

Outline of the Dual Value Theory

3.1 Introduction

As elaborated in Chapter 2, our methodology in ethics is *reflective equilibrium* – also called *the coherence model of moral reasoning*. Motivating this methodology is the assumption that no single level of ethical analysis – such as a foundational principle, mid-level principles, rules, or judgments about specific cases – deserves priority in our moral reasoning. It is true that considered judgments play a special role. Considered judgments are moral judgments that we have especially good reason to consider reliable; for this reason, an ethical theory may be constructed at least partly on the basis of considered judgments. But the considered judgments that form the starting points of ethical theorizing may be highly abstract and general, very specific and context-dependent, or of any level of intermediate generality.

What counts as an “ethical theory”? We think of an ethical theory as *a structure of general moral norms that helps to render specific ethical verdicts*. For example, utilitarianism at its core is a single ethical principle, the principle of utility – where the principle requires specification in terms of a particular account of utility, its scope, and the like. Once adequately specified, the principle of utility serves as the ultimate normative basis for all of the judgments that comprise utilitarian moral thinking. Another example is Gert’s rules-based approach, which comprises a system of ten rules and a decision-making procedure for dealing with conflicts among rules as the basis for sound ethical reasoning.¹ Here, the structure of rules plus the decision-making procedure constitutes the theory.

Our ethical theory cannot be stated so succinctly. This is because the method of reflective equilibrium reveals the best-justified theory to be more intricate than these examples suggest. Nevertheless, our ethical

¹ Bernard Gert, *Common Morality* (New York: Oxford University Press, 2004).

theory does have a discernible structure, so it is not just a collection of considered judgments. Understanding that structure will prove valuable for understanding how we specify our theory in the chapters that follow and how the theory would be naturally extended to topics that we lack the space to cover in this book. In this chapter, we sketch our ethical theory at two levels of generality: (1) fundamental values, principles, and scope; and (2) specification in terms of mid-level principles. Our arguments in favor of the theory, as well as more detailed specifications of its parts, comprise the rest of this book. In the final part of the chapter, we discuss how our theory relates to other familiar theories and concepts in moral philosophy, including consequentialism, deontology, virtue ethics, and the ethical theories of other bioethicists.

The theory we describe here is the one that we think applies to ethical reasoning in any area of life, not just bioethics. Bioethics is not special in the sense of needing its own ethical principles or methods. It is special in that it throws up a number of particularly difficult and important ethical questions, which a general ethical theory can help us answer. It may prove helpful to keep this in mind throughout the book. We emphasize those aspects of theory that are especially important for bioethics; illustrate them with examples from medicine, public health, health policy, and the like; and apply the theory to problems in bioethics. But, at its core, the theory is independent of the subject matter to which we apply it.

3.2 Fundamental Values, a Formal Principle, and Scope

At the highest level of generality, our theory consists of two broad substantive values, a formal distributive principle, and a scope determining the set of beings with moral status. The two broad values are *well-being* and *respect for rights-holders*. The distributive principle is *equal consideration* for all beings with moral status. And the scope is the set of *sentient beings*. Although our approach confers equal moral consideration on all sentient beings, it does not regard all sentient beings as rights-holders.

The value of well-being grounds duties to benefit others and prohibitions against harming them. Given the scope of the theory, moral agents have these duties with respect to all sentient beings. For those sentient beings who have moral status but not rights, equal consideration takes a consequentialist form, which means that the well-being of one individual can be traded off against the well-being of another. Moral agents have additional duties to those sentient beings who are also rights-holders. First, duties against harming rights-holders are much more stringent – it is not

sufficient to show that causing them harm would result in greater benefit for others. Second, rights-holders who are capable of acting autonomously have a set of autonomy rights whose exercise should be respected. Finally, rights-holders have positive rights held against individuals and institutions on the basis of distributive justice. Difficult bioethical questions frequently involve a tension between the two values of well-being and respect for rights-holders. Since our ethical theory regards both as important, how they should be balanced depends on working through the process of reflective equilibrium for particular cases.

Fundamental Values: Well-Being and Respect for Rights-Holders

The two substantive moral values of well-being and respect for rights-holders may also be understood as ethical principles: “Promote well-being” and “Treat rights-holders with respect.” A great deal must be done to explicate these principles before they can help guide action.

First, we must say what well-being and rights are. The account of well-being we defend is a subjective theory, in that what is in an individual’s interests is necessarily related to facts about that individual’s psychology. According to this account, both enjoyment and the satisfaction of “narrative-relevant” desires – that is, desires that are relevant to one’s life story – are prudentially good for an individual. Meanwhile, both suffering and the frustration of narrative-relevant desires are prudentially bad for an individual. In our view, reality has an amplifying effect. Enjoyment contributes more to one’s well-being when it responds to a state of affairs that (unlike, say, a delusion) actually obtains. In parallel, the fulfillment of desires contributes more to well-being when those desires are relevantly informed. What unifies enjoyment and desire-satisfaction in our account is the fact that both reflect the lived, self-caring perspective of a conscious subject.

A moral right is a justified moral claim that (1) imposes an obligation on one or more individuals and (2) ordinarily resists appeals to the common good as grounds for overriding the claim. For example, persons have rights against bodily trespass, which protect their interests in controlling their own lives, not being subject to harm, and the like. The right against bodily trespass imposes obligations on moral agents not to interfere with the bodies of other persons. Such obligations can be suspended (for example, if a competent person gives consent to be touched and so *waives* her right) or overridden by overwhelmingly strong reasons based on well-being (for example, in the case of mandatory reporting of certain infectious diseases where someone’s privacy right may be infringed in a limited way in order

to track the spread of disease and protect other members of society). It is worth noting that writers use the term “right” in different ways. In particular, it is common to say that if one individual owes a duty (or obligation) to another, then the latter has a right against the former.² This comprises only half of what we mean by a *right*, since it does not include the restriction on overriding the right for the common good. As we use the term, it is possible to owe duties to non-rights-holders.

