

*Nonmaleficence and Negative Constraints***4.1 Introduction**

The principle of nonmaleficence states a prohibition on causing harm to others in the absence of justifying circumstances. Among general ethical principles, *nonmaleficence* has the strongest claim to being self-evident.<sup>1</sup> If an ethical theory claimed that harming others had no tendency to be wrong, we would be justified in rejecting the theory on this basis alone. But to accept the obvious claim that harming others tends to be wrong leaves open several important issues. It leaves open exactly what constitutes a harm, when harming others is permissible, and the range of beings that fall within the scope of the prohibition on harm.

At a general level, the concept of harm is very familiar. We know that such actions as tormenting, assaulting, and killing others generally involve harming them. We know that harm is bad for an individual and that its opposite, benefit, is good for an individual. Relatedly, we generally think that harming someone for no good reason is unethical, and the fact that someone is attempting to harm another person can provide grounds for restricting the former person's liberty. In most instances, we are confident in our attributions of harm and our grasp of its practical significance.

However, the precise nature of harm is surprisingly difficult to capture. This is important because, depending on which account of harm is correct, an action may, or may not, cause harm, with apparent implications for whether the action is wrong. Consider these questions. Does it harm a patient to end his life if he is terminally ill, wants to die, and reasonably believes that he will never again enjoy an acceptable quality of life? What if the patient has entered an irreversibly vegetative state and had no settled preferences regarding continued life in this state? While it is obvious that

<sup>1</sup> For an excellent discussion of self-evidence in ethics, see Robert Audi, *The Good in the Right* (Princeton, NJ: Princeton University Press, 2004), 48–54.

sentient animals can be harmed, most obviously by being caused to suffer, what about nonsentient animals and plants? They are living organisms, and can be damaged or killed, but can they be harmed? In the case of beings who can be harmed, does harming them require *changing their condition*, making them worse off in some way? An affirmative answer might seem obvious, but any sense of obviousness dissipates in the context of “creation cases”: situations in which we bring someone into being. For example, one might hold that it harms a human being to bring her into an existence that involves unrelenting misery. Yet, since the individual did not exist before coming into existence, bringing her into existence did not change her condition.

Special circumstances in which harming others is sometimes regarded as permissible include self-defense, punishment, consent by the person harmed, or causing someone a minor harm in order to prevent a worse harm to the same individual. The latter two circumstances arise frequently in clinical care and research. More challenging are cases in which one person is harmed to save others from greater harms. These require some analysis of how nonmaleficence relates to individual rights.

Accepting the principle of nonmaleficence also leaves open the range of beings that fall within its scope. A small number of ethical theorists maintain that our obligations to refrain from harming ultimately concern only persons or human beings.<sup>2</sup> Because we have obligations of nonmaleficence to a set of beings that includes persons, and because sentient animals can be harmed – they can feel pain and, more generally, have an experiential welfare – a reasonable *presumption* is that obligations of nonmaleficence extend beyond persons and humanity to include sentient animals. This leaves open whether obligations of nonmaleficence are *equally strong* as applied to persons and to sentient nonpersons. Chapter 7, on moral status, defends the thesis that sentient beings have moral status, implying that the scope of nonmaleficence extends this far. It also addresses the question of the strength of the prohibition on harming by distinguishing sentient beings with rights of differing strengths and sentient beings without rights.

We begin this chapter by surveying some prominent accounts of the nature of harm. We defend 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. Though this account does not capture all the

<sup>2</sup> For a classic representative, see Immanuel Kant, *Lectures on Ethics*, trans. Peter Heath, ed. Peter Heath and J. B. Schneewind (Cambridge: Cambridge University Press, 1997).

ways in which the term “harm” is used, it does capture what matters about harm for the purposes of ethical analysis. Using this account we then specify nonmaleficence into several general moral rules corresponding to ways in which individuals can be harmed. These foundational reflections are followed by discussions of three areas of ethical concern in which rules concerning harm are prominent: (1) the ethics of torture, including health professionals’ involvement in torture; (2) the limits of permissible risk in pediatric research; and (3) the ethics of medical assistance-in-dying.

