Bioethics
A Coursebook

Compost Collective

https://www.openbookpublishers.com

©2025 Daan Kenis, Mayli Mertens, Franlu Vulliermet, Varsha Aravind Paleri, Yanni Ratajczyk, Emma Moormann, Christina Stadlbauer, Bartaku Vandeput, Nele Buyst, Lisanne Meinen, Kristien Hens, Ina Devos and Ilya Gordon Villafuerte (Compost Collective)

This work is licensed under a Creative Commons Attribution-NonCommercial 4.0 International (CC BY-NC 4.0). This license allows you to share, copy, distribute and transmit the text; to adapt the text for non-commercial purposes of the text providing attribution is made to the authors (but not in any way that suggests that they endorse you or your use of the work). Attribution should include the following information:

Attribution should include the following information:

Compost Collective, *Bioethics: A Coursebook*. Cambridge, UK: Open Book Publishers, 2025, https://doi.org/10.11647/OBP.0449

Further details about CC BY-NC licenses are available at http://creativecommons.org/licenses/by-nc/4.0/

All external links were active at the time of publication unless otherwise stated and have been archived via the Internet Archive Wayback Machine at https://archive.org/web

Digital material and resources associated with this volume are available at https://doi.org/10.11647/OBP.0449#resources

Information about any revised edition of this work will be provided at https://doi.org/10.11647/OBP.0449

ISBN Paperback: 978-1-80511-510-6
ISBN Hardback: 978-1-80511-511-3
ISBN PDF: 978-1-80511-512-0
ISBN HTML: 978-1-80511-514-4
ISBN EPUB: 978-1-80511-513-7

DOI: 10.11647/OBP.0449

Cover image by Bartaku Vandeput.
Cover design by Jeevanjot Kaur Nagpal

4. Health Care Ethics

Introduction

Bioethics is most commonly associated with questions on health and disease. Matters such as end-of-life care, abortion, or preimplantation genetic testing are often what first springs to mind when one hears about 'bioethics'. While these are important topics to be discussed on a societal scale, the clinical encounter and the patient-physician relationship similarly introduce important moral questions on the interpersonal level of healthcare. Illness and disease usually result in the patient presenting in a vulnerable state. Both practically and institutionally, we put our trust in healthcare professionals to provide the necessary care per our values. These conditions come with significant risks of exploitation.

Physicians, nurses, and other healthcare workers often face difficult ethical decisions on the other side of the sick bed. For example, how should they handle unresponsive patients, or patients whose treatment options have been exhausted? What about patients making 'irrational' decisions? How should care be prioritized under conditions of resource scarcity? These are but some of the moral dilemmas that arise in the clinic.

As the recent global pandemic acutely brought to light, matters of public health can also be relevant to the lives of (healthy) individuals. Debates on allocating hospital beds, vaccines and other therapeutic resources, lockdown measures, and vaccine mandates—while challenging for policymakers—also offer puzzling cases for ethicists.

Biomedical researchers, too, are often confronted with complex ethical dilemmas. Clinical trials require (healthy) participants and patients to voluntarily and selflessly put themselves at risk for research on disease mechanisms, novel treatments, or diagnostic tools. Even if the outcomes of such research (can) constitute social goods, biomedical research also raises questions on how to ensure that research participants are not exploited, who should benefit from such research, and which types of research should be prioritized.

In sum, all parts of biomedical science and healthcare introduce significant moral challenges and trade-offs. In this chapter, we will be looking at several topical issues in healthcare ethics. But before we get into some of the ethical frameworks that researchers and physicians rely on, it is worth pointing out that healthcare ethics is a

©2025 Compost Collective, CC BY-NC 4.0
https://doi.org/10.11647/OBP.0449.04

highly interdisciplinary field. In addition to philosophers, biomedical researchers and health care professionals—including physicians, primary care workers, and nurses—also often sit on committees for medical ethics. Outside of academia, various forms of activism are increasingly and significantly influencing bioethical discourse. As is to be expected, these different standpoints often offer radically different appraisals of traditional questions in bioethics—we will encounter some examples further on. A proper, flexible ethical toolkit is necessary to appreciate and accommodate such differing viewpoints into meaningful ethical decisions and guidelines.

Doing (medical) ethics

In chapter 2, we discussed how committing yourself to one moral theory, such as utilitarianism or deontology, is often ill-advised in complex, real-life cases. This is intuitively clear in ethical reflections on medical and public health questions, where our moral intuitions can take us in wildly different directions based on the specifics of the case. Recall, for example, the utilitarian surgeon who could save five lives by harvesting the organs of a single patient. Most would agree that this constitutes a clear transgression of the individual's autonomy—a deontological principle we generally consider important in medical decisions.

Compare this with public health cases where utilitarian reasoning sometimes may provide a sensible answer to an ongoing crisis. Think of measles eradication due to vaccine mandates, or the lockdown measures during the COVID-19 pandemic. In those instances, concerns over public health (i.e. protecting vulnerable populations or avoiding the collapse of our healthcare system) were considered sufficient to support the restriction of individual liberties.

Ethical questions generally do not allow for clear-cut answers. This is immediately evident when considering medical decisions. Patients may evaluate a particular medical result very differently or care about very different aspects of health and well-being. While it may seem clear that smoking increases the risk of lung cancer, for some, smoking a cigarette might be a welcome distraction from a stressful, busy day and might thus contribute to well-being. Or consider the following: for some prospective parents, trisomy 21 (a third copy of chromosome 21, often leading to developmental delays or intellectual disability in the form of Down's syndrome) might be a valid reason to opt for pregnancy termination; while for others, it might not be. For some, breast amputation and reconstruction is a reasonable response to an identified or assumed (genetic) risk for breast cancer, while others approach such situations differently. Even though we sometimes think we can conclusively determine what is in the patient's 'best interest', this might not always be clear. In sum, medical decisions are often complex and involve a variety of actors (physicians, patients and their families/community, etc.) all operating according to different value frameworks. Bioethical reflection needs to be able to deal with these (conflicting) values.

Ethical pluralism

As ethicists, we must have sufficiently sensitive tools to draw out these different features and viewpoints in a particular case. As such, most ethicists rely on multiple moral principles to come to a sensible conclusion. This stance is often referred to as *ethical pluralism*. In an earlier chapter, we saw how, according to Kant, certain maxims and specific universal laws are absolute: you have a *duty* to tell the truth even if a murderer is knocking at the front door. In contrast to the monist view (only using one ethical theory to assess a case), in applied ethics—of which medical ethics is a more specialized field—ethicists have devised multiple, specific principles to help make an ethical decision.

William David Ross (1877–1971) proposed that such principles are *prima facie* (Ross, 2002). These principles—which we will discuss in more detail later in this chapter—seem valid as duties *at first glance*; but when applied to concrete situations, they often must be weighed against one another. According to Ross, ethics is about how to act in *specific* situations. It follows that our duties might also depend on particular situations and relations between different actors. What might seem, at first glance, to be an ethical duty may be superseded by other aspects relevant to the case. Moreover, what might be a duty for one actor—for example, a doctor—may not be a duty for a patient. Therefore, rather than providing clear-cut answers to difficult questions, ethical reflection is complex and indecisive as to *whether we can* draw the correct conclusions. In conditions of uncertainty—as in most ethical cases—it is thus crucial that we (can) assess the complexities of a case before we come to decisions.

To facilitate decision-making and to consider all relevant aspects of a case, Ross proposed seven *prima facie* duties or principles:

- Duty of fidelity (promise-keeping)
- Duty of reparation (making up for prior wrongful acts)
- Duty of gratitude (being grateful for others' acts of kindness)
- Duty of justice (being fair)
- Duty of beneficence (benefiting or helping others)
- Duty of self-improvement (education or practice)
- Duty of non-maleficence (not harming others)

Principlism: the basic idea

In the wake of Ross' list, different subdisciplines of applied ethics have thought of specific principles or duties that capture important values relevant to their domain. Medical ethicists Tom Beauchamp and James Childress first developed a principled approach in bioethics (*principlism*). Beauchamp and Childress explicitly believed that utilitarianism and deontology were inadequate frameworks to effectively deal with the

complexities that arise in medical decision-making. Their book *Principles of Biomedical Ethics* (2013) proposes four *prima facie* principles that, to this day, make up the core of much of bioethical reasoning and theorizing. Let us briefly touch upon each of these:

1. *Non-maleficence* (consequentialist): You should not cause harm. Application: Do not give children medicine that might be effective for their complaints but will cause more significant problems.
2. *Beneficence* (consequentialist): You should do good. Application: Try to cure patients.
3. *Autonomy* (deontological): Respect for people's autonomy means respecting their choices and enabling them to make informed choices by providing objective and complete information.
4. *Justice*: Treat patients fairly. Application: Doctors should not discriminate based on gender or race.