The right against bodily trespass is an example of what is called a *negative right* – it imposes obligations on others to refrain from some action or actions. Other important negative rights in contemporary bioethics include rights to medical confidentiality and to nondiscrimination in the provision of services. Autonomous persons have additional *autonomy rights*, which include the power to waive others’ obligations. This is what underlies consent, whereby a competent patient is entitled to make a free decision, following adequate disclosure, to either authorize or reject an offer of a particular medical intervention. If a patient gives valid consent to treatment, they thereby waive their right against bodily trespass in this particular circumstance.

Our theory also asserts some *positive rights*, which impose obligations on others to perform some action or actions. For example, patients have the right to be provided with adequate health care as a matter of distributive justice. Such positive rights are typically limited, since resources are limited. The right to adequate health care does not entail that patients have rights to anything they please – for example, to care that is not medically indicated or to care that is so expensive relative to its efficacy that providing it would financially threaten the medical system. While it is important to recognize rights wherever there are solid moral grounds for asserting them and the obligations they entail can realistically be met, it is also important to acknowledge that what individuals have a claim to receive may depend on the resources available and the claims of others.

Well-being and respect are both important in our theory. Neither is supreme. Often, promoting well-being and respecting rights-holders will coincide. Sometimes, however, they will conflict. Where they conflict, there is no simple procedure for settling the conflict. Fortunately, the method of reflective equilibrium frees us from any notion that ethically right action must be *derived* from the most general ethical principles or values. Moral reflection at various levels of generality – and considerations

² See, e.g., Tom Beauchamp and James Childress, *Principles of Biomedical Ethics*, 7th ed. (New York: Oxford University Press, 2013), 371–372.

of coherence throughout – help to work out what is ethically defensible even when well-being and respect apparently conflict. For example, when a quarantine is morally justified, an individual's right to freedom of movement in public places is overridden by appeal to social utility, settling the conflict between well-being and respect (temporarily) in favor of the former. This resolution of the conflict rests on the conviction that the overall set of norms that incorporates this judgment – that it is permissible to override the right temporarily – is more plausible and coherent than the overall set of norms that includes the contrary judgment upholding the right and prohibiting a quarantine. Consider another example. No matter how much valuable information could be generated by a medical experiment that placed young children at high risk of death with no compensating medical benefits, we judge that such research would be unjustified. Children's right to adequate protection from harm would trump appeals to societal well-being. This way of settling the conflict between well-being and respecting rights – in favor of the latter – rests on the conviction that the overall set of norms that includes this verdict is more plausible and coherent than the overall set of norms that permits overriding pediatric subjects' rights to adequate protection. Because our ethical theory treats well-being and respect as its most fundamental values or principles but not as foundations from which ethical verdicts can be derived, their coequal status does not paralyze ethical analysis.

As the most general values in our ethical theory, well-being and respect for rights-holders call for careful analysis. We elaborate on respect for rights-holders in Chapter 7, "Moral Status." We devote Chapter 8 to the nature of individual well-being. In addition, questions of benefiting and harming, and analyses of individual rights, are integral to the development of our moral theory throughout the book.

A Formal Principle: Equal Consideration

In addition to featuring the general substantive values of well-being and respect, our ethical theory features a very general principle that indicates, for each individual who counts morally or possesses moral status, *how much* they count in relation to others. In our ethical theory the principle of *equal consideration of interests* – or *equal consideration* for short – plays this role. This principle asserts that everyone's interests are to receive impartial consideration. For example, other things being equal, causing x amount of suffering to A is as morally problematic as causing x amount of suffering to B, irrespective of the species, traits, or capacities of A and B. It would be

inconsistent with equal consideration to judge that it is worse to cause moderate suffering to A than to B just because A is a competent adult whereas B is a toddler or just because A is a human being and B is a pig. Note that equal consideration is about regarding *individuals* as equally deserving of moral concern (from an impartial standpoint), not treating each individual *the same*. For example, we take a moderate prioritarian view about distributive justice, according to which it is more important to benefit those rights-holders who are worse off than those who are better off. All else being equal, the interests of the worse off are therefore valued more highly. Nevertheless, we think that they receive equal consideration, since if their situations were reversed, priority would again be given to the ones who are worse off.

In asserting that everyone's similar interests are equally morally important, the principle of equal consideration needs a criterion for comparability of interests. The criterion is *prudential*: two interests are of similar magnitude when what is at stake for the individuals is roughly the same in terms of their well-being. So a human's interest in not suffering moderately may be assumed to be comparable to a turtle's interest in not suffering moderately because, we might say, moderate suffering is *moderately awful* whenever it occurs.³ In contrast, in ordinary cases a person with life plans has a much greater stake in remaining alive than a turtle who has very little sense of the future, so that equal consideration does *not* require attributing equal moral importance to the person's continued life and the turtle's continued life. It is far worse that the person dies than that the turtle dies. What equal consideration requires is that agents give equal moral weight to *prudentially similar interests* irrespective of what sorts of beings the interest-bearers are.

Equal consideration is an extremely general and abstract principle that is assumed – at least for application to persons or human beings – by a wide variety of moral theories. Utilitarianism incorporates equal consideration because the principle of utility gives equal or impartial consideration to all beings who have a welfare. Equal consideration is also assumed in Kantian ethics, libertarianism, and most other prominent deontological theories, although they usually limit the scope of this principle to humanity. Meanwhile, views that attribute rights to animals also assume a type of equal consideration for humans and animals.

³ Here we bracket the complication that suffering may be not only intrinsically bad for a subject but also instrumentally bad in thwarting some of their projects or valued activities, which may vary in value from individual to individual.

