

*Introduction***1.1 Why This Book?**

Bioethics is a young discipline. It has been recognized as a distinct area of scholarly investigation only since the late 1960s or early 1970s. The term “bioethics” apparently came into use in the early 1970s, with recognition of the discipline in the United States closely tied to institutional support through the founding of the Hastings Center in Hastings-on-Hudson, New York, in 1969 and the Kennedy Institute of Ethics at Georgetown University in Washington, DC, in 1971.<sup>1</sup> The Centre for Bioethics at the Clinical Research Institute of Montreal was established in 1976 as the first Canadian center focused on bioethics.<sup>2</sup> Driving discussions in the 1970s were such topics as the determination of human death, the ethics of birth control and abortion, fair allocation of kidney dialysis machines, the permissibility of forgoing life support for impaired newborns, and the rights of human research subjects. These debates took place in the press, in medical journals, and in new specialist academic journals, including the *Hastings Center Report* in the United States (from 1971) and the *Journal of Medical Ethics* in the United Kingdom (from 1975).<sup>3</sup> By the end of the decade the young discipline had its own encyclopedia.<sup>4</sup>

As we use the term, “bioethics” refers to the study of ethical issues that arise in medicine, in such allied fields as nursing, pharmacy, and public health, and in the life sciences. A theory of bioethics is a general framework for illuminating and ultimately addressing ethical issues that arise in

<sup>1</sup> See Warren Reich, “The Word ‘Bioethics’: Its Birth and the Legacies of Those Who Shaped It,” *Kennedy Institute of Ethics Journal* 4 (1994): 319–335; and Robert Martensen, “The History of Bioethics: An Essay Review,” *Journal of the History of Medicine* 56 (2001): 168–175.

<sup>2</sup> John Williams, “The Influence of American Bioethics in Canada,” *Journal International de Bioéthique* 20 (4) 2009: 95–105.

<sup>3</sup> For a history of bioethics in the United Kingdom, see Ruth Chadwick and Duncan Wilson, “The Emergence and Development of Bioethics in the UK,” *Medical Law Review* 26 (2018): 183–201.

<sup>4</sup> Warren Reich (ed.), *The Encyclopedia of Bioethics* (New York: Macmillan, 1978).

bioethics. Is euthanasia, or medical mercy killing, ever justified? Do people have a moral right to health care? Is it permissible to involve animals in experiments that seriously harm them in order to benefit humanity? A theory of bioethics will offer a set of ethical guidelines – and perhaps an explicit method for deploying them – to help people address such questions in a manner that (according to the theory) is likely to yield justified or correct answers.

There are a variety of bioethical theories already on offer. Some are tailored to the content of bioethics or, more specifically, medical ethics. Examples include the principle-based approach to bioethics associated with Tom Beauchamp and James Childress,<sup>5</sup> Tristram Engelhardt's libertarian bioethics,<sup>6</sup> and Robert Veatch's contract-based theory of medical ethics.<sup>7</sup> Some theories employed in bioethics are *ethical* theories – theories offering guidelines for addressing ethical questions in general rather than bioethical issues in particular. Examples of ethical theories that have been put to work in bioethics are utilitarianism,<sup>8</sup> Kantian ethics,<sup>9</sup> a rules-based common morality theory,<sup>10</sup> and virtue ethics.<sup>11</sup> Other theoretical approaches to bioethics are distinguished less by the content of their central ethical norms than by their methods for investigating ethical questions. Examples include casuistry, a type of case-based reasoning in historical context;<sup>12</sup> feminist ethics, which interrogates mainstream ethical thinking and theories with an eye toward exposing gendered and oppression-supporting assumptions;<sup>13</sup> and narrative bioethics, which views the exploration of stories as a means to ethical insight.<sup>14</sup>

Despite the richness of current offerings in bioethical theory, we perceive a substantial gap in the literature. A satisfying bioethical theory, in

<sup>5</sup> Tom Beauchamp and James Childress, *Principles of Biomedical Ethics*, 8th ed. (New York: Oxford University Press, 2019; first published 1979).