## 4.2 What Is Harm?

### *Three Accounts*

Accounts of the nature of harm can be divided into comparative and noncomparative accounts. *Comparative accounts* hold that one is harmed when one is made *worse off* in some way. If you kick a dog hard, causing substantial pain, you make the dog worse off – in at least two senses. Because of your action, the dog is worse off than she was before your assault, and she is worse off than she would have been had you left her alone. These two senses in which one can be made worse off correspond to two more specific comparative accounts of harm. According to the *historical* account, you harm someone if and only if you make them worse off *than they were before your intervention*. According to the *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*.<sup>3</sup>

As the dog example illustrates, these two accounts typically coincide – someone is made both worse off than they were beforehand and worse off than they would have been otherwise. But the accounts can diverge. Suppose a sick child needs a full dose of some medication, which her attending physician prescribes, but an inattentive pharmacist provides only a partial dose, which the child receives. The child’s condition improves, but less than it would have with a full dose. In this case, the pharmacist’s error seems to harm the child – even though she is not made worse off than she was before receiving the medication – because it makes her worse off than she would have been in the absence of the error.

<sup>3</sup> Kicking the dog for no good reason also makes her worse off than she *should have been*, indicating the possibility of a *normative* account of the relevant comparison. The possibility of defining harm partly in normative terms will resurface later when we consider harm by omission.

Cases like that of the inattentive pharmacist lead us to reject historical accounts of harm as inadequate for use in ethical analysis. First, they fail to capture the fact that harms can take place even though there is no historical worsening. These are often harms that matter for evaluating the ethics of a course of action. Second, without knowing what would have happened without the event that caused the harm (a counterfactual matter), it is impossible to assess the magnitude of the harm. How bad was it for the child that she only received a partial dose? That depends on the difference between the course of her illness following the partial dose and what its course would have been with the full dose. Assessing the magnitude of harm is often an essential part of an ethical analysis.

Proponents of *noncomparative accounts of harm* claim that making someone worse off is not a necessary condition for harming them. Instead, harming someone involves putting them in an intrinsically bad, or harmful, state. According to one formulation of this alternative account, “An action harms a person if the action causes pain, early death, bodily damage, or deformity to her. . . .”<sup>4</sup> Kicking harms the dog because it puts her into a state of pain. One reason for adopting a noncomparative account of harm is that it explains a common intuitive reaction to some creation cases that comparative accounts appear unable to explain. Consider a child who is brought into existence with a significant, genetically determined disability that causes her substantial pain. According to the noncomparative account the child is harmed by the act of procreation because she is thereby put into harmful states like pain. Comparative accounts seem unable to draw this plausible conclusion, because being brought into existence does not make the child worse off. She is not made worse off than she was beforehand because she *wasn't* beforehand. She is not made worse off than she would have been otherwise because she *wouldn't have been* otherwise.

We do not think that it is necessary to adopt a noncomparative account of harm to deal with these cases. In Chapter 10, when we discuss the “nonidentity problem,” we deny that the child in this sort of case is harmed. The child could not have existed in a better state because the disability was caused by genetic factors and any child who came into existence with a different genome would have been a different individual.

<sup>4</sup> Elizabeth Harman, “Can We Harm and Benefit in Creating?,” *Philosophical Perspectives* 18 (2004): 93. Seanna Shiffrin presents a different noncomparative account of harm in “Wrongful Life, Procreative Responsibility, and the Significance of Harm,” *Legal Theory* 5 (1999): 117–148.

The parents did not harm their child even if they anticipated that the child would come into existence with a significant disability.

For us, the central problem with noncomparative accounts of harm is that they seem unable to make sense of the harm of death.<sup>5</sup> Nearly everyone agrees that death (ordinarily) harms the individual who dies, and we concur. Since an afterlife might be an improvement over earthly existence, the common assumption that death is a harm tacitly assumes that there is no afterlife and that death involves our going out of existence. The problem for noncomparative accounts is that nonexistence cannot be an intrinsically bad state of an individual because it is not a state of the individual at all. Rather, someone's having died entails that the individual does not exist. By contrast, a counterfactual account of harm can explain why death is bad for the individual by looking at how the individual would have fared were they not to have died. If they would have continued to live a life worth living, then death harmed them. (The historical account of harm also has difficulty explaining the harm of death because, in order to do so, it would have to compare the condition of an individual right before death with the condition of the individual upon becoming dead – but, once death occurs, there is no individual in any condition.)