Equal consideration is compatible with various theories because it is a formal principle – where what is under consideration is rather abstract, namely, moral importance or how much one's interests matter morally. Indeed, equal consideration is compatible with a nihilism according to which *nothing* matters morally. A nihilist could say, "I give everyone's comparable interests equal moral weight – to wit, none." Thus, the significance of the formal principle of equal consideration depends on substantive values. As we have seen, the two most general substantive values in our ethical theory are well-being and respect (for rights-holders). So, in effect, at a very general level we answer the question "*What* is substantively at issue in ethics?" with the substantive answer "well-being and respect." Again at a very general level, we answer the question, "*How* are individuals to be regarded with respect to these fundamental substantive values?" with the answer "equally, in the sense of giving equal moral weight to individuals' comparable interests." A third question concerns scope: "Who is subject to such equal consideration?"

Scope: Sentient Beings

Some ethical theories limit equal consideration to persons or human beings. If such theories address the moral status of animals or nonpersons, they assert that these individuals have either less moral status or no moral status at all. By contrast, as we discuss in Chapter 7, our theory maintains that all beings who have interests – namely, sentient beings – have moral status and that all beings with moral status fall within the scope of equal consideration.

The value of well-being applies to all beings who have a welfare: sentient beings. Thus all sentient beings, in our view, at least deserve equal *consequentialist* consideration. But the value of respect, in our view, is best understood as *respect for rights-holders*. Again, we use the term "rights" somewhat strictly so that rights-holders have moral claims that generally may not be overridden by appeals to utility. Our theory attributes rights of full strength to persons and not to sentient nonpersons. Persons are defined as beings who have "narrative self-awareness" or "narrative capacity" (the ability to understand the parts of their lives as forming a sort of story). Derivatively, we also attribute rights to human beings who are not persons in this sense but are expected to develop into such persons. Moreover, because the normative importance of self-awareness over time is not limited to its appearance in full-blown narrative self-awareness, we attribute *rights of partial strength* to animals who have nontrivial temporal self-awareness that falls short of narrative self-awareness.

The view that all sentient beings are entitled to equal consideration represents a radical departure from other leading works in bioethical theory. One implication, lying outside the purview of bioethics but exceptionally important to everyday choices and agricultural policy, is that modern industrial animal husbandry – “factory farming” – is ethically indefensible. A second broad implication is that the traditional presumption in favor of conducting animal studies before proceeding to clinical trials involving human subjects is morally backward:⁴ there should be a significant presumption *against* involving (sentient) animal subjects in research, in view of their moral status and the fact that nearly all biomedical research involving animals seriously harms them without compensating benefits.

3.3 Mid-level Principles

In Chapters 4–6 we use the method of reflective equilibrium to specify the two values – in light of the formal principle of equality and our views regarding scope – in the form of substantive “mid-level” ethical principles. Here, we describe some key mid-level principles as they relate to non-maleficence, beneficence, distributive justice, and autonomy rights.⁵ Since what we say about them is at a level of specificity that allows moral verdicts to be drawn for particular cases, our specification of the fundamental values into these mid-level principles is a vital part of our overall theory. Thus, we think of the principles not as ones whose content we have *derived* from the two values but as integral components of the theory itself.

Nonmaleficence

The principle of nonmaleficence states that it is *pro tanto* wrong to harm others – meaning that it is wrong to harm others unnecessarily or without sufficient justification. We believe that nonmaleficence is so deeply plausible as to require no justification. Indeed, one would be justified in rejecting any ethical theory that denied that harming others tended to be wrong, on the basis of this denial alone. Although the principle of non-maleficence is virtually self-evident, its scope is not. Our argument that all and only sentient beings have moral status implies that the scope of nonmaleficence is much broader than humanity or the set of persons,

⁴ This assumption is stated in “The Nuremberg Code,” reprinted in *Trials of War Criminals before the Nuremberg Military Tribunals* (Washington, DC: US Government Printing Office, 1948), article 3.

⁵ We discuss them at length in Chapters 4–6.

extending to sentient beings more generally. In view of this scope, the principle of nonmaleficence as we understand it establishes a moral presumption against harming sentient beings: They should not be harmed unless there is a special justification for harming them. Our further argument, that many sentient beings have rights, guides judgments about the sorts of justifications for harm that are acceptable. For rights-holders, showing that an action would provide greater benefits to others is not sufficient justification for inflicting harm. The magnitude of the benefits must exceed a higher threshold to count as justification.

In stating that it is *pro tanto* wrong to harm others, nonmaleficence appeals to the concept of harm. What is harm? As discussed in Chapter 4, the concept is somewhat elusive. After critically evaluating several leading conceptions, we suggest the following definition for use in ethical analysis: *A, an agent or event, harms B if and only if A makes B worse off than B would have been in the absence of the event or A's intervention.* This analysis fits our intuitions about whether or not harm has occurred in a broad range of cases. For example, if A punches B, then B is harmed because he is now worse off than he would have been if B had left him alone. Likewise, suppose B has a chronic condition and was about to receive drugs but A steals them. B may be no worse off than he was before – he is still not getting treatment – but he is worse off than he would have been if A had left him alone, so A has harmed him.

The principle of nonmaleficence underlies various rules of thumb pertaining to distinct types of harm. The following moral rules have their basis in nonmaleficence:⁶

1. Do not cause pain, suffering, or other experiential harm.
2. Do not kill.
3. Do not cause illness, injury, or disability.
4. Do not deprive of goods or opportunities to which the individuals deprived have legitimate claims.
5. Do not impose excessive risk of harm.

Beneficence

Because we believe distributive justice and beneficence are more closely connected than common morality and most prominent theories of

⁶ For ease of formulation, we are leaving out such standard qualifications as “unless there is adequate justification for doing so,” whose purpose is to allow for exceptions to the unqualified rules.

bioethics suggest,⁷ we explore both principles together in Chapter 6. Nevertheless, the *concepts* of beneficence and distributive justice are very different. And there are types of justice other than distributive justice that are important to bioethics. Moreover, beneficence applies in principle to all beings with moral status, whereas distributive justice applies only to rights-holders. So the ethical principles that relate to beneficence and distributive justice merit separate discussion before any connection between them is forged.