<sup>6</sup> H. Tristram Engelhardt, Jr., *The Foundations of Bioethics* (New York: Oxford University Press, 1986).

<sup>7</sup> Robert Veatch, *A Theory of Medical Ethics* (New York: Basic, 1981).

<sup>8</sup> See, e.g., R. M. Hare, "A Utilitarian Approach," in Helga Kuhse and Peter Singer (eds.), *A Companion to Bioethics*, 2nd ed. (Oxford: Blackwell, 2009), 85–90.

<sup>9</sup> See, e.g., Onora O'Neill, *Autonomy and Trust in Bioethics* (Cambridge: Cambridge University Press, 2002).

<sup>10</sup> Bernard Gert, Charles Culver, and K. Danner Clouser, *Bioethics* (New York: Oxford University Press, 1997).

<sup>11</sup> See, e.g., Justin Oakley, "Virtue Ethics and Bioethics," in Daniel Russell (ed.), *The Cambridge Companion to Virtue Ethics* (Cambridge: Cambridge University Press, 2013), 197–220.

<sup>12</sup> See, e.g., John Arras, "Getting Down to Cases: The Revival of Casuistry in Bioethics," *Journal of Medicine and Philosophy* 16 (1991): 29–51.

<sup>13</sup> See, e.g., Susan Wolf (ed.), *Feminism and Bioethics* (New York: Oxford University Press, 1996).

<sup>14</sup> See, e.g., Hilde Lindemann Nelson (ed.), *Stories and Their Limits* (New York: Routledge, 1997).

our estimation, would achieve several aims. First, it would provide a high-quality discussion of ethical theory and methodology in ethics, recommending an approach that holds up well under critical scrutiny. Second, an adequate bioethical theory would avoid the narrowness of normative vision that one finds in some theories – with excessive focus on, for example, hypothetical agreement in contract theories, on liberty in libertarianism, and on moral rules in the rules-based common morality approach. Third, a fully adequate contribution to bioethical theory would probe and integrate areas of philosophical theory that are relevant to ethics but tend to receive little coverage in the bioethical theory literature. These include the nature of harm, the nature of well-being, models of moral status, personal identity theory, and the “nonidentity problem.” We are not aware of any previous book on bioethical theory that achieves all three aims. In pursuing them, we think we have arrived at a normative vision of bioethics that is wider in scope, but is also more progressive than many works in bioethical theory. For example, our theory treats nonhuman animals with the moral seriousness they deserve and accepts extensive institutional and individual obligations to the global poor.

This book is intended for several overlapping audiences. It is intended for upper-level undergraduate and graduate classes in bioethics, applied ethics, or ethical theory. It is intended for scholars from any discipline who are interested in these areas. Most generally, this book is intended for readers – ranging from those with little theoretical background to specialists – who are interested in bioethical theory or a vision of normative bioethics that covers a broad array of important issues.

Some limitations are inevitable for a book that attempts to cover so much ground. In particular, we acknowledge geographical restrictions in the scope of many discussions of specific issues in the book. Our theory is intended to be universal – correct (or incorrect) everywhere. But its application must be sensitive to context. We lack the expertise to speak to the contexts in which many health care practitioners operate. We therefore mostly draw our examples from the Anglophone high-income countries with which we are most familiar: Canada, the United Kingdom, and the United States. In addition, although our book discusses the relationship between our theory and other broad types of ethical theory, it does not offer a thorough introduction to ethical theory. Finally, while we have made every effort to achieve accessibility while seeking philosophical depth and precision, our book is unlikely to be suitable for high school students and perhaps for lower-level undergraduates.

## 1.2 Plan for the Book

Each of the nine chapters that follow this introduction engages theoretical issues. Each of Chapters 4–10 also includes two or three substantial applications of the theoretical material developed in the chapter to concrete issues of bioethical interest.