Having rejected historical and noncomparative accounts of harm, we endorse a counterfactual account of harm as the one best suited for ethical analysis. We think that defining harm using such an account can explain why we care about harm in the first place – because making someone worse off than they would otherwise have been is typically a bad thing to do to them. As noted above, the counterfactual account also has the advantages that it can plausibly include failures to benefit as harms in certain cases, explain why death can be a harm, and allow us to calculate the magnitude of the harm that someone incurred. Moreover, the historical account of harm may be understood as *implicitly presupposing* the counterfactual account. Why, in determining whether Joe was harmed, should we ask how Joe *was faring* just before the putatively harmful event? Because how he was faring suggests (in ordinary cases) how he *would be faring* in the absence of that event. Despite these strengths of the counterfactual account, it faces several important challenges. Addressing them will show how the account is useful in ethical analysis, make the view's implications more precise, and allow us to acknowledge that we are biting some intuitive bullets by accepting it.

<sup>5</sup> See Ben Bradley, "Doing Away with Harm," *Philosophy and Phenomenological Research* 85 (2012): 390–412, at 400–401.

*Challenges to the Counterfactual Account**Preemption Cases*

Suppose that Dan dies of a brain aneurism precisely at noon. One second later his body is riddled with bullets shot by a firing squad. Had the aneurism not occurred, Dan would have died a second afterward. So the aneurism does not make Dan worse off than he would have been in its absence, which on the counterfactual account entails that the aneurism does not harm Dan. Many writers consider this implication counter-intuitive.<sup>6</sup>

One theoretical option for those who want to preserve their intuition that the aneurism harms Dan would be to adopt a hybrid comparative account of harm. On such a view, an event would constitute a harm if it made someone worse off than they would otherwise be *or* if it made them worse off than they were before. Though we are sympathetic to this hybrid account as a way to capture those events we intuitively judge to be harmful, we think that it is unnecessary and unhelpful for the purpose of ethical analysis.

The historical account is unnecessary because the counterfactual account can identify harms in the cases in which they matter ethically. In Dan's case, for example, there would seem to be no reason to expend resources to prevent the aneurism, if we cannot save him from the firing squad. Why? Because he is not, on balance, made worse off because of the aneurism.<sup>7</sup> On the other hand, if we could prevent both, that would provide a reason to do so: the aneurism *plus* the bullets make Dan worse off than he would otherwise be by depriving him of life.

Consider another case, in which unconscious Suraj will lose a leg unless his infected toe is amputated. Removing the toe harms him on the historical account, since he is – let us assume – worse off than he was just before the amputation, but benefits him on the counterfactual account, since he would otherwise go on to lose his leg from the infection's spreading. Now, one might interpret the historical account such that it reaches the same verdict about amputating Suraj's toe: doing so confers an overall benefit insofar as it removes a dangerous infection that could spread. But this interpretation depends on counterfactual reasoning.

<sup>6</sup> See, e.g., Michael Rabenberg, "Harm," *Journal of Ethics & Social Philosophy* 8 (3) (2015): at 8–10.

<sup>7</sup> Could there be something wrong with killing someone who was about to die anyway? Absolutely, but it is not helpful to think in terms of harming them. Killing someone who is about to die might violate their rights, although it does not, on balance, make them worse off.

Why might it be plausible to think that Suraj would be better off *after the amputation than he was beforehand* (a historical comparison)? Because, *if the toe had not been removed, the infection would have led to Suraj's losing his leg* (a counterfactual comparison). Even where it delivers plausible verdicts, the historical account of harm is unnecessary, and often, as in this case, it depends on the counterfactual account. Returning to practical ethical concerns, what should we do for Suraj? Barring very unusual circumstances, we should operate and remove the toe, since he will be much worse off if we do not. Labeling the amputation as a harm does nothing to help the ethical analysis, which should be concerned only with what would otherwise happen to Suraj.<sup>8</sup>

By rejecting the historical account and a hybrid comparative account we end up denying that some intuitively “harmful” events (like the aneurism) are actually harmful. We accept this consequence. Our definition of harm has a particular function. The definition is not intended to capture all common language uses of the term “harm.” Instead, it is intended to capture the part of our usage that is important for ethics.

### *Omission Cases*

A second challenge to the counterfactual analysis of harm denies that making someone worse off than they would otherwise have been is a *sufficient* condition for harming that individual. One argument motivating this challenge appeals to our judgments about “omission cases.” Suppose you buy a drum that you intend to give to Manuel, which would please him greatly. But Kevin sees you leaving the store and persuades you that you would enjoy it if you kept it for yourself. So you do, and never mention the earlier intention or the omission to Manuel.<sup>9</sup> The counterfactual account suggests that you harm Manuel by not giving him the drum, because you make him worse off than he would have been had you given it to him. Yet it seems implausible that your decision not to give Manuel the drum harms him.