Beneficence requires agents to take positive steps to promote the well-being of others. One may promote others' well-being by conferring a benefit upon them (e.g., giving them cash) or by preventing them from being harmed (e.g., by giving them access to effective vaccines or rescuing them from a fire). Sometimes these categories of beneficent measures are difficult to distinguish – for example, when conferring a benefit (e.g., giving food) removes or prevents a harmful condition (e.g., starvation). But just as instances of harming involve making someone worse off than they otherwise would have been, instances of beneficence involve *making someone better off* than they would have been otherwise.

Once we begin to investigate its content more concretely, it becomes apparent that beneficence really splinters into several distinct mid-level principles. We find it helpful to distinguish (1) general beneficence, which is a nonspecific moral obligation to help others in need; (2) a duty to rescue; and (3) special obligations of beneficence that attach to positions within special relationships (e.g., as parents have toward their children and caretakers have toward their pets) or to professional roles (e.g., physician, lawyer, teacher).

General beneficence is the moral obligation to contribute significantly, relative to one's ability and over the course of a lifetime, to assist individuals in need. Potential beneficiaries might include people who are homeless or malnourished, political refugees, victims of human trafficking, and so on. They also include nonhuman animals – such as homeless companion animals and animal victims of organized fighting. General beneficence is what ethicists call an *imperfect* obligation. This means that a moral agent has discretion over to whom and how the obligation is discharged: to what particular causes or individuals, at what particular times, and in what particular forms (e.g., money, volunteer services, blood donations). It is

⁷ See, e.g., Beauchamp and Childress, *Principles of Biomedical Ethics*; H. Tristram Engelhardt, Jr., *The Foundations of Bioethics* (New York: Oxford University Press, 1986); and Bernard Gert, Charles Culver, and K. Danner Clouser, *Bioethics*, 2nd ed. (New York: Oxford University Press, 2006).

appropriate that this obligation leaves such discretion both because there is far more need in the world than any one agent can hope to address and because there are limits to how much one must sacrifice in order to meet this obligation.

In contrast to classical libertarians, we claim that general beneficence is genuinely *obligatory*, not an ideal that is beyond the call of duty. In contrast to common morality, we claim that this obligation is *fairly strong*. At the same time, in contrast to act-utilitarians and other maximizing forms of consequentialism, we claim that the demands of general beneficence are *limited*. That is, we deny that general beneficence requires agents to do everything they can to promote the best results or make the world a better place.

In addition to the imperfect obligation of general beneficence, our ethical theory recognizes *the duty to rescue*, which is a *perfect* obligation. This means that it is morally binding on specific occasions. The duty of rescue requires an agent to provide a benefit to another when the benefit is very large and the agent can do so at a sufficiently low cost to himself. A paradigm scenario featuring the duty to rescue is one in which a lone passerby spots a nearby child, apparently drowning, in a lake into which the passerby can safely wade to rescue the child. Assuming that the costs to the passerby are just inconvenience and soiled clothing, it seems obvious that she has a duty to attempt to rescue the child.

There are further obligations to benefit that are grounded in special relationships or particular professional roles. Parents have special obligations to house, feed, protect, and nurture their children. One generally has stronger obligations to help friends in distress than to help individuals with whom one has no special relationship. You have perfect obligations to feed and protect your pet, get him veterinary attention as needed, and so on. In professional settings, one's role often generates special obligations to benefit. Physicians, for example, have obligations to provide health benefits for their patients and plausibly to strangers too, as when someone needs urgent medical attention on a plane or boat. Teachers have special obligations to provide their students certain educational benefits and to be available to advise them on educational matters. Firefighters have obligations, while on duty, to fight fires within some geographical area.

Let us consider physicians' beneficence-based obligations in greater detail. Doctors clearly have obligations to benefit patients with whom they have a physician-patient relationship. Moreover, if a physician happens

upon a stranger who has just collapsed, she, like any other person who could help, has a duty to attempt to do so (if there is no extraordinary reason why she cannot). Indeed, as a physician, she has a duty to provide medical attention, whereas a layperson may just have a duty to call for assistance. These obligations of beneficence are relatively straightforward.

Matters are less clear when a possible duty of beneficence is very burdensome to discharge, many people need to be rescued, or the reason rescue cases are so costly or so common is that other people are not doing their duty. A physician from a wealthy country may struggle over whether to volunteer to combat a dangerous epidemic in a nation with weak medical and public health infrastructure. In the United States, a physician may have to decide whether to treat indigent, uninsured patients with no expectation of receiving payment. A third physician may live in a part of the world where many people living in slums have serious diseases but could be helped with relatively low-cost medical care.

In the face of such ambiguity about a physician's obligations of beneficence in cases such as these, we emphasize two distinctions: (1) between individual and institutional obligations and (2) between ideal and nonideal background circumstances. Consider again the American physician who must decide whether to treat uninsured patients for free. He faces this dilemma only because of a nonideal background in which the United States, a wealthy nation that can afford to provide universal health care, has failed to meet its obligation to do so. Were the institutional obligations met, US citizens and residents would collectively meet the need of the otherwise-uninsured *indirectly*, by paying taxes that are used by the federal government to ensure universal access to health care.

Given that the United States does not meet this institutional obligation to ensure universal access, does a physician have a beneficence-based obligation to treat uninsured patients who cannot pay? We argue that in such nonideal circumstances those who are in a position to help have much greater obligations than they would in a just world (where the costs of helping would be spread more widely). Not only does a US physician have the duty to attempt to rescue someone who collapses right in her presence; her perfect duties of beneficence require providing a substantial amount of care without charge or at a minimal cost. For example, a physician might charge patients on a sliding scale according to their ability to pay or set aside a certain number of appointments per week in which she provides free care to indigent or uninsured patients.