As these are *prima facie* principles, they often conflict and should be weighed against one another. For example, chemotherapy causes damage to the human body. If we strictly follow these duties as if they were obligations, the principle of non-maleficence would not permit the use of this type of treatment. However, most would agree that it is more important to try to cure someone. In this case, then, the principle of beneficence overrides the principle of non-maleficence. We might also encounter tensions between beneficence and autonomy. Consider the following: if a patient is unconscious and needs to undergo a surgical procedure to save their life, the physician cannot respect the patient's autonomy. However, in this case, the doctor may perform the procedure without consent; the duty to save a life is more important than respecting the patient's autonomy.

Principlism: comments and concerns

Before we move onto more specific applications of medical and health care ethics, it is worth reflecting on the presuppositions of principlism. Although 'the principles' to this day can be considered the primary 'toolkit' in much bioethical literature, they have received ample criticism within other disciplines and traditions. Two important criticisms have been widely expressed. First, postcolonial and feminist authors have expressed concerns over the presentation of the principles as the universal basis for moral reasoning and decision-making. Indeed, Beauchamp and Childress suggest that their principles accord to a *common morality*. They argue against the relativism gaining ground in ethical theorizing, suggesting that the principles are actually based upon something we all share—namely, a "set of norms shared by all persons committed to morality" (2013, p. 3).

Postcolonial and feminist thinkers have suggested that the principles presented by Beauchamp and Childress are deeply *situated*. These concerns reflect an argument

against the representation of principlism as a rational and universal core of morality. Instead, these authors argue that the principles are committed to a distinctly *Western* view of what is of value. In particular, the emphasis on (individual) autonomy implies a commitment to specific views on moral agents and the role of communities. In some cultures, decision-making is not merely a personal endeavour but may be collaborative or even delegated to an elder. Researchers conducting clinical trials in such cultures may face the problem that the ethical procedures imposed by Western ethics committees—requiring each participant to sign their own informed consent—may be foreign to the local population. *Decolonizing ethics* presents a meaningful alternative to principlism. It entails developing and deploying more indigenous approaches to professional ethics rather than approaching local issues through a Western lens. In her article *Decolonizing Ethics* (2018), Amohia Boulton describes her work in a tribally-owned health research centre in New Zealand. The researchers at that centre use research principles from Māori protocols rather than Western protocols. She writes that there are Māori values or ethics that all Māori understand throughout the country. These include intrinsic or implicit principles such as *Whanaungatanga* (kinship or relationship), *Ahwi* (to cherish), and *Kotahitanga* (solidarity), which guide how people work together as Māori and how they treat one another. There are also explicit principles that are written down in strategic documents. These include *Rangatiratanga* (self-determination), *Manaaki Tangata* (care of all people), *Hauora Tangata* (health of the people, interpreted holistically), *Mātauranga* (knowledge), and *Ngākau Tapatahi me te Aurere* (working with integrity leads to achievement of purposes).

A second oft-cited worry is that the principles might be too abstract to offer actual guidance in decision-making; they do not describe *how* to act. Beauchamp and Childress (2013) defend the abstract nature of the principles by arguing that their lack of concrete content is precisely what allows us to apply and specify them according to the details of the relevant situation. Nevertheless, as we have seen in the context of care ethics, ethical concerns often involve concrete needs and relations with specific others. The 'objective' approach of the principles leaves little room for questions on personhood and experience.

Proponents of narrative ethics offer one potential alternative or supplement to principlism. Narrative ethics emphasizes the importance of storytelling, voice, and the first-person perspective. Authors such as Rita Charon (2002) and Arthur Frank (2013) recognize that human beings often understand and communicate moral experiences through storytelling. In healthcare, narratives can include the stories of patients, families, and healthcare providers. Such stories may offer richly detailed, personal, and contextualized accounts of the situation at hand and, as such, can provide valuable insights into the experiences of and relations between those involved. In addition, this approach is sensitive to the importance of listening. This is important, since patients often do not 'feel heard' by their caregivers. In sum, narrative approaches generally highlight the fact that, in ethical reasoning, we are dealing with questions

of personhood, culture, and (personal) history—domains that are significantly underrepresented in the abstract and universalist approach of principlism.

Concerns have been expressed over narrative ethics as well. Some wonder whether stories are too subjective to feature in ethical (and clinical) reasoning. Others suggest that patients may not always have a clear idea of what they value, and that outside forces and narratives can influence personal stories. Some worry that given the high-speed, high-stress environment in which medical decision-making occurs, allowing stories 'to breathe' might be overly time-consuming.

Other critics offer a more measured response and suggest that a principlist approach can be fruitful to ethical reasoning, but indicate that the list provided by Beauchamp and Childress is too limited. The principles of autonomy, beneficence, non-maleficence, and justice may not be sufficient to deal with all ethical problems that might occur. They should be supplemented by dignity, integrity, vulnerability, and solidarity as essential principles in bioethical decision-making. Despite these worries, principlism with its emphasis on autonomy is still a central framework which a lot of ethical reasoning is based on, within the clinical context and biomedical research.

Now that we have a clear view of the ethical toolkit at the disposal of the medical or clinical ethicist, we can delve into some important topics within health care ethics.

Medical and clinical ethics: the patient-physician relationship

Starting with a range of questions we can categorize under the rubric of 'clinical ethics', let's take a closer look at some moral difficulties arising in the patient-physician relationship. This relationship is morally significant, not least due to its asymmetrical nature. Patients often present themselves to physicians in a state of (physical or mental) vulnerability. The physician, in turn, is situated as the expert, conditional for the receival of proper care or treatment. As such, they occupy the powerful position of effectively standing between the patient and their access to appropriate care. This imbalance evidently raises some fundamental ethical questions.

Paternalism and informed consent

Not that long ago, the relationship between doctor and patient could be characterized as *paternalistic*: patients (mostly) had to follow *doctor's orders*. We generally do not accept these overt forms of medical paternalism anymore. Current procedures and legislation in contemporary medical practice encode an important role for patients. Today, it is generally deemed unacceptable for physicians to act on behalf of their patients. One of the most important tools against medical paternalism is the requirement for *informed consent*.

The practice of informed consent stipulates that before patients can be admitted to medical procedures, they must agree (verbally or in writing) to the proposed

treatment plan. As such, informed consent effectively ensures that patients can engage in autonomous decision-making without the coercion or influence of healthcare professionals.

Merely having the choice to accept or refuse treatment does not suffice to speak of proper consent. To enable genuine self-determination concerning treatment decisions, patients must have (or be given) access to all relevant information, including the benefits, the (physical, psychological, and potentially financial) risks of treatment, and potential alternatives to the proposed therapeutic action. In addition to respecting the patient's decision, physicians and other healthcare professionals should communicate openly and transparently with the patient so as not to impeach their autonomy and, relatedly, so as not to counteract the principles of beneficence and non-maleficence. Additionally, Onora O'Neill (2002) suggests that in order for informed consent to involve truly autonomous decision-making, patients should have access to meaningful alternatives. Sadly, this is often not the case, which raises questions on whether informed consent truly enables autonomy or merely acts as a legal tool to waive the responsibility of clinicians and hold patients accountable instead.

Substituted judgement

Whatever your stance on the ethics of informed consent, sometimes a particular medical situation renders it impossible to ask for direct consent. For example, children, unconscious patients, or those with severely diminished mental capacities may not be in a position to sign the relevant forms, or they might not be capable of understanding all the relevant information. Where does that leave us with regard to their health care?

Individuals who are not competent to consent are, of course, also eligible for medical care. In those cases, an appointed guardian should consent by proxy. This decision can be based on a *substituted judgement standard* when values or interests are known—for instance, when someone is in a coma, but their spouse knows what they would have wanted. If this is not possible, a proxy should consent with the best interest of the patient or research participant in mind.

Nevertheless, every effort should be made to inform the patient anyway. For example, those who are legally unable to consent should be asked for informed *assent*. Children should be asked for their opinions about research participation or treatment, which should be considered and featured in the final decision.