This objection identifies a challenge for accounts of harm generally, not only counterfactual accounts. In fact, we can construct parallel cases for historical and noncomparative accounts. For example, suppose that your neighbor Manuel is out in the sun without a hat or sunscreen. You could

<sup>8</sup> There may be cases in which someone is benefited on one dimension but harmed on another (for example, an operation saves their sight but leaves an ugly scar). Still, we can capture what matters on the counterfactual account by comparing what would have happened with and without the operation. The historical account does not help.

<sup>9</sup> This case is similar to one presented in Bradley, “Doing Away with Harm,” 397.

cross the road and offer him a hat, but you don't – he is capable of getting his own. Manuel gets sunburned. As a result of your omission, Manuel is worse off than he was before and so the historical account suggests that you harmed him. Moreover, as a result of your omission, Manuel is in a harmed state, so the noncomparative account also suggests that you harmed him. But, as with the drum example, it is implausible to say that you harmed him.

The challenge here is thus not about which account of harm we adopt, but how we should attribute harms in omission cases. We do not want to hold you responsible for Manuel's being harmed by the sun, because you do not cause the harm. However, we do, plausibly, want to hold the negligent pharmacist responsible for failing to prescribe a high enough dose of medication to her patient. How to distinguish causing an event from failing to prevent it is a difficult philosophical question. We do not attempt to solve it in this book. For the purposes of ethical analysis we rely on intuitive judgments about which behaviors are acts and which are omissions. *In terms of responsibility for harms, we hold someone responsible only if they cause the harm or if they have an ethical duty to prevent it from occurring.*

#### *Fair Competition Cases*

Another potential counterexample to comparative accounts of harm represents what we might call "fair competition cases." Suppose Al is excelling in a calculus class. Eager to get into a top doctoral program, he realistically believes that a letter from his professor stating that Al was the top student in the class will significantly improve his prospects. Unfortunately for Al, Jill (but only Jill) outperforms him, and their honest professor is unwilling to write that Al was the top student. Through her excellent work, Jill has made Al worse off than he would otherwise have been. But we might hesitate to attribute harm here. Not only does it sound a little odd to say that Jill's excellent performance *harmed* Al; surely, in doing work superior to Al's, Jill did nothing pro tanto (presumptively) wrong. Yet it is widely believed that it is pro tanto wrong to harm other persons. If we retain this belief and the intuition that Jill does nothing wrong, then we must judge that Jill's excellent performance does not harm Al, in which case making someone worse off is not sufficient for harming him.

In response, we accept that Jill may do nothing wrong but nevertheless harms Al. Disadvantaging someone, with no ill intent, in a fair competition is a plausible example of a harm where the perpetrator does not act even pro tanto wrongfully. Our account is intended to capture when



something is a harm and what the magnitude of the harm is in ways that are relevant to ethical analysis of cases involving harm. It does not entail that all harms are even *pro tanto* wrongful. Our account is descriptive, not normative, so it is a separate matter to decide how to evaluate a specific harm. Note that the claim that Jill's action is not at all wrongful is consistent with the thesis that *intentionally* harming is always *pro tanto* wrong; had Jill intended to set back Al's dreams, it is plausible to think that her behavior would have been morally problematic. It is also consistent with the thesis that basing admission to doctoral programs on relative rankings of students within a program is unfair, in which case Al would be wronged, but not by Jill.

### 4.3 Types of Harm and Preliminary Specifications of Nonmaleficence

The principle of nonmaleficence states, at a first approximation, that it is *pro tanto* wrong to harm others. This approximation will be qualified in several ways in our discussion, including for fair competitions, where consent has been given, and to take account of cases where a *risk* of harm is imposed. To say that harming others is *pro tanto* wrong means that the fact that an act will harm another individual is a reason not to do it, but a reason that can be outweighed by sufficiently important opposing reasons. Nonmaleficence may be specified into the following relatively general rules, each concerned with a way in which individuals can be harmed:

1. Do not cause pain, suffering, or other experiential harm.
2. Do not cause illness, injury, or disability.
3. Do not deprive of goods or opportunities to which the individuals deprived have legitimate claims.
4. Do not kill.<sup>10</sup>

These are good rules of thumb for everyday life. As our later discussions show, they work less well to guide action in difficult cases. In such cases it can be important to understand the underlying harms that the rule prohibits. Here we comment briefly on each rule.