Distributive Justice

Understood in a highly abstract or formal way, justice may be considered a single moral principle: the principle that requires moral agents to give others *their due*. But, considered so formally, justice provides no actual direction because one needs to have a substantive idea of what different individuals are due. As soon as we turn our attention to substantive mid-level principles of justice, it becomes apparent that there is considerable disagreement about these principles – and, therefore, about what justice actually requires.

Given our commitment to equal consideration, as discussed earlier, our substantive approach to justice must be compatible with equal consideration – that is, with giving equal moral weight to individuals' prudentially comparable interests. This is a very significant commitment when it comes to our dealings with animals, since they are so often treated in ways that give much less than equal consideration to their interests. It is also sufficient to ground a requirement of nondiscrimination among persons – for example, on the basis of race, class, religion, or sexual orientation. Although widely accepted today – at least in liberal democracies – as a requirement of justice, this broad acceptance is a result of hard-won battles in civil rights movements. Moreover, the acceptance in principle does not always correlate with actual practice (e.g., in police's differential treatment of persons of different races), and there continue to be some disputes about what is required in principle (e.g., to instantiate nondiscrimination for transgendered persons). Even so, equal consideration, due to its lack of specificity, is compatible with a broad array of substantive principles of justice as they pertain to persons or human beings.

Let us here distinguish four kinds of justice and clarify the focus of our investigation. *Retributive justice* gives responsible agents what they are due in light of their wrongful acts. As such, this type of justice concerns punishment, a topic that falls outside the purview of this book. *Restorative justice* gives appropriate compensation to individuals who have been wronged so as to “restore” them, in some sense, to their state of well-being prior to being wronged. *Distributive justice*, on which we primarily focus, gives individuals what they are due in the form of benefits (e.g., income, health care access) and burdens (e.g., tax obligations, jury duty) independently of anyone's prior wrongdoing. In morally ideal circumstances, then, neither retributive justice (which responds to agents' having done wrong) nor restorative justice (which responds to individuals' having been wronged) would be relevant. But distributive justice would be

relevant so long as there were benefits and burdens to distribute. Finally, *procedural justice* is a matter of the fairness of the process by which decisions are made.

In our view, only rights-holders have claims on the basis of distributive justice. Though we acknowledge a wider set of rights-holders than human persons, for reasons of space we further restrict our discussion of justice in this book to human persons. We address questions of distributive justice within a single state and also internationally. Domestically, we defend a liberal egalitarian view of distributive justice according to which it is presumptively unjust if one person is worse-off than another person as a result of factors beyond their control. When scarce societal resources are allocated among individuals, we balance two goals: giving higher priority to people who are worse off and maximizing the total benefits that are distributed. Thus our distributive principle is a form of *moderate prioritarianism*.

The differences in life prospects between individuals in different countries are vast. For example, life expectancy for a child born in Sierra Leone in 2018 was fifty-four years. For a child born in Japan it was eighty-four years.⁸ We regard these unchosen differences as problematic in just the same way as differences within a country. We therefore defend a *cosmopolitan* view of global distributive justice: the principles of justice apply in the same way across states as they do within states.

Here we note two important implications of our approach to justice. First, we argue that people have a right to access affordable health care. As a matter of distributive justice, this right is limited in certain ways. In particular, it is limited by the availability of resources, so that individuals do not have a right to every intervention that would be beneficial – moderate prioritarianism applies to the distribution of health care resources too. Our views about global justice also imply that richer countries (and richer individuals) have substantial duties to provide poorer countries with the means to provide their people with adequate health care. Second, we argue that the current international intellectual property system that grants twenty-year patents on novel pharmaceuticals is unjust. The international community has an obligation to permit poorer countries to purchase cheaper, generic versions of patented medicines and to adopt different methods for incentivizing innovation in medicine.

⁸ The data come from the World Bank, *World Bank Open Data* (available at <https://data.worldbank.org/indicator/SP.DYN.LE00.IN>; accessed September 22, 2020).

Autonomy Rights

Autonomy means self-rule. More specifically, autonomous or *competent* individuals are capable of deliberating about their options in light of their own values and priorities, reaching a decision on the basis of such deliberation, and acting accordingly.

Autonomy is of special importance because of its close connection with both well-being and respect for rights-holders. When a competent person exercises their capacity for autonomous action, doing so tends to promote their well-being. Autonomy has *instrumental* value insofar as competent persons tend to know their interests better than other people do, with the consequence that competent persons' self-governance tends to promote their well-being more effectively than paternalistic interventions into their affairs. Autonomous action can also itself be a source of well-being, since many people desire to and enjoy being able to act free of others' control. The value of autonomy in terms of its contribution to well-being can also ground obligations to assist persons to actualize their capacity for autonomous decision-making and action. For example, a mental health professional may have an obligation to foster their patient's confidence, clarity of thinking, and other capacities that enable autonomous choice. Likewise, good parents recognize an obligation to nurture their children's development into autonomous adults.

The capacity for autonomous action is also the ground of autonomy rights, such as the rights to bodily control that allow competent persons to give or refuse consent to medical interventions. Strictly speaking, "respect for autonomy" means respecting the exercise of these rights. Other things being equal, it is wrong to interfere with competent persons' decision-making or control their action through deception, motivational manipulation, or coercion. As we discuss in Chapter 5, awareness of these threats to autonomous decision-making helps to illuminate the conditions of valid consent and of appropriate surrogate decision-making.

Competent individuals may or may not act autonomously in any given case. We analyze autonomous action in this way: *An agent A performs action X autonomously if and only if (1) A performs X (i) intentionally, (ii) with sufficient understanding, (iii) sufficiently freely of controlling influences; and (2) A decided, or could have decided, whether to X in light of A's values.* Condition (2) implies that only beings who have values can act autonomously. Many beings who can act intentionally on the basis of their desires lack the capacity to stand back and evaluate their desires in light of values. A bird, for example, might fly intentionally to her nest on the basis of her desires and perceptions, but she does not thereby fly autonomously.