The patient's best interest

Other complications arise when physicians believe that patients act in opposition to what is in their best interest. A physician might think a patient should not undergo a risky operation that will only have marginal or even adverse effects on their quality of life, yet the patient is willing to take the risk.

Some authors have suggested that health care professionals can *nudge* their patients. Nudging refers to practices intent on influencing the patient's decision. For example, the 'framing effect' is a well-known nudging technique. Social psychology suggests that, for example, communicating that 25% of patients experience complications, rather than saying 75% do not, might push the patient's decision towards refusing treatment. While nudging can help patients reconsider what is in their interests, it is generally agreed that it conflicts with autonomy. Whether this constitutes an unethical transgression of informed consent is up for debate. Some ethicists, taking a utilitarian approach, think nudging can be morally permissible if the benefits clearly outweigh the risks. Others argue that nudging involves misleading the patient and is thus not permissible under informed consent requirements.

Another complication arises when physicians *perceive* that patients are making decisions based on 'irrational beliefs'. An example often repeated in bioethics courses is the refusal of a life-saving blood transfusion based on religious beliefs by a patient raised as a Jehovah's Witness. Some ethicists suggest that irrational beliefs—which might include religious beliefs, in their view—can affect a patient's capability to consent, and physicians should not merely respect the patient's judgement in such cases. Instead, physicians have a moral obligation to engage the patient rationally and discuss all the medically relevant aspects of the decision. In short, in addition to respecting patient autonomy, physicians should not be mere passive compliers to patient decisions. Instead, they are morally obligated to act as normative guides to help patients make the right decision.

As feminist scholars have emphasized, what makes a particular belief 'irrational' is highly contingent and based on contextual, social, and cultural factors. Thus, it might be that what seems irrational to a physician is of genuine importance to a patient. Therefore, in debates on autonomy and clinical decision-making, we should exercise caution in quick attributions of incompetence or irrationality. Nevertheless, conceiving informed consent not merely as a contractual obligation but as an opportunity for dialogue and deliberation aimed at mutual understanding seems to be a fruitful approach to the aforementioned concerns.

Liver transplant

You are an ethicist at a major hospital in Brussels. The transplant team has requested an ethics consultation regarding Marco, a thirty-nine-year-old Italian man with acute liver failure. Marco had spent several years working in South America before recently relocating to Belgium to stay with his cousin. He has no health insurance, no official residency in Belgium or Italy, and no current source of income. Marco's medical history includes substance abuse, though he states he has been sober for the past nine months. His condition is critical, and he needs a transplant within forty-eight hours.

A matching liver has become available, but the procedure and follow-up care are expected to cost over €150,000, a sum the hospital would likely have to absorb. The administration is hesitant, as this expense represents the entire annual charity care budget, which typically supports dozens of uninsured patients. Concerns have also been raised about Marco's ability to adhere to the strict post-transplant regimen without insurance or stable living arrangements. Marco's cousin has committed to providing him with a place to live and helping with his recovery, but cannot contribute financially. She passionately advocates for him, citing his sobriety and determination to rebuild his life as reasons why he deserves this opportunity.

What would the best course of action be in this case?

Ethics of medical AI

An important topic in contemporary medical ethics pertains to the use of artificial intelligence (AI) algorithms across a variety of medical applications. AI is increasingly used to optimize health expenditure and resource allocation, in diagnosis and risk prediction, and in patient and hospital management. Unsurprisingly the introduction of technically complex, highly-performant algorithms in sensitive contexts such as healthcare, biomedical research, and public health raises important ethical questions. This section surveys three central topics in contemporary AI ethics: ethical principles and regulations, algorithmic bias, and the role of AI in clinical decision-making.

AI ethics and governance

At present, various governing bodies are setting up systems for the governance of AI. How this is approached differs across the world. For this coursebook, we limit our discussion of regulations to the European context. The EU's approach to AI governance is to promote the uptake of 'human-centric and trustworthy' AI systems, which serve humanity and the common good (human-centric), and are lawful, ethical, and robust (trustworthy).

The two most relevant EU policy documents on AI are the non-binding *2019 Ethics guidelines for trustworthy AI of the High-Level Expert Group on Artificial Intelligence* (AI HLEG), and the binding *AI Act* (Regulation (EU) 2024/1689). We discuss these in order.

Ethics guidelines for trustworthy AI of the High-Level Expert Group on Artificial Intelligence

The AI HLEG guidelines set out a framework for trustworthy AI, stipulating seven requirements rooted in ethical principles, and European fundamental rights.

AI HLEG lists four ethical principles: respect for human autonomy, prevention of harm, fairness, and explicability. To make the guidance more concrete, the expert

group has translated the ethical principles into seven requirements to be continuously evaluated and addressed throughout the development, deployment, and use of trustworthy AI. We will briefly discuss these principles and their related requirements in more detail. For complete descriptions of the principles and requirements, we refer to the AI HLEG guidelines document.

Respect for human autonomy: Humans interacting with AI systems should be able to retain self-determination, and AI systems' work processes should be subject to human oversight. This principle translates to the requirement of *human agency and oversight.* Users should be given sufficient information on the AI system to make autonomous decisions. Moreover, they have the right to involve a human in the decision-making process (i.e. 'human-in-the-loop') if they would be significantly affected.

Prevention of harm: AI systems cannot (exacerbate) harm to the dignity and integrity (mental and physical) of human and non-human beings, and the natural environment. Specific attention must be paid to vulnerable people and contexts of asymmetry in power or information. This principle requires developers to strive for *technical robustness and safety, privacy and data governance* and *societal and environmental wellbeing*. AI systems should be reliable, accurate, secure, and resilient to attacks, and they should have fallback plans in place and guarantee privacy or data protection. Data collected about individuals cannot be used to unlawfully or unfairly discriminate against them.

Fairness: The development, deployment, and use of AI systems should be fair. This means ensuring equal and just distributions of benefits and costs, and preventing unfair bias, discrimination, and stigmatization against individuals or groups. Individuals should be able to effectively contest AI decisions, and redress mechanisms should be in place in case of harm. The importance of fairness is stressed by the various requirements implicated, including *privacy and data governance, diversity, non-discrimination and fairness, societal and environmental wellbeing* and *accountability.* All affected stakeholders (humans, non-humans, society, the environment) should be considered and involved throughout the entire AI system's process. Clear and transparent oversight procedures should prevent unfair bias in datasets and development. AI systems should also be user-centric, accessible, and have equitable outcomes for persons regardless of age, gender, ability, or other characteristics. Finally, mechanisms should be in place to ensure responsibility and accountability for AI systems and outcomes before and after their development, deployment, and use.

Explicability: Explicability encompasses the need to have transparent AI processes, open communication on the systems' capabilities, purposes, and specific aspects, and the ability to explain the AI process and its decisions to those affected (explainability). This principle translates to *transparency* requirements for all elements relevant to AI systems. Datasets, technical processes, and outputs should be traceable and explainable, including related human decisions. Humans should be informed when interacting with an AI system and, where necessary, have the option for human interaction instead.

These principles are reminiscent of the bioethical principles of autonomy, non-maleficence (*harm prevention*), and justice (*fairness*), with the addition of the AI-specific principle of *explicability*. They are also subject to similar critiques. Tensions and conflicts may arise between these principles, sometimes in such a way that an acceptable balance cannot be achieved. For example, sometimes explicability can affect an algorithm's performance, which may lead to preventable harm. As such, these ethical principles have to be interpreted and translated into workable tools and applications.

AI Act (Regulation (EU) 2024/1689)

While the AI HLEG Ethical guidelines are helpful for AI systems' governance, they are legally non-binding. Regulation (EU) 2024/1689 (AI Act) lays down a risk-based legal framework for AI governance. In contrast to one-size-fits-all governance frameworks, risk-based frameworks aim to counteract overregulation by setting requirements and obligations proportional to the risk. In this way, the EU strives to find an optimal balance between (the benefits of) AI-related innovation and protection of health, safety, and fundamental rights against harmful effects. The approach may also be more flexible to govern quickly changing AI technologies, as it links obligations to the potential harms and risks, not specific technical classifications.