Chapter 2 presents our methodology for bioethics. A methodology, although theoretical, is not itself a theory. A theory of bioethics features *substantive* norms, or action-guides, such as “Respect the informed, voluntary choices of autonomous decision-makers” and “Do not deceive patients or prospective research participants.” These norms have ethical content. A methodology for bioethics, by contrast, is a structured *process* for arriving at such norms. One methodology, sometimes called “deductivism,” recommends starting with (purportedly) self-evident or rationally provable general ethical principles and then, in view of relevant facts, deriving specific conclusions about right conduct. A second methodology, featuring a type of inductive reasoning, involves carefully examining a variety of specific cases, intuitively judging the right answer in each case, and then generalizing on the basis of these specific judgments to more general principles, which can then be applied to new cases. Our methodology is distinct from each of these.

The method we embrace has been called both “reflective equilibrium” and “the coherence model of ethical justification.” It is neither purely deductive (justifying specific moral judgments on the basis of general principles) nor purely inductive (justifying general norms on the basis of confident judgments about specific cases). Instead, it provisionally accepts “considered moral judgments” – judgments taken to be especially reliable in virtue of their inherent plausibility, stability, and low likelihood of being motivated by biases – at any level of generality. Some considered judgments will be very general. For example, the principle of nonmaleficence – which states that it is wrong to harm others in the absence of special justifying circumstances – is a very general considered judgment. Others will be judgments about a specific case. For example, the judgment that the act of running a particular dog fight is wrongful is a specific considered judgment about a particular case. Still others will be of intermediate generality. For example, the judgment that rape is wrong is a considered judgment of intermediate generality. According to the method of reflective equilibrium, a moral judgment is justified if it is part of an overall view of ethics that we could accept, upon reflection, on the basis of its incorporation of considered judgments, its overall plausibility and coherence, and

various other theoretical virtues such as simplicity, comprehensiveness, and explanatory power. The discussion of Chapter 2 explains how the critical tools of this method entail that what we intuitively believe about an ethical issue may sometimes have to be rejected. For instance, this will be the case with intuitive judgments that seem prejudiced or misinformed and fail to cohere with judgments that persist in reflective equilibrium. The methodology is therefore capable of generating radical and surprising conclusions, even though it starts with our existing moral beliefs.

Chapter 3 presents an overview of the ethical theory that emerges from our application of the method of reflective equilibrium. Our theory acknowledges two fundamental and irreducible values: well-being and respect for rights-holders (or “respect” for short). Accordingly, we refer to it as our *dual value theory*. Some prominent alternative ethical theories are grounded in a single value. For example, utilitarianism may be understood as being grounded in utility or well-being alone. Kantian ethics may be regarded as grounded entirely in respect for persons. Bernard Gert’s rule-based common morality approach appears to be grounded entirely in nonmaleficence, insofar as all of its ten rules require the avoidance of some kind of harm.<sup>15</sup> Other theories are pluralistic. For example, W. D. Ross’s theory of prima facie duties features multiple moral obligations that are treated as irreducible to any more basic norm<sup>16</sup> while William Frankena’s ethical theory features two ultimate principles, beneficence and justice.<sup>17</sup>

The two fundamental values in our theory, well-being and respect for rights-holders, inform the theory’s scope – that is, who has moral status. Our answer is: all beings who have a welfare have moral status. This means all sentient beings – beings who are capable of having pleasant or unpleasant experiences. All sentient beings are entitled to a form of equal moral consideration. We understand such equal consideration in consequentialist terms, meaning that the well-being of some sentient beings may be traded off for the well-being of others. Both equal consideration and the permissibility of consequentialist trade-offs are important in considering obligations of nonmaleficence as they apply to nonhuman animals. Equal consideration entails a moral prohibition against many of the harms that humans have traditionally felt entitled to impose on animals. The prerogative of trade-offs permits sacrifices of some animals’ well-being for the greater good.

<sup>15</sup> See Bernard Gert, *Morality* (New York: Oxford University Press, 1988).

<sup>16</sup> W. D. Ross, *The Right and the Good* (Oxford: Oxford University Press, 1930).