To head off a possible confusion: it is not necessary, for an action to be autonomous, that the agent actually reflect on and endorse the action in terms of the agent's own values; many or most autonomous actions in everyday life are not preceded by such reflection. What *is* necessary for someone's action to be autonomous is that the agent be capable of such evaluation. For example, Dan may choose to eat a bowl of sugary cereal with milk and do so autonomously but without any reflection. If someone were to point out an inconsistency between his choice of breakfast and his values (relevantly here: good nutrition and abstaining from factory farm products), Dan could change his behavior to bring his choice in line with his values. If Jeri is addicted to nicotine and wishes she did not have a desire to smoke, considering the habit contrary to her value of healthful living, she may smoke intentionally on the basis of her desire, but her smoking may be compulsive rather than autonomous.

Many human beings lack the capacity for autonomous action due to immaturity or substantial cognitive incapacity, yet are capable of performing intentional actions in accordance with their desires and beliefs. The same is true of many animals. We use the term *nonautonomous agents* to refer to all beings – whether human or nonhuman – who can act intentionally but not autonomously. Although the choices of nonautonomous agents do not have to be respected as a matter of rights, these individuals can still have an interest in liberty and freedom from controlling influences. Other things being equal, an eight-year-child has an interest in freedom of movement, as does a dog – it is good for children, dogs, and other nonautonomous agents to enjoy themselves as they see fit and do things they want to do. Yet, frequently, other things are not equal. Considerations of well-being require restricting children's and dogs' liberty in order to keep them safe from potential kidnappers or cars on the highway. For nonautonomous individuals, this sort of paternalistic interference is justified when the benefits outweigh the harms.

Contrast with Principles of Biomedical Ethics

Because the mid-level principles we have described fall under similar categories to the four principles featured in Tom Beauchamp and James Childress's prominent textbook, *Principles of Biomedical Ethics*, it may be helpful to identify some of the main respects in which our treatment of these principles differs from the approach of Beauchamp and Childress.

Our chapter on nonmaleficence, unlike the corresponding chapter in *Principles*, includes an extensive exploration of the concept of harm,

canvassing several theoretical options and issues before arriving at our own analysis.⁹ In addition, differences between Beauchamp and Childress's (partly sketched) model of moral status and our own model of moral status entail differences in how nonmaleficence is to be interpreted in relation to sentient animals. One clear difference is that, whereas Beauchamp and Childress assert degrees of moral status such that some beings with moral status deserve less than equal consequentialist consideration,¹⁰ our commitment to equal consideration denies this assertion. Consequently, there is a stronger presumption against harming sentient animals on our view than on theirs.

There are several important differences between our approach to beneficence and justice and that of Beauchamp and Childress.¹¹ Unlike these authors, we find beneficence and distributive justice to bear significant overlap in their substantive moral demands. This is related to the fact that we acknowledge moderately strong obligations of general beneficence (as explained above), whereas Beauchamp and Childress remain neutral on the strength of such obligations.¹² Relatedly, Beauchamp and Childress treat classical libertarianism, which expressly denies that people have general obligations of beneficence, as one among several respectable theories of justice; by contrast, we reject and attempt to refute libertarianism. The principle of distributive justice at which we arrive is a form of moderate prioritarianism, whereas Beauchamp and Childress embrace several distributive principles without clarifying whether they can be integrated.

Finally, our treatment of autonomy differs from Beauchamp and Childress's approach in several ways. Although our analysis of autonomous action significantly converges with theirs,¹³ only ours requires that agents have the ability to evaluate their prospective actions in terms of their own values. This feature of our analysis substantially limits the class of agents who are capable of acting autonomously – as opposed to merely intentionally – and coheres with our ascription of rights only to beings with substantial self-awareness. In addition, we analyze the conditions for valid consent – arguably the most central autonomy-related concept in bioethics – differently from Beauchamp and Childress. For example, whereas they, like many other commentators, require significant comprehension as

⁹ Their analysis of the concept of harm is limited to four paragraphs and discusses neither recent nonstandard analyses nor challenges to standard analyses such as theirs (*Principles of Biomedical Ethics*, 153–154).

¹⁰ *Ibid.*, 82–85.

¹¹ *Ibid.*, chap. 7.

¹² *Ibid.*, 203–206.

¹³ *Ibid.*, 104.

a condition of valid consent,¹⁴ we do not, instead placing great emphasis on adequate disclosure and the avoidance of deception.

3.4 Other Ethical Theories and Concepts

It is common today to divide ethical theories into three broad categories: consequentialism, deontology, and virtue ethics. Our approach may be understood in terms of consequentialism and deontology, both of which rely on principles that prescribe certain actions. However, in addition to specifying action-guiding norms, we can also identify moral virtues that correspond to some of these norms. As we understand moral virtues, their primary role is to support moral conduct by complementing agents' ability *to know what actions are right* (such knowledge being facilitated by such action-guides as principles and rules) with the strength of character that makes them more likely *to act accordingly*. Here we describe the key features of these three categories of ethical theory and note how our theory relates to them.

Consequentialism

Consequentialism is the class of ethical theories that converge on the general idea that *right action is that which is expected to produce the best (or good enough) results*. Utilitarianism is the most prominent type of consequentialist theory. It features a single supreme principle: the principle of utility. The principle of utility directs agents to act in ways that can be reasonably expected – either directly or via utility-promoting rules – to maximize well-being or utility. Important utilitarian thinkers in bioethics include Joseph Fletcher, a theologian and pioneer of the newly recognized discipline of bioethics, and the renowned philosopher Peter Singer.¹⁵ We agree with utilitarians that the well-being of individuals is a central value to be promoted, but disagree that it is the only general value at the heart of ethics. For this reason, we cannot commit to the principle of utility, which directs agents to *maximize* – rather than merely promote – well-being or utility.

Consequentialist theories differ along two dimensions. First, they differ regarding how consequences matter morally. For example, in contrast to

¹⁴ Ibid., 124.