<sup>10</sup> One might include an additional rule, "Do not deprive of liberty," insofar as interfering with individuals' liberty tends to harm them. But we understand such harm in terms of the content of rules 1 and 3. Note that a moral presumption against interfering with people's autonomous choices is relevant only to those who can act autonomously and raises moral concerns that are independent of nonmaleficence, as discussed in Chapter 5.

The first rule concerns *experiential harms*, which take the form of unpleasant, disagreeable, or aversive states of consciousness. The second rule concerns the harm of *dysfunction*, or loss of functioning, which can occur in the absence of either experiential harm or death. A person could lose a great deal of mental functioning through being pleasantly drugged and experience no unpleasant consequences of their new condition. According to most accounts of well-being, one would be harmed by this loss of functioning (see Chapter 8). The third rule concerns losses of valuable experiences (e.g., enjoyments), resources (e.g., property), and opportunities (e.g., job prospects). These losses may or may not occur in conjunction with experiential harm or dysfunction. Stealing, for example, can harm someone even if that individual never learns of the theft and is not harmed in any of the other three major ways noted here. As argued in the previous section, it is implausible that all instances in which someone causes someone else to lack goods or opportunities are cases of harming (consider omission cases), and there are even some types of harm that are not even *pro tanto* wrong (as in fair competition cases). Thus, in order to fashion a useful rule, we limit it to cases where the object of deprivation is something *to which one has a legitimate claim*. So, for example, even though Manuel has no claim on you to give him a drum, you would harm him if you stole a musical instrument from him.

The fourth rule focuses on the harm of *death*, that is, loss of life – a conceptually complex matter that merits more sustained reflection than the harms associated with the first three rules. The harm of death is distinct from experiential harm because (we assume) dead bodies have no experiences. The dying process can involve considerable experiential harm, yet most people regard death as bad for the person who dies over and above the harms involved in dying. We go to great lengths to avoid death and most people are willing to undergo considerable experiential harm and illness to prolong their lives. For example, chemotherapy for cancer is often debilitating and extremely unpleasant for the patient, yet where it offers a reasonable prospect of cure almost all patients accept it.

To some, the idea that death is a harm to the decedent seems odd. The individual who dies will not be around to experience any effects of death, and so in what sense is there a subject of the harm?<sup>11</sup> We find it plausible to understand the harm of death in terms of instrumental disvalue: death

<sup>11</sup> As Epicurus wrote: “So death . . . is nothing to us, since so long as we exist, death is not with us; but when death comes, then we do not exist” (“Letter to Menoeceus,” in *The Stoic and Epicurean Philosophers*, ed. W. J. Oates, trans. C. Bailey [New York: Modern Library, 1940], 30–34).

deprives the one who dies of whatever goods their life would have contained had they lived.<sup>12</sup> The counterfactual account of harm then allows us to say how bad it is for someone that they die. The disvalue of someone's death is a function of what goods their life would have contained had they not died. A young man killed in a car crash might be greatly harmed because he would otherwise have lived a flourishing, healthy life for the next fifty years. An elderly man who dies of a stroke at age eighty-five would likely be harmed much less by death because he misses out on less by dying.

Another complication relating to the harm of death is worth exploring here. It seems plausible that the death of a young adult is, all else equal, worse for them than the death of an older adult. According to one possible extension of this reasoning, the death of a baby is even worse for her than the death of the young adult. After all, the baby likely misses out on even more valuable life than the young adult. This reasoning can be extended back to the earliest point at which we are individuated, so that the very worst death would be that of a sentient fetus, presentient fetus, embryo, or zygote (depending on which account of personal identity is correct, as discussed in Chapter 9). To many, this seems implausible. Jeff McMahan writes:

If we thought that the death of a fetus or infant [were] as serious a misfortune as the death of an older child or adult, we would have to think of the vast number of spontaneous abortions that occur as a continuing tragedy of major proportions. We would surely mobilize ourselves, as a society, to lower the prenatal death rate. Yet the level of social spending on the prevention of spontaneous abortion remains exceedingly low – lower by far than the social investment in the search for a cure for diseases, such as AIDS, that result in far fewer deaths. The explanation for this is that we simply do not regard the death of an embryo [or] fetus as a serious misfortune. . . .<sup>13</sup>

The solution that McMahan proposes, and which we endorse, is to distinguish two factors that are relevant to the harm of death. The first, mentioned already, is the total good (the amount of valuable life and its quality) of which the decedent is deprived. The second is the connection between the individual who dies and the valuable life of which they are deprived. Most adult humans have strong psychological connections to themselves in the future and the past. They are aware of themselves as

<sup>12</sup> Thomas Nagel, "Death," *Noûs* (1970): 73–80.