The AI Act outlines four risk levels. First, systems may be of *unacceptable risk* when "considered a clear threat to safety, livelihoods and rights of people" (European Commission, n.d.). Categorization systems using biometric information for the inference of sensitive and protected characteristics—such as race, political opinions, and sex life—fall under this category. These systems are banned under current regulation. *High risk* systems have the potential to cause significant harm to the health, safety, or fundamental rights of individuals or the environment if they fail or are misused. These systems are strictly regulated. *Limited risk* systems have specific transparency obligations due to, for instance, the risk of manipulation or deceit. Developers and deployers must ensure that users are aware that they are interacting with AI. Generative AI applications (e.g. GPT, CoPilot, and Claude) generally fall into this category. All other AI systems (e.g. spam filters) are considered *minimal risk* and are not subject to any mandatory legal requirements or obligations.

Governance of medical AI

The AI HLEG ethical guidelines and the AI act are directed at all AI systems, not specifically medical AI. Consider an AI model implemented in a (non-invasive) wearable health tracking device (e.g. a smartwatch), which continuously monitors vital signs such as blood pressure, skin temperature, and heart and respiration rate. These vital signs are input for an AI model which would alert the

patient and a trusted contact person (family member, nurse, etc.) when medical attention is needed.

How do the AI HLEG ethical principles and requirements apply to this AI system? Which aspects of this system could be ethically problematic? How would you address it? In which risk category of the AI Act would you place this AI system?

Algorithmic bias and justice

As we saw earlier in this chapter, one of the central principles in biomedical ethics is justice. Physicians, other healthcare professionals, and institutions ought to treat patients fairly and refrain from discrimination based on social categories such as gender, race, religion, or socio-economic status. As some recent controversies have shown, medical AI systems deployed across various contexts of healthcare put the matter of justice and *fairness* at the forefront of current ethical discussions.

In 2019, Obermeyer and colleagues published a study in *Science* exposing significant racial bias in a widely used healthcare algorithm for identifying 'high-risk patients' for additional clinical management. The algorithm disproportionately excluded Black patients from needed care, even when they had more severe health conditions than their White counterparts. Another example is the GAIL Risk Score algorithm used in breast cancer risk assessment. Although widely used, its performance on younger populations, non-Western patients, and atypical breast cancers is increasingly shown to be subpar. These issues point toward persistent biases in medical algorithms. *Algorithmic bias*—or systemic distortions or unfair influences in AI decision-making processes disproportionately favouring or disadvantaging particular individuals or groups—is a particularly pressing issue for healthcare.

The roots of algorithmic bias often lie in the data used to train and test AI systems. A common mantra in computer science—'garbage in, garbage out'—captures this problem succinctly: if your data is of low quality, expect low quality results. In the context of healthcare, the problem of missing and low-quality data evokes a larger history of ethical misconduct. Historical underrepresentation and exploitation of marginalized groups—such as women, people of colour, and disabled people—in biomedical research has resulted in datasets poorly reflecting the diversity of real-world populations. These and similar dynamics have resulted in a lack of reliable data today, which impacts the performance of medical algorithms. However, bias is not merely a matter of *missing* data. Most health data is, and historically has been, generated in the context of routine healthcare. As is well-documented, healthcare professionals often hold implicit biases towards marginalized patients. These biases, which can lead physicians to dismiss the concerns of certain groups, are embedded into datasets and perpetuated by AI systems.

'Garbage-in, garbage-out' tells only part of the story of bias in AI, however. In their book *Data Feminism*, Catherine D'Ignazio and Lauren F. Klein (2020) show that data are never truly 'raw' but rather are shaped and moulded by the social and political context in which they are generated. As philosopher Gabbrielle Johnson (2021) has shown, existing (unjust) social and political structures infiltrate algorithmic systems in surprising ways. Recall the resource allocation algorithm discussed by Obermeyer and colleagues. The algorithm did not directly use the patients' race as a feature in determining health needs. Instead, their analysis revealed that healthcare spending was taken as a measure for health needs. On the face of it, it makes sense to conclude that those patients spending more on healthcare, generally, have higher care needs. However, because race is a predictor of improper access to care in the US healthcare system, the algorithm implicitly incorporated these racial biases. Mechanisms like these demonstrate how societal biases seep into AI systems, even when sensitive categories like race or gender are excluded.

Clinical decision-support systems and the patient

While algorithms can be used for resource allocation, the applications we will likely face most directly are clinical decision-support systems (CDSS). These AI-driven tools aid physicians in diagnosis, risk prediction, and in making treatment decisions. Though CDSS offer potential for improving healthcare and can process vast amounts of health data, identifying patterns beyond individual practitioners' capabilities, their introduction into the clinical encounter raises several ethical questions as well.

The clinical encounter typically provides a space where patients actively engage with healthcare professionals, seeking explanations for treatment recommendations and advice based on their specific circumstances. This highlights the importance of *transparency* and *explainability* for medical AI – healthcare professionals must be able to understand and explain AI-generated recommendations to patients, particularly when they disagree with the system's conclusions. This transparency is also essential for maintaining patient *autonomy* and ensuring informed consent. In this context, the implementation of CDSS should be viewed through and designed according to a *collaborative* lens (e.g. human-in-the-loop), where such systems provide additional, explainable input to support clinical decision-making, but are not relied on exclusively.

Another concern involves algorithms potentially dominating patient-physician dialogue. AI systems' perceived 'objectivity' may diminish the patient's voice in clinical decision-making. Philosopher Miranda Fricker (2007) uses the term *epistemic injustice* to describe situations in which a patient's knowledge and testimony about their own condition(s) is dismissed. Such dismissal is ethically problematic for two reasons: it undermines human dignity by denying the patient's role in knowledge creation about their own health, and it can lead to practical harm when important patient concerns are overlooked in treatment decisions. We risk amplifying these existing concerns if

we allow CDSS recommendations to overshadow patient experiences. AI systems can only capture particular aspects of disease, primarily those that are easily quantifiable. As such, while an AI system is well-suited to incorporate genomic data or the results from blood tests, it will struggle in capturing the social and experiential realities of illness. Effective clinical decision-making must consider not only pathophysiological factors but also patients' values, social relationships, and lived experiences with illness and treatment. Failing to incorporate these qualitative aspects while deferring to AI recommendations would represent a new form of *medical paternalism*. As we saw throughout this chapter, AI ethics borrows heavily from principles in biomedical ethics. While these are important tools for AI assessment, more fundamental questions may need to be addressed as well. As feminist philosopher Alison Adam reminds us in her book *Artificial Knowing* (*1998*), the implementation of AI is preceded by important social and cultural questions: what role *can* and do we *want* AI to play in our structures and institutions? Adam – while critical – points toward the potential for AI systems to provide care, alleviate critical workers from cumbersome tasks, and democratize our access to knowledge. AI is not necessarily a threat to existence, nor a solution to all our problems, but first and foremost, it is a useful tool that can mean something *for all of us*.

Reproductive ethics

Introduction to reproductive ethics

In reproductive ethics, people have expanded upon Beauchamp and Childress' principles and applied them to ethical questions in conception, childbearing, and rearing.

On a more fundamental level, it is essential to note that opinions on the status of the unborn child heavily influence this debate. When does an embryo or foetus become a person? For some, this happens right after conception, as everything is available for the embryo to become a person. For others, this is when the embryo has implanted, the nervous system has started developing, or when the foetus is viable after twenty-four weeks. The law on abortion in The Netherlands, for example, takes the foetus' viability as a starting point and forbids abortion after twenty-four weeks.

For some philosophers, a human is only a person after birth. It goes without saying that how one perceives the 'person status' of unborn life heavily influences what one believes can morally be done to it. For example, people who believe an embryo is a person or a potential person right after conception may object to embryos being thrown away during fertility treatment. In the case of embryo selection, which is offered to prospective parents who know they carry a genetic disease, several embryos are created and checked for genetic mutations. Only an embryo without the mutation is transferred back to the prospective mother's womb. The rest are not used for fertility purposes. Some have argued that in this case, gene editing an embryo to correct the

mutated gene (with technologies such as CRISPR/CAS9) might be better, as it would not necessarily involve creating embryos which will not be used. So, this would be a solution for people who object to creating more embryos than needed. Of course, this does not change the fact that many non-viable embryos were used or even explicitly created to be tested on and then destroyed in order to develop this technique.

Besides debates on moral agency, reproductive ethics also focuses on particular biomedical technologies related to (human) reproduction. As we will see in chapter 6, developments in epigenetics raise new questions on reproductive and parental autonomy. Embryo selection is another recent technological development that requires us to reassess longstanding debates in reproductive ethics. Embryo selection is a technique to select embryos without genetic defects by conducting a genetic test on the *in vitro* embryo (official name: Preimplantation Genetic Testing). The technique has advanced to make whole-genome sequencing of the embryo's genome possible. Bioethicists have considered the extent to which such embryo selection should be allowed. Maybe it should be allowed to prevent serious harm, such as specific congenital genetic diseases. Perhaps it should also be allowed to select embryos likely to develop diseases later in life, such as Alzheimer's.