¹⁵ See, e.g., Joseph Fletcher, *Situation Ethics* (Louisville, KY: Westminster John Knox Press, 1966); and Peter Singer, *Rethinking Life and Death* (New York: St. Martin's, 1994).

utilitarian theories, other types of consequentialist theories either interpret the value at stake in the “best results” in different terms (for example, taking into account both well-being and priority to the worst-off), moderate the demand for the *best* results so that an agent is required to do *enough* to promote well-being rather than having to maximize it, or both.

Second, consequentialist theories differ in terms of how they think the norms that guide our behavior should take into account the consequences that ultimately matter. *Direct* consequentialism asserts that agents ought to act in such a way that can be expected, on each occasion, to produce the best (or, in some versions, good enough) results. Direct utilitarianism – which is often called *act*-utilitarianism – instructs the agent to act in ways that can be expected to maximize well-being. This directive, however, would seem to justify some actions that conflict with our considered judgments. For example, it would apparently justify the discreet murder of a hospital patient in order to salvage his organs to save several other individuals. Further, direct utilitarianism would justify the most barbaric torture of an animal if the enjoyment of a large number of sadistic spectators outweighed the harm to the animal and no other activity could be expected to offer a greater gain in overall well-being. These apparent implications of direct utilitarianism contradict our considered judgments that killing an innocent person as a means to save others wrongfully violates the victim’s rights and that torturing an animal for fun cannot be justified. By contrast, *indirect* consequentialists believe that in the long term the good is best promoted by complying with certain rules and constraints.¹⁶ These might include a prohibition on killing the innocent without their consent and a rule against tormenting animals for entertainment. Even if the immediate consequences of murdering an innocent appear to be positive, on balance, given human nature it would ultimately lead to worse consequences if we endorsed making these judgments on a case-by-case basis.¹⁷

Given its emphasis on individual rights, our ethical theory is inconsistent with direct forms of consequentialism. It might, however, be consistent with a form of indirect consequentialism. Consequentialism offers a principled basis for identifying the rare justified exceptions to well-fashioned rules and their corresponding rights that we accept. To return to an earlier example, it tends to promote well-being and the best results for society if we grant people a right to freedom of movement in public

¹⁶ For a recent example, see Brad Hooker, *Ideal Code, Real World* (Oxford: Clarendon, 2000).

¹⁷ See R. M. Hare, *Moral Thinking* (Oxford: Clarendon, 1981).

spaces, but this right is appropriately overridden whenever imposing a quarantine on infectious individuals is justifiable – and it is plausible that the basis for overriding the right is utility or the public welfare. A further reason to think our approach is consistent with indirect consequentialism is that the consequentialist commitment to promote well-being can straightforwardly justify the moderately strong imperfect obligation to help those in need that our theory embraces. Our hesitation in asserting that our ethical theory is consistent with indirect consequentialism is the fact that this assertion depends on the speculative empirical claim that general acceptance of and compliance with the norms we defend would, in fact, be conducive to the best results in the long term.

Deontology

Deontology is the broad class of ethical theories that agree with consequentialism (against virtue ethics) that right action is the most central concept in ethics but disagree with the consequentialist thesis that criteria for right action concern only its results. Deontologists hold that moral duties constrain what we are permitted to do, even when our actions would produce the best results. For example, our duty not to kill innocent people without their consent means that we may not murder one person even if we expect thereby to save several people's lives. Deontologists also generally hold that morality includes "options" that permit agents to pursue their own projects and interests, rather than do everything they can to bring about the best results.¹⁸ Thus, deontology is *stricter* than consequentialism in that it sets constraints on the pursuit of the best results or any other ends. Deontology is *more permissive* than utilitarianism – and other types of consequentialism that require acting so as to bring about the *best* results – in not requiring agents to do all they can to make the world a better place. Frances Kamm is an important contributor to bioethics whose thinking is distinctively deontological.¹⁹

The most influential deontological theory derives from the work of Immanuel Kant.²⁰ Kantian ethics features a single supreme principle, the Categorical Imperative. One way to understand the substance of the Categorical Imperative's content is in terms of respect: one must always

¹⁸ See, e.g., Samuel Sheffler, *The Rejection of Consequentialism* (Oxford: Oxford University Press, 1982).

¹⁹ See, e.g., Frances Kamm, *Bioethical Prescriptions* (New York: Oxford University Press, 2013).

²⁰ *Grounding for the Metaphysics of Morals*, 3rd ed., trans. James Ellington (Indianapolis, IN: Hackett, 1993; first published 1785).

act in ways that are consistent with respecting persons and never in ways that treat them as mere means to an end. Moreover, the fact that the Categorical Imperative does not require one to do everything one can to promote well-being – since one can live in compliance with this principle while devoting much time and energy to one's own personal projects – might also be understood in terms of respect. At issue here is respect for a moral agent, who has sovereignty over their own life and may pursue their own ends, consistent with the Categorical Imperative, rather than facing a never-ending obligation to make the world a better place. An important neo-Kantian contributor to bioethics is Onora O'Neill.²¹

We agree with Kantians that respect is a central value at the heart of ethics. We disagree with Kantians by denying that respect constitutes a supreme principle that always overrides the promotion of well-being in cases of conflict. Moreover, we recast respect for persons as respect for *rights-holders*, a conception that leaves open whether rights-holders are limited to persons, include all sentient beings, or comprise a range of beings wider than the set of persons but narrower than the set of sentient beings. We defend a version of the latter, intermediate view of rights-holders.