<sup>13</sup> Jeff McMahan, *The Ethics of Killing* (Oxford: Oxford University Press, 2002), 165–166.

temporal beings: I am the person I was yesterday and will be tomorrow and so forth. They make plans for the future, have fears about it, and have desires regarding it. In other words, it matters to them what will happen in their future. Less cognitively developed individuals have much weaker connections to their futures. A creature who is sentient, but not self-conscious, cannot conceive of itself as someone who persists through time. It might have desires and fears, but not plans. Some sentient beings might even be so cognitively simple that they are not really connected to their futures at all. How bad it is for someone to miss out on future goods, then, is a matter of their connection to those future goods, and this varies with their cognitive capacities. In general, it is worse for a mature human to die than for a mature tortoise to die, even if death deprives each of comparable quantity and quality of life, because the human, we believe, is more closely psychologically connected to the life of which they are deprived.

When we combine these two factors that are relevant to the harm of death, we get a *gradualist* account of how bad death is for typical humans. Early in development, before they are sentient, embryos and younger fetuses miss out on a great deal if they die, but since they are not yet sentient, there is no psychological connection to their future lives and so it is not actually bad for them if they die. Once sentience is reached, death starts to be bad for the decedent, but older fetuses and young infants are much less psychologically connected to their future selves than typically developing older children and adults are. Thus, how bad it is for a human to die typically increases during early childhood, peaks close to the point where all the psychological capacities relevant to the harm of death are present, and then declines as the human ages because they miss out on decreasing amounts of valuable life.<sup>14</sup>

Turn now to a different complication that frequently arises in assessing how bad a potential harm is and that motivates an additional rule. While some types of actions are inherently or nearly always harmful, other types of actions sometimes are and sometimes are not harmful. If someone at a crowded beach throws a rock without aiming at anything in particular, whether this action causes harm depends on whether the rock strikes someone. Both the agent's choice and luck play a role in determining whether harm occurs. Because such a choice is subject to moral evaluation – most likely, in this case, that throwing the rock was irresponsibly

<sup>14</sup> For discussion of how we might put numerical values on the harm of death, see Joseph Millum, "Putting a Number on the Harm of Death," in Espen Gamlund and Carl Tollef Solberg (eds.), *Saving People from the Harm of Death* (Oxford: Oxford University Press, 2019), 61–75.

reckless – the scope of nonmaleficence includes not only actions that are expected to harm others but also actions that *risk harm to others*. It is generally wrong to cause such harms as death, suffering, injury, and loss of goods or opportunities to which one has a legitimate claim. But it is also wrong to cause *excessive risk* of these and related harms to others. This is true even where the risky act does not eventuate in harm. Throwing a rock at a crowded beach may be excessively risky even if the rock does not hit and harm anyone.

One might think we should add this rule to the four general rules identified earlier: *Do not impose risk of harm*. But this rule would be implausibly strict. Nearly everything we do imposes some risk of harm on someone. Every time one opens a storm door while exiting a house, for example, one takes an extremely slight risk of injuring some toddler who wandered up and happens to have his head next to the door. The other rules operate well as rules of thumb, so that in everyday life we can mostly follow these rules and only occasionally have to consider whether acts falling under them really involve harms or whether the harms are justified. But for risks of harm this rule would not even be helpful as an approximation. The imposition of risk is so ubiquitous in life – including in medicine – that we find it more helpful to formulate our fifth general, nonmaleficence-based rule this way:

5. Do not impose excessive risk of harm.

A judgment as to whether an imposed risk is excessive should take into account both (1) the degree of risk and (2) whether or not, in the circumstances, taking that degree of risk is justified. The degree of risk associated with a possible harm is, in turn, a function of two factors: (a) the likelihood (probability) that the harm will occur and (b) the magnitude of the harm, if it occurs. For example, injection with a local anesthetic has a high probability of transient pain at the injection site and soreness once the anesthetic has worn off. There is a very low probability that it will cause a life-threatening allergic reaction. Taking into account the factors of both likelihood and magnitude, these two risks of local anesthesia – even in combination – are generally considered to be very low. Whether imposing some degree of risk is justified will depend on various factors, including what other expected harms and benefits will result from the action. A good illustration of this point is available in the context of pediatric research, as discussed later.