Julian Savulescu (born 1963), an Oxford ethicist, has suggested the principle of *procreative beneficence* (*2001*). He has fine-tuned this principle in various articles. However, it boils down to this: prospective parents should, in principle, and if possible, choose the embryo that will develop into the child that potentially has the best life. He gives the example of IQ: if it is possible to select, *in vitro*, an embryo with an IQ of 140, one should do that. Many would intuitively feel that it may be better, or even a duty, to select the embryo that would not develop severe diseases if you have a choice. However, for characteristics such as IQ, this is much less intuitive. Can we decide how well someone will experience life based on their genotype? Does someone with a higher IQ necessarily have a better life? For whom is this better, for the person themself or their society? Many authors have criticized the principle of procreative beneficence, and the challenges associated with a utilitarian approach are also applicable here.

John Robertson (1943–2017) has taken a deontological approach to reproductive ethics, with the principle of procreative liberty (or reproductive autonomy) (1983). This principle states that anyone has the right to reproduce or not to reproduce. This means that people can choose whether they want to have children. For Robertson, this also means people may decide on the children they want. If prospective parents find it important that their children are musical, they may choose an embryo with a genetic propensity for perfect pitch. Not anything goes, however, as there are some limitations. The choices must not harm the resulting child. Robertson would have objected to parents choosing a child with a disease. This approach has its drawbacks. How can it be determined whether the parents' choice would harm the resulting child? What counts as a disease? In a gendered world, could the gender of a child be seen as a good reason not to want that particular child? Is it fair to children that they are picked based on the parents' preferences? Should parents not accept children as they are?

Reproductive rights and justice are also a central concern in feminist and disability scholarship and activism. Calling for reproductive justice, Black and Indigenous feminist thinkers have highlighted how, apart from the right to not have a child (which is put forward in the abortion debate), other rights are equally important: the right to raise your own child, and the right to do so in a safe and healthy environment. Silliman and colleagues (2004), for example, refer to the reproductive violence bestowed upon enslaved and Indigenous people, and the many ways these histories still play out today. Other reproductive justice scholars, such as Sigrid Vertommen and colleagues (2022), emphasize the social, cultural, and political assumptions underlying assisted reproductive technologies, and the individuals—often women—whose bodies and reproductive tissues are essential in medically assisted reproduction but who are nevertheless often neglected (surrogates, egg and sperm donors, etc.). In ethical debates on reproductive technologies, it is vital to be aware of who is being excluded from the discussion, but also who is encouraged or discouraged to procreate. One example is found in the research of Virginie Rozée and colleagues (2024), who analysed European fertility clinic websites and concluded that medically assisted reproduction is presented as primarily a matter for white, cisgender, and heterosexual women.

Towards a disability bioethics

A prominent debate concerns the ethics of 'choosing disability' in the context of preimplantation genetic testing and prenatal diagnosis in general. Some disability scholars suggest that the selection of embryos without disabilities is eugenic—i.e. aimed at 'improving' the population through genetic selection. Feminist bioethicists have lamented the decontextualization of the debate in much ethical theorizing (see for example Scully 2023). Arguing against simplified views of reproductive autonomy, they suggest that we should scrutinize to what extent pregnant people are genuinely free to choose, given the overall discrimination of particular social identities and the lack of social support. Another example of the tension between a feminist stance in favour of choice and sensitivity to disability justice is the letter exchange between Eva Kittay (born 1946) and her son Leo. Kittay (2019) emphasizes the autonomy of women while at the same time agreeing with her son that in our current society, some reasons for not choosing a future with a disabled child are more informed or better than others.

While most of these debates primarily concern the 'acceptance' of disability, the 'active preference' for a disabled child is also discussed. An oft-cited example involves a Deaf couple actively seeking out a sperm donor with the right kind of genetic deafness. While many were quick to condemn such behaviour, their decision might also lead us to reconsider our intuitions on disability. As disability scholar and bioethicist Jackie Leach Scully (born 1961) notes, these debates are ultimately premised on whether we consider (the choice for) disability harmful and, if so, whether it is severe enough to outweigh procreative autonomy (2008, 2022). As

demonstrated by the responses to the aforementioned case, many clearly take this stance. We might, however, want to consider how, for many disabled people, 'disability' is deemed to be *merely* a difference, instead of necessarily a *bad* difference, to paraphrase philosopher Elisabeth Barnes (2018).

These applications of the principles in the context of reproductive decisions and technologies show that even if moral theories and principles are fine-tuned to tackle specific questions, merely choosing a principle and applying it is not enough: cases require weighing up principles and sensitivity to the particular circumstances in which questions arise. Notably, the concerns of the disability community also serve as a stark reminder of how intuitions may vary wildly across social groups. In addition to being contextually sensitive, our ethical theorizing should also involve the input of the relevant stakeholders. Maybe one of the primary tasks of the bioethicist is to make sure that those voices are heard that have traditionally been left out of the debate, but that are often most affected.

Public health ethics

As we have pointed out earlier, questions on public health also fall within the purview of health care ethics. Public health interventions such as public sanitation, fluoridization of community water, and vaccine policies have historically been some of the most effective ways of improving or maintaining population health. Two more recent examples are the lockdown measures and the vaccination campaigns in the early stages of the COVID-19 pandemic. Public health initiatives differ from traditional, clinical health care since they primarily aim to advance health at the *level of the population* through enacting top-down policies and measures affecting individual citizens. It is evident that this might raise some important ethical questions.

Triage and resource allocation

Vaccine distribution

A global pandemic is happening due to a newly mutated virus. Infection can be lethal regardless of age, and all sorts of patients are being admitted to emergency care facilities, putting enormous pressure on the healthcare system. Countries decide on different policies: some have ordered lockdowns, while others expect that people will take appropriate measures themselves to ensure everyone's safety. In some countries, people keep living as if nothing had changed; while in other countries, people take precautions.

After a few months, scientists have succeeded in developing an effective and safe vaccine, the only available treatment for the infection. However,

given its complexity, the vaccine is scarce, and only limited amounts can be produced in highly specialized facilities. The distribution of the vaccine becomes a critical task for global health, and the World Health Organization (WHO) sets up a working group to decide on the allocation of the vaccines for each country. As an expert on the matters of vaccines, infectious diseases, and public health, you are, of course, invited on this committee. What would be a just distribution? Which criteria do you use to justify your choice? Which factors do you take into consideration?

Health care and research funding are not unlimited as they compete with resources for other social goods. As a result of such competing interests, policymakers, researchers, and clinicians often have to focus on one issue over another. More extreme conditions of scarcity put these issues at the centre of ethical decision-making. A recent example was the allocation of hospital beds and respirators in the early days of the Covid-19 pandemic.

During the pandemic, treatments, hospital beds, and even health care professionals were in short supply. This resulted in policymakers and physicians having to make difficult decisions on who gets a bed at the intensive care unit, who receives a respirator, and who doesn't. Evidently, this may lead to significant moral distress in health care workers. In response, many countries developed so-called *triage protocols* to guide physicians in decision-making.

The general idea of 'triage' is utilitarian. Given the limited resources we have at our disposal, how and according to which criteria can we maximize the lives saved? What seems like simple calculus raises some important questions about which principles to follow in triage. Most guidelines take the 'requirement of critical care' as uncontroversial. Instead of merely following a 'first come, first serve' approach, we generally would want to only offer a bed to those patients who genuinely need it. Discussions among ethicists concern primarily which *other* criteria to rely on to prioritize certain patients over others.

Age is often taken as one of those criteria. While relatively uncontroversial, some ethicists have suggested the *ageism*—discrimination based on age—presupposed here can be discriminatory and, therefore, unjustifiable. More controversially, some have suggested that doctors and politicians, for example, should receive preferential care since they might significantly impact general well-being. Again, others suggest that lifestyle should be featured in the decision-making process. Should a lifelong smoker, for example, receive a respirator? Most of these latter criteria are (rightfully) controversial.