Deontology, as noted, sets constraints on the pursuit of ends but allows moral agents considerable freedom to decide what to do within those constraints. A modern example of a *pure* deontology might feature absolute rights, which correspond to absolute rules binding all moral agents. For example, libertarians, such as bioethicist Tristram Engelhardt,²² assert rights to life, liberty, and property that cannot be overridden by appeals to utility. Unlike libertarians, neo-Kantians accept that individuals have an imperfect obligation to help those in need, but they generally assert some absolute constraints such as against killing innocent persons without their consent. Our theory, in contrast, holds that nearly all such constraints have thresholds such that they may be overridden if the gain in overall well-being from doing so is sufficiently great. In brief, constraints, yes, but not absolute ones.²³ Hence our view is a form of *moderate* deontology. Further, because our approach embraces stronger obligations to promote the good than many deontologists acknowledge, our deontology is

²¹ See Onora O'Neill, *Autonomy and Trust in Bioethics* (Cambridge: Cambridge University Press, 2002).

²² See Engelhardt, *Foundations of Bioethics*.

²³ As we note in Chapter 4, in practice we regard certain rights – including rights against torture, rape, and enslavement – as absolute. There are, we maintain, no actual situations in which these acts are justified by the expectation of uniquely valuable consequences.

moderate in terms of how much freedom it gives persons to pursue their own projects rather than assist others in need.

Virtue Ethics

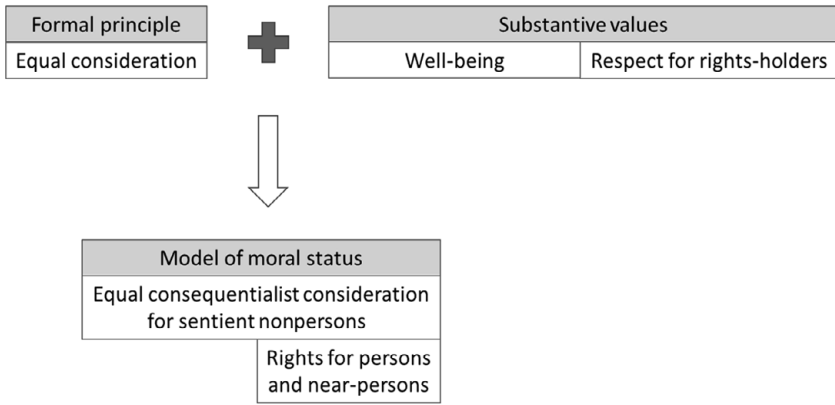
Virtue ethicists take moral character, rather than the results of actions or their consistency with moral rules, to be fundamental to ethics. Virtues are character traits such as courage, honesty, and generosity. A virtuous agent is someone who possesses these character traits such that they recognize which situations call for, say, courage – as opposed to being timid or overly rash – and are motivated to act accordingly. Edmund Pellegrino was an important contributor to bioethics who represented the tradition of virtue ethics.²⁴

Though our view does not take virtues as foundational, we acknowledge that they are vital parts of an ethical life. An ethical theory like ours that captures a defensible normative structure in the form of principles, rules, obligations, and rights can help moral agents identify the right action to perform or the right policy to support. But knowing what is right is one thing, doing it quite another. It is possible to know what is right and fail to do it due to insufficient motivation, weakness of will, being overwhelmed by peer or institutional pressure to do something else, or major character defects – that is, moral vices. *Moral virtues*, as we understand them, are character traits that facilitate right conduct. Also important are *intellectual virtues*, traits of mind that facilitate good thinking, including good ethical thinking. One might think that mastery of a good theory would make intellectual virtues unnecessary in ethics, since the theory would tell you the right action or policy. This reasoning would miss the important point that any theory needs to be *applied*, which requires intellectual work (e.g., “Would saying X be deceptive?”), and sometimes needs to be *interpreted* (e.g., “What does equal consideration for animals imply about the permissibility of killing them?”). Both moral virtues and intellectual virtues are crucial to living ethically.

Which traits of character and mind are virtues? Beginning with traits of character – moral virtues – we can distinguish two types. *First-order virtues* have a moral content that guides action. They include respectfulness, benevolence, compassion, honesty, discretion, fairness, and passion for justice. *Second-order virtues* do not directly guide action. They function

²⁴ See, e.g., Edmund Pellegrino, “Toward a Virtue-Based Normative Ethics for the Health Professions,” *Kennedy Institute of Ethics Journal* 5 (1995): 253–277.

1. Fundamental values



2. Specification

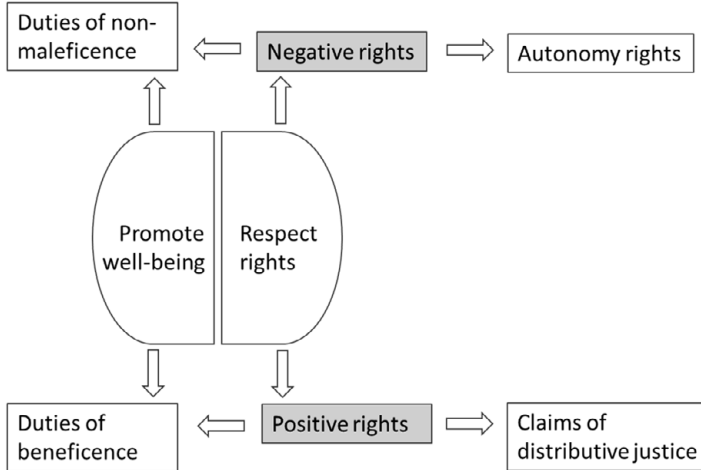


Figure 3.1 The dual value theory.

to make it more likely that one will do the right thing when it is difficult to do so. These virtues include conscientiousness or integrity, moral courage, mindfulness (which counteracts impulsivity), and strength of will. Meanwhile, intellectual virtues or traits of mind that are important to ethical living include intellectual honesty, open-mindedness, clarity of thinking, perceptiveness, the capacity to master and remember salient information, and sound judgment or practical wisdom. We believe that the importance of both second-order moral virtues and intellectual virtues in ethical living tends to be underappreciated.

This completes a sketch of our ethical theory. A fuller picture will emerge with the details of later chapters. Figure 3.1 gives a pictorial representation of the theory.