<sup>17</sup> William Frankena, *Ethics* (Englewood Cliffs, NJ: Prentice-Hall, 1963).

In our theory, in addition to having obligations of nonmaleficence, and sometimes beneficence, to sentient beings, moral agents have distinct obligations to rights-holders. Rights, as we understand them, protect individuals from consequentialist trading-off of their interests for the greater good. These protections are not absolute and so may sometimes be overridden when the consequences of doing so are sufficiently important. In our theory, those who have full-strength rights are individuals with *narrative capacity*: the capacity to form narrative identities or temporally structured mental stories about their own lives. Beings who lack this narrative capacity, but still have some significant self-awareness over time, have rights of partial strength. This group likely includes animals such as dogs and monkeys. While rights-holders, in our approach, enjoy special protections in connection with obligations of nonmaleficence, they also have entitlements related to the principles of distributive justice.

Chapter 4 examines the principle of nonmaleficence, which states a prohibition on causing harm to others in the absence of justifying circumstances. The chapter begins with a type of theoretical exploration that is generally lacking in other books on bioethical theory: an investigation of the nature of harm. After surveying several leading accounts of the nature of harm, and noting challenges to each, it defends a counterfactual account: you harm someone if and only if you make them worse off than they would have been in the absence of your intervention. Next we specify nonmaleficence into several general moral rules corresponding to ways in which individuals can be harmed.

Following these foundational reflections, the chapter explores three areas of practical ethical concern in which rules concerning harm figure prominently: (1) the ethics of torture (an important, instructive issue even if not squarely within bioethics), (2) the limits of permissible risk in pediatric research, and (3) the ethics of medical assistance-in-dying. We find that nonmaleficence supports a right not to be tortured that should never, in practice, be overridden, and children's right to adequate protection, which allows children to be exposed to some net risks for the sake of valuable scientific knowledge, while placing a ceiling on this risk. In exploring medical assistance-in-dying we find a conflict, in some circumstances, between two rules: "Do not cause pain, suffering, or other experiential harm" and "Do not kill." We argue for the permissibility of physician-assisted suicide and euthanasia in some cases but also for an array of safeguards to protect against error and abuse. We find that death is not always a harm. Perhaps surprisingly, whether or not death is a harm for an autonomous individual, given their circumstances, is partly determined by their values.

Chapter 5 analyzes autonomy. The concept of autonomy has played a pivotal role in bioethics. Recognition of the importance of patient autonomy – in relation to informed consent, patient rights, and the value of people making their own decisions about medical care – has transformed medical practice and clinical research, distinguishing contemporary medical ethics from the far more paternalistic medical ethics that preceded it. According to our analysis, an agent, A, performs a particular action autonomously if and only if (1) A performs the action (a) intentionally, (b) with sufficient understanding, and (c) sufficiently freely of controlling influences, and (2) A decided, or could have decided, whether to perform the action in light of A's values.

The importance of autonomy, we argue, may be understood both in its contribution to individual well-being and in terms of the intrinsic moral importance of an individual's sovereignty over their own life. We discuss how autonomy grounds certain rights and then construct a taxonomy of ways in which someone's autonomy can be interfered with. The chapter briefly explores two justifications for interfering with someone's autonomous actions: paternalistic justifications and the prevention of harm to others. Autonomous individuals have the power to waive some of their rights by giving consent. Because this is a crucial concept for bioethics, we identify necessary and sufficient conditions for valid consent. Our analysis of valid consent departs from most others in the literature in the way it understands the condition of comprehension.

When someone lacks competence to make their own decisions, someone else must decide for them. This brings us to the topic of surrogate decision-making, where we introduce and defend a novel "reasonable subject" standard. The chapter's final two sections take up practical applications of our theoretical reflections: the right to refuse medical treatment and the ethics of direct-to-consumer advertising of pharmaceuticals.