What seems perhaps less contentious, and more objective is relying on the likelihood of survival post-hospitalization. In addition to epistemological questions

such as 'how do we assess this likelihood?' and 'where do we set the threshold?', this raises some clear moral quandaries. Tolchin, Hull, and Kraschel (2021) note that while comorbidities (other conditions such as obesity, diabetes, or hypertension that affect the clinical outcome of COVID-19 infection) seem reasonable and objective parameters to base triage on, comorbidities often also track social determinants of health. Suppose we take obesity as a condition leading to a lower rank on the respirator list, and obesity is correlated to lower socioeconomic status. In that case, our 'objective' protocol might disproportionately affect marginalized groups, raising justice-related concerns.

One way to express post-hospitalization survival is through quality-adjusted life years or QALYs. QALYs represent the number of years lived in 'full health'—i.e. without disability. The general idea is that while some treatments might prolong life, they might also lead to disvalued states of well-being. So, instead of just measuring the effect of a particular intervention on life expectancy, these metrics are calibrated in terms of quality of life. Quality is a notably subjective term, of course. To arrive at a widely supported notion of quality, QALY's are thus generally measured by surveying the general public on how they would value particular, hypothetical health states. Like other forms of utilitarian calculus, policymakers rely on these evaluations to maximize the calculated number of QALYs gained through a specific intervention. As philosopher Laura Cupples (2020) rightfully points out, the idea of QALYs is built on the *ableist* assumption that rational people would prefer a shorter life in an able-bodied state than a longer one lived with disability. Cupples suggests that this is further corroborated by how QALYs are measured. Given their situatedness as (generally) able-bodied, the general public might not have a nuanced view of life with disability. Instead, Cupples argues—in line with feminist epistemologists—that we should primarily ask disabled people to evaluate these states since their testimony from experience might be more objective than that of the general public.

We might also give up on looking for a utilitarian calculus altogether. Strict egalitarians, for example, argue that no differences between patients can be operationalized as reasonable criteria for differential care allocation. Instead, they favour randomized processes – as in clinical trials, as we will see later – since these exclude external factors from triage. Finally, prioritarians intentionally favour those patients who are worse off. Instead of pursuing likely benefits, as a utilitarian might, prioritarians position sickness or socio-economical disadvantage as relevant features in prioritizing care.

Prevention and health promotion

Another increasingly important topic in public health ethics is prevention. Many commentators suggest that one of the central challenges of contemporary healthcare is the increasing costs due to an aging population and a steep rise in chronic conditions. Unsurprisingly, these past few decades have been marked by a push towards a preventive model for healthcare.

Policy-makers support these efforts by enacting health promotion policies. Banning cigarettes from public spaces, promoting exercise and healthy diets through public campaigns, and national screening or vaccination programs are all examples of such top-down health promotion policies.

Many of these programs generally frame health promotion and prevention in terms of 'making the right choices'. Relatedly, we often find conditions like diabetes, obesity, and cardiovascular diseases referred to as 'lifestyle' diseases. It is suggested then that these conditions are (to large extents) avoidable. Health promotion campaigns often aim to change the behaviour of individuals so that they adopt healthier lifestyles.

As philosopher Per-Anders Tengland (2016) argues, while such interventions may lead to better outcomes from a public health standpoint, such campaigns are often paternalistic. Indeed, in encouraging citizens to make particular health-related choices and behavioural changes, professionals assume and impose specific understandings of the relevant problems and good health-related behaviour. One reason why health promotion campaigns often fail to attain the desired outcomes is that the (lifestyle) issues identified are considered less relevant or important to people. Instead, people care about their quality of life, which is not entirely reducible to their health.

Additionally, these campaigns might also foster stigmatization of particular behaviours or bodies. Being overweight may be considered risky or irrational behaviour which the individual is to be blamed for. This might, in turn, divert attention from social explanations such as limited access to healthy food or open space for exercise, which might equally affect one's opportunities to follow such directives.

Finally, conceiving of prevention in terms of behavioural change holds individual patients responsible for conditions which might be better tackled by addressing social causes. Food deserts—areas where healthy, affordable food is scarce—are often located in poorer areas. Instead of investing in promotional campaigns, communal development and improving access to nutritious, inexpensive meals might be a more effective way to improve public health.

Research ethics in biomedical research

Research on human subjects is imperative to gain insights into the pathophysiological mechanisms of disease, discover and validate new treatments, and monitor their effects on patients. It is also clear that these research practices may be subject to relevant ethical questions. For example, we have already seen that respecting patient autonomy entails informing them about the treatment or medical procedure and acquiring consent.

In most countries, researchers must submit prospective research and trial designs to an Institutional Review Board or Ethical Commission. Sometimes, researchers feel that these requirements are burdensome. They suggest that these ethical constraints hinder research and, thus, scientific progress. Throughout the rest of this chapter, we will examine some of these requirements and how good

science and ethical science can, should, and generally do go hand in hand. But first, to clarify where these ethical requirements come from, we will give you some examples of where research went wrong.

Why research ethics matters

The Tuskegee syphilis experiment ran from 1932 to 1972 in Tuskegee, Alabama. It aimed to study the natural development of syphilis (Jones, 1993). The experiment enrolled 600 Black men, 399 with syphilis and 201 without syphilis. The men with syphilis were not told that they were part of an experiment or that they had syphilis. They were told they were treated for 'bad blood'. The participants did not receive any treatment. As part of their participation, they did receive free medical exams, meals, and burial insurance. In 1947, penicillin became the drug of choice to treat syphilis, but researchers still did not offer it to participants. In the experiment, 128 subjects died, 40 women contracted syphilis, and 19 children were born with congenital syphilis. In 1972, The Washington Post reported on the experiment, and in 1973, there was a class-action lawsuit. In 1974, there was a ten-million-dollar settlement, and the US government promised lifetime medical benefits and burial services to all living participants. In 1997, President Clinton apologized on behalf of the Nation, and in 2004, the last participant died.

The atrocities of the nazi experiments are well known. One example is the nazi freezing experiment. In 1941, German soldiers were confronted with cold weather on the Eastern Front (Annas, 1992). So, Ernst Holzlöhner and Sigmund Rascher wanted to know how much cold humans could tolerate. They performed 360–400 experiments on 280–300 Jews in Dachau and Auschwitz. Participants had to sit in cold water to see how long they would last and how they could be 'reheated'. Approximately 100 participants died during these experiments.

Clinical research ethics

While these examples are primarily historical, it is essential to note that even today, research participants—especially those who belong to marginalized groups—are vulnerable to exploitation for scientific (or financial) gain. In their text *What Makes Clinical Research Ethical* (2000), Emanuel, Wendler, and Grady (2000) summarize this intrinsic issue of biomedical research well: "By placing some people at risk of harm for the good of others, clinical research has the potential for exploitation of human subjects. Ethical requirements for clinical research aim to minimize the possibility of exploitation by ensuring that research subjects are not merely used but are treated with respect while they contribute to the social good" (p. 2701). Before we discuss the principles they laid down, let us try to find some of them ourselves.

Discussion: If you were to design an ethical code of conduct for researchers conducting experiments with human persons, what should be included?

Emanuel, Wendler, and Grady devised seven requirements to aid in assessing ethical research:

- Value
- Scientific validity
- Fair subject selection
- Favourable risk-benefit ratio
- Independent review
- Informed consent
- Respect for enrolled subjects

The first requirement, *value*, states that a treatment, intervention, or theory should improve health and well-being or increase knowledge. To be evaluated as such is necessary for research to be ethical, considering the finite resources and risk of exploitation. In short, since funds are limited and could serve several other socially relevant goals, researchers should think hard about what their research is for. Determining what is of 'value' depends on our social aims and interests. While traditionally, clinical value was determined by those conducting research—primarily (male) scientists and clinicians—ethicists and policymakers are increasingly pushing for participatory research, where those most affected should have a say in identifying the research questions, aims, and outcomes.

Scientific validity requires that research be methodologically rigorous—meaning that accepted scientific principles and methods, including statistical techniques, should be used to produce reliable and valid data. While this might seem a clear-cut, scientific issue, ethical questions also feature here. For example, in clinical trials a p-value lower than 0.05 signifies statistical significance, but we may want to reflect on the consequences of accepting such an error rate. In some cases, falsely identifying 5% of patients as at risk of cancer might not be morally permissible; while a 5% error rate might be acceptable in an influenza test.

Additionally, clinical research should have *clinical equipoise*, meaning that research comparing therapies must have an honest null hypothesis: clinical researchers must genuinely not know which treatment is better. Also, placebos should not be used when conventional treatment is available. For example, since we already have good drugs for managing diabetes, new drugs need to show benefits compared to those established treatments.

