, Eric M. Meslin, Harold T. Shapiro, eds. *Belmont Revisited: Ethical Principles for Research with Human Subjects* (Washington, DC: Georgetown University Press, 2005).
3. Tom L. Beauchamp and James F. Childress, *Principles of Biomedical Ethics*, 1st edition (New York: Oxford University Press, 1979).
4. H. Tristram Engelhardt, Jr., *The Foundations of Bioethics* (New York: Oxford University Press, 1996), 2nd edition; Jay Katz, *The Silent World of Doctor and Patient* (New York: The Free Press, 1984); James F. Childress, “The Place of Autonomy in Bioethics,” *Hastings Center Report* 20 (January/February 1990): 12–16; Rebecca Kukla, “Conscientious Autonomy: Displacing Decisions in Health Care,” *Hastings Center Report* (March–April 2005): 34–44.
5. Ruth R. Faden and Tom L. Beauchamp, *A History and Theory of Informed Consent* (New York: Oxford, 1986).
6. Thomas Percival, *Medical Ethics; or a Code of Institutes and Precepts, Adapted to the Professional Conduct of Physicians and Surgeons* (Manchester: S. Russell, 1803), 165–166. Percival’s work was the pattern for the American Medical Association’s (AMA) first code of ethics in 1847.
7. Compare Bernard Gert, *Morality: Its Nature and Justification* (New York: Oxford, 2005).
8. See, for example, Sidney Bloch and Peter Reddaway, *Soviet Psychiatric Abuse: The Shadow over World Psychiatry* (Boulder, CO: Westview Press, 1984), esp. ch. 1.
9. Paul S. Appelbaum, “The Parable of the Forensic Physician: Ethics and the Problem of Doing Harm,” *International Journal of Law and Psychiatry* 13 (1990): 249–259, esp. 250–251.
10. Edmund Pellegrino and David Thomasma, *For the Patient’s Good: The Restoration of Beneficence in Health Care* (New York: Oxford University Press, 1988); Pellegrino, “The Four Principles and the Doctor-Patient Relationship: The Need for a Better Linkage,” in Raanan Gillon, ed. *Principles of Health Care Ethics* (Chichester, England: John Wiley & Sons, 1994).
11. Peter Singer, “Living High and Letting Die,” *Philosophy and Phenomenological Research* 59 (1999): 183–187. Peter Singer, *Practical Ethics*, 2nd edition (Cambridge: Cambridge University Press, 1993); Shelly Kagan, *The Limits of Morality* (Oxford:

Clarendon Press, 1989); Richard W. Miller, "Beneficence, Duty, and Distance," *Philosophy & Public Affairs* 32 (2004): 357–383.

12. For diverse accounts of justice in biomedical ethics, see Norman Daniels, *Just Health Care* (New York: Cambridge University Press, 1985); Daniels, *Just Health* (New York: Cambridge University Press, 2006); Madison Powers and Ruth Faden, *Social Justice: The Moral Foundations of Public Health and Health Policy* (New York: Oxford University Press, 2006); Allen Buchanan, "Health-Care Delivery and Resource Allocation," in *Medical Ethics*, ed. Robert Veatch, 2nd edition (Boston: Jones and Bartlett Publishers, 1997); Allen Buchanan, Dan Brock, Norman Daniels, and Daniel Wikler, *From Chance to Choice: Genetics and Justice* (New York: Cambridge University Press, 2000).

13. Sissela Bok, *Common Values* (Columbia, MO: University of Missouri Press, 1995), 13–23, 50–59.

14. G. J. Warnock, *The Object of Morality* (London: Methuen & Co., 1971), esp. 15–26; John Mackie, *Ethics: Inventing Right and Wrong* (London: Penguin, 1977), 107ff.

15. Henry S. Richardson, "Specifying Norms as a Way to Resolve Concrete Ethical Problems," *Philosophy and Public Affairs* 19 (Fall 1990): 279–310; Richardson, "Specifying, Balancing, and Interpreting Bioethical Principles," *Journal of Medicine and Philosophy* 25 (2000): 285–307; David DeGrazia, "Moving Forward in Bioethical Theory: Theories, Cases, and Specified Principlism," *Journal of Medicine and Philosophy* 17 (1992): 511–539; and DeGrazia and Tom L. Beauchamp, "Philosophical Foundations and Philosophical Methods," in *Methods of Bioethics*, ed. D. Sulmasy and J. Sugarman (Washington, DC: Georgetown University Press, 2001), esp. 33–36.

16. Richardson, "Specifying, Balancing, and Interpreting Bioethical Principles," 289.

17. Norman Daniels, "Wide Reflective Equilibrium and Theory Acceptance in Ethics," *Journal of Philosophy* 76 (1979): 256–282; Daniels, "Wide Reflective Equilibrium in Practice," in L.W. Sumner and J. Boyle, *Philosophical Perspectives on Bioethics* (Toronto: University of Toronto Press, 1996).

18. DeGrazia and Beauchamp, "Philosophical Foundations and Philosophical Methods"; and David DeGrazia, "Common Morality, Coherence, and the Principles of Biomedical Ethics," *Kennedy Institute of Ethics Journal* 13 (2003): 219–230.

19. See, for example, B. Gert, C. M. Culver, and K. D. Clouser, *Bioethics: A Return to Fundamentals* (New York: Oxford University Press, 1997), ch. 4; John H. Evans, "A Sociological Account of the Growth of Principlism," *Hastings Center Report* 30 (Sept.–Oct. 2000): 31–38; Earl Winkler, "Moral Philosophy and Bioethics: Contextualism versus the Paradigm Theory," in L. W. Sumner and J. Boyle, eds. *Philosophical Perspectives on Bioethics*; John D. Arras, "Principles and Particularity: The Role of Cases in Bioethics," *Indiana Law Journal* 69 (1994): 983–1014; and Carson Strong, "Specified Principlism," *Journal of Medicine and Philosophy* 25 (2000): 285–307.