Chapter 6 explores both distributive justice and beneficence – which we believe to be more closely related than is generally appreciated. Justice involves giving individuals what they are due. *Distributive justice* governs the distribution of valuable resources (e.g., income), the distribution of burdens (e.g., taxes), and the granting of certain legal rights (e.g., the right to marry). Beneficence concerns agents' duties to benefit other individuals. The *imperfect duty of beneficence* is a duty to contribute substantially, relative to one's ability, to assist individuals in need over the course of one's life. We consider it an advance over much prominent work in bioethical theory that our theory unequivocally supports such a duty. The *perfect duty of beneficence* or *duty of rescue* is a duty agents have to

provide large benefits to others when they can do so at relatively low cost to themselves. There are, in addition, *special duties of beneficence* that attach to agents in virtue of their roles and relationships, such as the special duties of clinicians to their patients or parents to their children.

The chapter highlights two crucial distinctions: (1) between the ideal and the nonideal and (2) between how institutions should be arranged and how individuals should act. *Ideal theory* concerns the organization of just social and international institutions. It tells us how individuals ought to act against a background of just institutions and on the assumption that other people will act rightly. *Nonideal theory* concerns what ought to be done when the institutional background is not just and other people cannot be relied upon to act as they should. Most saliently, we understand nonideal theory to address what particular actors – both states and persons – should do in the world as it is now.

Regarding institutions, domestically, we defend a relatively generic *liberal egalitarian* view about distributive justice: unchosen differences in individual advantage within a society are *prima facie* unjust. Justifications for differences in individual advantage might include a great gain in overall well-being, the need for such differences to secure a fundamental right, or the fact that the inequality results from voluntary, informed decisions. Globally, we endorse a form of *cosmopolitanism*: similar principles of justice apply internationally as apply domestically. Regarding individuals' obligations in our nonideal world, we defend extensive duties of beneficence, albeit consistent with considerable leeway for people to prioritize their own projects.

The practical consequences of our theoretical views for institutional arrangements and the obligations of individuals are far-reaching. Among other implications, we argue that national governments should ensure that all their residents have access to affordable health care and that the international community ought to amend the global intellectual property regime that governs pharmaceutical patents.

As discussed in Chapters 4–6, morality generates obligations related to nonmaleficence, autonomy, distributive justice, and beneficence. But to whom are these obligations owed? Who has rights that correspond to such obligations? To address these questions is to engage the concept of *moral status*. Chapter 7 examines moral status in depth.

The discussion begins with the concept of moral status, formally unpacking its elements before commenting on its usefulness. It proceeds to a sketch of our account of moral status. Our account embraces equal consequentialist consideration for all sentient beings while ascribing the



stronger protection of rights to those with narrative capacity (full-strength rights) and to those with nontrivial temporal self-awareness that falls short of narrative capacity (partial-strength rights). This account of moral status is neutral with respect to species in the sense that membership in *Homo sapiens* is, in itself, neither necessary nor sufficient for moral status or rights.

The final three sections explore ethical implications for research involving human embryos, rodents, and great apes. We defend a very liberal position with respect to embryo research, a relatively restrictive approach to rodent research (one that accords equal consequentialist consideration to rodents' interests while permitting their use on utilitarian grounds), and a prohibition of invasive, nontherapeutic research involving great apes.

Chapter 8 explores the nature of individual well-being. The chapter examines subjective value theories, which understand well-being in terms of the experiences or judgmental authority of the individual subject, and more objective theories, which understand individual well-being partly in terms of factors that are independent of the subject's experiences or authority. We then sketch our preferred approach, a type of subjective theory.

According to our theory, both enjoyment (positively experienced mental states) and the satisfaction of narrative-relevant desires (desires whose satisfaction makes a difference to one's life story) are prudentially good for an individual. Suffering and the frustration of narrative-relevant desires are prudentially bad for an individual. Critically, in our view, contact with reality – as contrasted with illusion or delusion – plays an amplifying role. Enjoyment is better for someone when they are taking pleasure in something real. Likewise, the fulfillment of desires is prudentially better when those desires are informed and rational. Enjoyment and desire-satisfaction are unified in a single coherent account of well-being in that both reflect the lived, self-caring perspective of a conscious subject. The chapter's final three sections address three areas of practical concern: (1) the relationship between disability and well-being, (2) decision-making for impaired newborns, and (3) decision-making for patients in irreversibly unconscious states.