The third requirement, *fair subject selection*, entails that scientific objectives (not vulnerability or privilege) and the potential distribution of risks and benefits

should form the basis for selecting communities to study and the inclusion criteria for individual subjects. This means that the study must be generalizable to the populations that will use the intervention. It also means that those who bear the risks and burdens of research should be able to enjoy the benefits (*distributive justice*). This does not mean, however, that healthy controls should be excluded from all biomedical research. Improvement in the representation of children and women in clinical trials is still possible and necessary. Children are considered a vulnerable population. They are often underrepresented in clinical trials, although they still receive prescriptions for drugs that are not tested on children. Hence, they are *therapeutic orphans* because they are either denied the use of many new treatments or exposed to drugs that have bypassed rigorous regulatory evaluation. Women are also underrepresented in clinical trials, except for clinical trials investigating reproductive organs (this is termed *bikini medicine*, because it focuses on what is covered by a bikini). However, differences related to other aspects of the female body and endocrine system might also affect other drugs. Indeed, 80% of drugs withdrawn from the market are withdrawn because of side effects in women. For example, a dosing issue for women was only discovered after Ambien was on the market for twenty years, leading to early-morning car accidents in which women were predominantly involved.

A *favourable risk-benefit ratio* requires that risks must be minimized, potential benefits must be maximized, and benefits should outweigh the risks. This sounds straightforward, but balancing risks and benefits is complex and controversial. Can payment count as a benefit? How fair is it to balance societal benefits and burdens/risks to individuals? How do we define risk?

The *precautionary principle* states that if an action or policy has a suspected risk of causing harm to the public or the environment in the absence of scientific consensus (that the action or policy is not harmful), the burden of proof that it is not harmful falls on those taking the action that may or may not be a risk. Nevertheless, opinions on this principle are divided. To some, it is unscientific and an obstacle to progress. To others, it is an approach that protects human health and the environment.

The requirement of *independent review* entails that unaffiliated individuals review the research and approve, amend, or terminate it. This ensures that the potential impact of conflicts of interest is minimized, and social accountability is ensured. If you are conducting research, you must have this approved by an ethics committee. An ethics committee often comprises of physicians, specialists, nurses, ethicists, and philosophers. Some authors go further and suggest that institutional review boards must be sufficiently diverse. Relying on the argument from situatedness we saw earlier, they indicate that if moral intuitions and knowledge are related to who we are, a more diverse group is likely to better identify potential issues in research proposals.

This sixth requirement is *informed consent*. This means that research participants should be accurately informed about the research's purpose, methods, risks, benefits, and alternatives. They should understand this information and its bearing on their

personal clinical situation (if applicable). They should consent to participate in the research voluntarily (without outside pressure) and, as noted earlier, be competent to consent. The randomization of assignment to treatment or a control group should be explained well to participants. Particular attention should go into averting the *therapeutic misconception. The therapeutic conception occurs when patients or research participants hold the* mistaken belief that their participation in a clinical trial will lead to personal benefit. This is clearly not always the case, for example, when patients are offered a placebo treatment, or the intervention turns out to be ineffective. Informed consent also means ensuring that participants know they are primarily participating in research to contribute to scientific knowledge rather than their own benefit.

Informed consent in biomedical research is sometimes presented as a blanket statement. The participant signs a document once at the beginning of the study. Afterward, they donate their samples or data to the researchers, which the research community can use forever without any restrictions. This is sometimes called the 'sign here to consent forever'-model. This model has become more contentious in the wake of the controversies surrounding the HeLa cell line (Skloot, 2010). The HeLa cell line was taken from African American woman Henrietta Lacks and distributed to research labs worldwide. The HeLa cell line is still used to this day. Neither Lacks (who eventually succumbed to cervical cancer) nor her family would receive compensation for the highly lucrative tissue sample taken from her.

In the current context, where health care research increasingly involves gathering large amounts of data, we must revisit such narrow interpretations of consent. As clinical data is increasingly used for secondary purposes (e.g. biobanks or reusing clinical datasets), this raises questions such as: to whom does this data belong? Should we consider data as donated? Or should participants have a say in what kind of secondary research their data are used for?

A recent example of an issue with informed consent concerns the *Havasupai tribe* in Arizona (Van Assche, Gutwirth, and Sterckx 2013). In 2003, the University of Arizona gathered blood samples from Havasupai members. The goal was to investigate the high incidence of type 2 diabetes—itself linked to historical food shortages due to forced relocation of the tribe by the US government—amongst the Havasupai people. The tribe members received oral information about the focus of the research project on diabetes, after which they willingly participated in the study and provided blood samples. In the written informed consent form, however, the purpose of the study was described more vaguely ("study the causes of behavioural/medical disorders"), so the research scope was not limited to diabetes only. One of the researchers involved had already obtained a research grant to study genetic causes of the (assumed) high incidence of schizophrenia within the Havasupai tribe. As a result, the tribe's genetic material, blood samples, and biomedical data were also used and shared with other researchers to research inbreeding and schizophrenia.

Additionally, the samples were used to trace the Havasupai genetic origin, contradicting their own cultural origin story, without seeking permission from the

tribe. All these additional research aims were not adequately disclosed to Havasupai Tribe members. The Havasupai eventually sued the University of Arizona for invading personal and 'cultural and religious' privacy and causing harm and distress. Their blood samples were returned, and participants were financially compensated.

Examples like the Havasupai make us reconsider the type of consent we can demand from research participants. Some alternative models for informed consent are *tiered consent* (e.g. 'I consent to this research but not further studies') or *dynamic consent* (the participant has access to a digital platform to check what kind of research their samples are used for, and can revoke their consent accordingly). However, does *individual* informed consent suffice? Should community considerations not be brought into perspective?

The last principle is *respect for enrolled patients*. This means that a patient's privacy should be protected and they should be allowed to withdraw. It also means that their well-being should be monitored, and they should be informed about potentially relevant research outcomes for themselves and in general. The *European Union General Data Protection Regulation* (GDPR) came into effect to improve data privacy and protection in May 2018. The GDPR has the following principles: (1) consent for data usage and storage should be obtained; (2) if a breach of security and privacy has occurred, participants should be notified promptly; (3) participants have a right to access their data; and (4) participants have a right to be forgotten (to delete their data).

There is a right to *data portability*: the data subject has the right to receive personal data concerning them, and privacy should be part of the design. Also, researchers should designate potential data protection officers who are aware of the regulations.

Research in developing countries deserves special attention. We have seen that placebos should not be used in clinical trials if a known treatment is available. However, which standards of care should apply here, those of developed or developing countries? Some have argued that a placebo is justified in clinical trials in developing countries, even if the treatment is available in developed countries but not in developing countries. At the same time, we could argue that we owe more care to research participants in developing countries. In their 1997 publication, "*Unethical trials of interventions to reduce the perinatal transmission of HIV in developing countries*", Lurie and Wolfe argue that certain clinical trials with Zidovudine in developing countries were unethical. Indeed, in 1994, there was the discovery of a significant reduction in HIV transmission from mother to child after treatment with Zidovudine (25% to 8%). However, this treatment was expensive (over $800 per pregnancy)—the WHO decided that a less costly alternative was needed for developing countries. A shorter treatment with Zidovudine was proposed. A double-blind, placebo-controlled trial with two arms was executed: one arm was a placebo, and the other was a shorter treatment with Zidovudine 076. The argument was that using the placebo arm was warranted here because the standard of care in the developing country was 'no treatment'. However, Lurie and Wolfe argued that this justification was invalid since

the treatment was available in developed countries: foetuses were potentially exposed to HIV if they were in the placebo arm, which could have been prevented.

Conclusion

Although medical ethics is often reduced to a set of hot button issues such as end-of-life care and designer babies, this chapter showed that clinical encounters, public health, and biomedical research give rise to a variety of ethical questions. In order to address such a wide swath of complex situations, bioethicists generally rely on a set of prima-facie principles such as autonomy, non-maleficence, beneficence, and justice to assess their stakes and weigh the various values involved. While frameworks such as principlism often function as a useful starting point, throughout the chapter we have also seen examples of their limitations. Although still a foremost part of the bioethicist's toolbox, these principles themselves are situated in a specific historical, cultural, and social context. As such, they need to be enriched by considering the voices of different traditions and social positions. The need to consider existing inequalities and involve stakeholders from a variety of cultural or social backgrounds became even clearer when we discussed reproductive issues, public health, and algorithmic justice. Bioethicists play an important role in guiding policy-making and public discourse on these topics. As such, they have the opportunity or maybe even the obligation to ensure that all relevant voices are heard.