From a theoretical standpoint, it is helpful to specify a principle such as nonmaleficence into some general rules, as above. It is, if anything, even

more helpful to engage a moral principle and one or more of the rules that specify it with an important problem in practical ethics. This sort of engagement is often called an “application” of a principle or theory. Such applications can both illuminate the concrete problem area and extract some of the content and moral significance of the general principle. With this in mind, the remainder of the chapter discusses three areas of practical ethical concern in which nonmaleficence plays a prominent role: the right not to be tortured, nontherapeutic pediatric research, and medical assistance-in-dying. The discussion of torture brings out several important issues, including the question of when preventing harm to others can justify inflicting harm and the distinction between what is justified in principle versus what is justified in practice. Torture is, in a sense, the easy case, since we argue that it is never actually justified. Pediatric research is more philosophically challenging, since sometimes risking harm to children will be justified on the basis of preventing harms to others. Both the right not to be tortured and nontherapeutic pediatric research are topics where our conclusions are not dependent on any particular account of harm – much of the time, in bioethics, the competing accounts of harm coincide in their verdicts. The case of medical assistance-in-dying is one where it matters how we conceptualize harm, since it requires judgments about the magnitude of harm caused by someone’s death, and the counterfactual account is uniquely capable of explaining the harm of death and assessing its magnitude.

#### 4.4 The Right Not to Be Tortured

The United Nations characterizes torture as follows:

[a]ny act by which severe pain or suffering, whether physical or mental, is intentionally inflicted on a person for such purposes as obtaining from him or a third person information or a confession, punishing him . . . , or intimidating or coercing him or a third person . . . when such pain or suffering is inflicted by or at the instigation of . . . a public official or other person acting in an official capacity.<sup>15</sup>

Torture violates the nonmaleficence-based rule against causing pain, suffering, and other experiential harm. Torture also frequently causes injuries, such as broken limbs, and sometimes lasting physical or psychological

<sup>15</sup> United Nations General Assembly, “Convention against Torture and Other Cruel, Inhuman or Degrading Treatment or Punishment” (adopted December 10, 1984; available at [www.un.org/documents/ga/res/39/a39r046.htm](http://www.un.org/documents/ga/res/39/a39r046.htm)), Article 1.

disability, such as mangled fingers or posttraumatic stress disorder. In these instances, torture violates the rule prohibiting the infliction of injury, illness, and disability. Another feature of torture – coercion – violates the victim’s autonomy rights, as discussed in Chapter 5, but here we focus on torture in relation to nonmaleficence.

In view of the enormous, inherent harmfulness of torture, nonmaleficence supports a moral obligation on the part of governments and individuals not to torture. Major international bodies have interpreted the moral obligation not to torture as absolute – as having no exceptions. They have also generally regarded the injunction against torture as a *legal* obligation within the body of international law that rests on a *human right* – a fundamental moral right common to all human beings. In 1948, the United Nations prohibited torture in the Universal Declaration of Human Rights.<sup>16</sup> In 1984, the UN adopted the Convention against Torture and Other Cruel, Inhuman or Degrading Treatment or Punishment (the source of the above definition of “torture”).

Despite these prohibitions, torture remains widespread. Over a five-year period, Amnesty International reported on torture in three-quarters of the countries in the world.<sup>17</sup> The use of torture by state officials ranged from (presumed) isolated instances to routine practice. For example, it is now widely accepted that the US government, following the September 11, 2001, terrorist attacks, practiced torture in the name of national security. Suspected terrorists held in Guantanamo Bay, Cuba were subjected to such forms of torture as waterboarding (simulated drowning); mock executions; sexual humiliations of various kinds, including rectal forced-feeding; being forced to stand, handcuffed, for up to forty continuous hours; menacing with dogs; sleep deprivation for as long as a week; hypothermia (in one case to the point of death); and threats to kill or rape prisoners’ family members, including children.<sup>18</sup> Justice Department legal counsel wrote memos in 2002 and 2005 in an attempt to provide legal cover for these practices, which the administration called “enhanced interrogation techniques” rather than “torture.” In 2014, a scathing report by the Senate Intelligence Committee asserted that the post-9/11 torture

<sup>16</sup> United Nations General Assembly, “Universal Declaration of Human Rights” (1948; available at [www.un.org/en/documents/udhr/index.shtml](http://www.un.org/en/documents/udhr/index.shtml)).