The topic of Chapter 9 is *personal identity theory*. Two concepts of personal identity are important for bioethics, but need to be kept distinct. First, *numerical* identity is the relationship an individual has to themselves in being one and the same individual over time. Second, *narrative* identity involves a person's self-conception or self-told story about herself and her life. We explore four approaches to numerical identity from the

philosophical literature: person-based accounts, biological accounts, mind-based accounts, and a social account. The chief findings of the theoretical investigation are that (1) person-based accounts and the social account are implausible accounts of numerical identity, and (2) the biological accounts and mind-based accounts are *both* plausible, motivating a pluralistic approach to personal identity. Our account is pluralistic in the sense that we hold that policies and practices should be consistent with the assumption that the biological and mind-based accounts are both reasonable and worthy of accommodation.

Equipped with these theoretical resources, the discussion turns to three areas of practical application. First, we neutralize some common concerns about human enhancement through biomedical means. Second, we investigate and ultimately vindicate the authority of advance directives in cases of severe dementia. Third, we enter the controversy over the definition of death and associated questions about unilateral discontinuation of life support and vital organ procurement. We find that proper resolution of these issues turns primarily on practical considerations other than the nature of death. The overarching lesson of these practical investigations *deflates* the role of personal identity theory in bioethics. Contrary to the claims of most bioethics scholars who have invoked personal identity, after we have narrowed down the theoretical options to genuinely plausible accounts, the latter do not have far-reaching implications in bioethics.

The final chapter of the book, Chapter 10, addresses the ethics of procreative decision-making. It begins by defending a negative right to procreative autonomy on the basis of more general rights of autonomous agents to control their own bodies. These rights, like other autonomy rights, are limited in scope by potential harm to others. Nevertheless, the negative right to procreative autonomy supports allowing the use of a wide range of procreative technologies. We also contend that while people's interests in procreating may ground claims to assistance on the basis of justice, they have no special weight compared with other interests and so do not qualify as positive rights.

From procreative autonomy the discussion turns to the ethics of making decisions that affect which humans come into existence (or to term). These divide into fixed-identity decisions and identity-determining decisions. The former occur when one chooses whether or not to bring a specific individual into the world – for example, a decision to terminate a pregnancy because the fetus has spina bifida. The pivotal question in fixed-identity cases is whether and in what circumstances abortion is ethically permissible. We contend that presentient fetuses are not harmed by death

and that killing them is permissible. Once sentience emerges – probably no earlier than twenty-eight weeks’ gestational age – it becomes plausible that death ordinarily harms the fetus. Yet we argue (on the basis of a “gradualist” view of the harm of death) that, due to the weak psychological connections between the fetus and its possible future, the harm of death to the sentient fetus is relatively small. It follows on our view that terminating pregnancy even in the late stages of pregnancy can be justified when there is a weighty reason to do so.

Identity-determining decisions determine *which* of several possible individuals will come into being. A couple might attempt to get pregnant now or – concerned about an outbreak of an infectious disease that might affect a fetus – postpone their attempt for several months. The sperm and egg that would be part of conception now will not be the same gametes that would be involved in a conception a few months later, so the decision about whether or not to delay determines which of two possible individuals will come into being. Identity-determining cases are ethically complex when the individuals who could come to exist differ substantially in their expected quality of life. Many philosophers believe it would be wrong to bring into existence someone whose life would go worse than that of another individual who could, with little cost, be brought into existence instead. But it is hard to understand this judgment once we note that there is no actual individual whose life is made worse by such a decision. We contend that, in at least a subset of these “nonidentity” cases, it is permissible to cause the existence of someone whose life will go worse than that of another possible individual.

The chapter’s final two sections apply our theoretical conclusions about fixed-identity and identity-determining decisions to two practical issues: the use of medical technologies for sex selection and public health measures in the context of a Zika virus outbreak.