Bibliography

Adam, Alison, ed. 1998. *Artificial Knowing: Gender and the Thinking Machine*. London New York: Routledge.

Annas, George J., and Michael A. Grodin, eds. The Nazi Doctors and the Nuremberg Code: Human Rights in Human Experimentation. Oxford University Press, 1992.

Barnes, Elizabeth. 2018. *The Minority Body: A Theory of Disability*. Studies in Feminist Philosophy. Oxford: Oxford University Press.

Beauchamp, Tom L., and James F. Childress. 2013. *Principles of Biomedical Ethics*. 7th edition. New York: Oxford University Press.

Boulton, Amohia. 2018. "Decolonising Ethics: Considerations of Power, Politics and Privilege in Aotearoa/New Zealand". *Southern African Journal of Social Work and Social Development* 30 (1). https://doi.org/10.25159/2415-5829/3825

Camporesi, Silvia, and Maurizio Mori. 2021. "Ethicists, Doctors and Triage Decisions: Who Should Decide? And on What Basis?". *Journal of Medical Ethics* 47 (12): e18–e18. https://doi.org/10.1136/medethics-2020-106499

Charon, Rita, and Martha Montello, eds. 2002. *Stories Matter: The Role of Narrative in Medical Ethics*. Reflective Bioethics. New York: Routledge.

Cupples, Laura M. 2020. "Disability, Epistemic Harms, and the Quality-Adjusted Life Year". *IJFAB: International Journal of Feminist Approaches to Bioethics* 13 (1): 45–62. https://doi.org/10.3138/ijfab.13.1.03

D'Ignazio, Catherine, and Lauren Klein. 2020. *Data Feminism*. Cambridge: MIT Press.

Emanuel, Ezekiel J., David Wendler, and Christine Grady. 2000. "What Makes Clinical Research Ethical?". *JAMA* 283 (20): 2701. https://doi.org/10.1001/jama.283.20.2701

Engelen, Bart. 2019. "Ethical Criteria for Health-Promoting Nudges: A Case-by-Case Analysis". *American Journal of Bioethics* 19 (5): 48–59. https://doi.org/10.1080/15265161.2019.1588411

European Commission. n.d. "AI Act | Shaping Europe's Digital Future". Accessed 11 January 2025. https://digital-strategy.ec.europa.eu/en/policies/regulatory-framework-ai

Fricker, M. 2007. *Epistemic Injustice: Power and the Ethics of Knowing*. Oxford: Oxford University Press.

Frank, Arthur W. 2013. *The Wounded Storyteller: Body, Illness, and Ethics, Second Edition*. Chicago: University of Chicago Press. https://press.uchicago.edu/ucp/books/book/chicago/W/bo14674212.html

Future of Life Institute. 2024. "High-Level Summary of the AI Act | EU Artificial Intelligence Act". *EU Artificial Intelligence Act*, 27 February 2024. https://artificialintelligenceact.eu/high-level-summary/

HLEG. 2019. "Ethics Guidelines for Trustworthy AI". Brussels: European Commision.

Kara, Mahmut Alpertunga. 2023. "Is It Possible to Allocate Life? Triage, Ageism, and Narrative Identity". *The New Bioethics: A Multidisciplinary Journal of Biotechnology and the Body* 29 (4): 322–39. https://doi.org/10.1080/20502877.2023.2261735

Johnson, G. M. 2021. "Algorithmic bias: On the implicit biases of social technology". *Synthese 198* (10): 9941–9961. https://doi.org/10.1007/s11229-020-02696-y

Jones, James H. 1993. *Bad Blood: The Tuskegee Syphilis Experiment. New and expanded edition*. Cambridge: Free Press.

Kittay, Eva Feder. 2019. *Learning from My Daughter: The Value and Care of Disabled Minds*. New York: Oxford University Press.

Lurie, P., and S. M. Wolfe. 1997. "Unethical Trials of Interventions to Reduce Perinatal Transmission of the Human Immunodeficiency Virus in Developing Countries". *The New England Journal of Medicine* 337 (12): 853–56. https://doi.org/10.1056/NEJM199709183371212

Obermeyer, Ziad, Brian Powers, Christine Vogeli, and Sendhil Mullainathan. 2019. "Dissecting racial bias in an algorithm used to manage the health of populations". *Science* 366 (6464): 447–453. https://doi.org/10.1126/science.aax2342

Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence and amending Regulations (EC) No 300/2008, (EU) No 167/2013, (EU) No 168/2013, (EU) 2018/858, (EU) 2018/1139 and (EU) 2019/2144 and Directives 2014/90/EU, (EU) 2016/797 and (EU) 2020/1828 (Artificial Intelligence Act).

Robertson, John A. 1983. "Procreative Liberty and the Control of Conception, Pregnancy, and Childbirth". *Virginia Law Review* 69 (3): 405–64. https://doi.org/10.2307/1072766

Ross, W. D. 2009. *The Right and the Good*. British Moral Philosophers. Oxford: Clarendon Press.

Rozée, Virginie, Anna De Bayas Sanchez, Michaela Fuller, María López-Toribio, Juan A. Ramón-Soria, Jose Miguel Carrasco, Kristien Hens, Joke Struyf, Francisco Guell, and Manon Vialle. 2024. "Reflecting sex, social class and race inequalities in reproduction? Study of the gender representations conveyed by 38 fertility centre websites in 8 European countries". *Reproductive Health* 21 (150). https://doi.org/10.1186/s12978-024-01890-2

Savulescu, J. 2001. "Procreative Beneficence: Why We Should Select the Best Children". *Bioethics* 15 (5–6): 413–26. https://doi.org/10.1111/1467-8519.00251

Scully, J. L. 2008. *Disability Bioethics: Moral Bodies, Moral Difference*. Lanham: Rowman & Littlefield.

——. 2022. "Being Disabled and Contemplating Disabled Children". In *The Disability Bioethics Reader*, edited by Joel Michael Reynolds and Christine Wieseler. London: Routledge.

——. 2023. "Feminist Bioethics". In *The Stanford Encyclopedia of Philosophy*, edited by Edward N. Zalta and Uri Nodelman. Stanford: Stanford University. https://plato.stanford.edu/archives/fall2023/entries/feminist-bioethics/

Silliman, Jael, Marlene Gerber Fried, Loretta Ross, and Elena Gutiérrez. 2004. *Undivided Rights: Women of Color Organize for Reproductive Justice.* Chicago: Haymarket Books.

Skloot, Rebecca. 2010. *The Immortal Life of Henrietta Lacks*. Crown Publishing Group.

Soled, Derek. 2021. "Public Health Nudges: Weighing Individual Liberty and Population Health Benefits". *Journal of Medical Ethics* 47 (11): 756–60. https://doi.org/10.1136/medethics-2020-106077

Tengland, Per Anders. 2016. "Behavior Change or Empowerment: On the Ethics of Health-Promotion Goals". *Health Care Analysis* 24 (1): 24–46. https://doi.org/10.1007/s10728-013-0265-0

Tolchin, Benjamin, Sarah C. Hull, and Katherine Kraschel. 2021. "Triage and Justice in an Unjust Pandemic: Ethical Allocation of Scarce Medical Resources in the Setting of Racial and Socioeconomic Disparities". *Journal of Medical Ethics* 47 (3): 200–202. https://doi.org/10.1136/medethics-2020-106457

trail. 2024. "EU AI Act: Risk-Classifications of the AI Regulation". Accessed 11 January 2025. https://www.trail-ml.com/blog/eu-ai-act-how-risk-is-classified

Van Assche, Kristof, Serge Gutwirth, and Sigrid Sterckx. 2013. "Protecting Dignitary Interests of Biobank Research Participants: Lessons from Havasupai Tribe v Arizona Board of Regents". *Law, Innovation and Technology* 5 (1): 54–84. https://doi.org/10.5235/17579961.5.1.54

Vertommen, Sigrid, Bronwyn Parry, and Michal Nahman. 2022. "Global Fertility Chains and the Colonial Present of Assisted Reproductive Technologies". *Catalyst: Feminism, Theory, Technoscience* 8 (1): 1–17. http://doi.org/10.28968/cftt.v8i1.37920