<sup>17</sup> Amnesty International, “Torture in 2014: 30 Years of Broken Promises” (available at [www.amnestyusa.org/files/act400042014en.pdf](http://www.amnestyusa.org/files/act400042014en.pdf)).

<sup>18</sup> US Senate, Senate Select Committee on Intelligence, *Committee Study of the Central Intelligence Agency’s Detention and Interrogation Program* (declassified December 9, 2014; available at [www.intelligence.senate.gov/study2014/sscistudy1.pdf](http://www.intelligence.senate.gov/study2014/sscistudy1.pdf)).

techniques were more brutal than the CIA acknowledged to the Bush administration – so brutal, in fact, that some CIA personnel tried to have them halted, only to be rebuffed by senior officials who demanded their continuation. The report also went to great lengths to refute the agency's leading arguments justifying the harsh forms of interrogation: that the torture was both necessary for and effective in thwarting terrorist plots and capturing senior Al Qaeda figures, including Osama bin Laden.<sup>19</sup>

Involvement in torture by physicians and other medical professionals is also widespread.<sup>20</sup> Physicians may be asked to devise methods of torture, to assess whether a prisoner can be tortured without fatal consequences, to treat a prisoner so that torture can continue, and to falsify records and otherwise act so as to conceal torture. In the case of Guantanamo Bay, there is ongoing controversy over the role of psychologists and physicians in the design of interrogation methods, the treatment of tortured detainees, and the neglect or concealment of evidence of torture.<sup>21</sup> The UN's resolution on *Principles of Medical Ethics* states:

It is a gross contravention of medical ethics, as well as an offence under applicable international instruments, for health personnel, particularly physicians, to engage, actively or passively, in acts which constitute participation in, complicity in, incitement to or attempts to commit torture or other cruel, inhuman or degrading treatment or punishment.<sup>22</sup>

In Chapter 3, we analyzed moral rights as valid moral claims that protect important interests and ordinarily trump appeals to the general welfare. The moral right not to be tortured is a valid moral claim, grounded in nonmaleficence, that protects an individual's important interests in avoiding terrible pain and distress, injury, disability, and so on. Thus, to say that someone has a right not to be tortured implies that appeals to the general welfare are not – at least ordinarily – sufficient to justify overriding their

<sup>19</sup> Ibid. See also Mark Mazzetti, "Senate Panel Faults C.I.A. over Brutality and Deceit in Terrorism Interrogations," *New York Times* (December 10, 2004), A1.

<sup>20</sup> Jesper Sonntag, "Doctors' Involvement in Torture," *Torture* 18 (2008): 161–175.

<sup>21</sup> Vincent Iacopino and Stephen Xenakis, "Neglect of Medical Evidence of Torture in Guantanamo Bay: A Case Series," *PLoS Medicine* 8 (2011): e1001027; Task Force on Preserving Medical Professionalism in National Security Detention Centers, *Ethics Abandoned* (New York: Institute on Medicine as a Profession and the Open Society Foundation, 2013); David Hoffman et al., *Independent Review Relating to APA Ethics Guidelines, National Security Interrogations, and Torture* (Chicago: Sidley Austin LLP, 2015).

<sup>22</sup> United Nations, *Principles of Medical Ethics* (1982): Resolution 37/194 (Geneva: United Nations, 1982). Whether and under what conditions medical professionals may treat tortured prisoners in the face of possible complicity with the torturing regime is a complex question that we cannot address here (see Chiara Lepora and Joseph Millum, "The Tortured Patient: A Medical Dilemma," *Hastings Center Report* 41 [3] [2011]: 38–47).



moral claim not to be tortured. Imagine that police in a particular district calculated that they could produce more public benefit than harm, overall, by torturing gang members and using elicited information to make key arrests of gang leaders, thereby greatly reducing gang violence and drug trafficking in the district. Even if their calculations were accurate, this appeal to the public welfare would not justify torturing the gang members. As we noted in Chapter 3, *sometimes* rights give way when the difference in the net gain in public welfare becomes sufficiently large. This motivates the question: *Is it ever permissible to override someone's right not to be tortured?* We argue, "Yes, but only in principle; never in practice."

The classic case that is invoked to support the permissibility of torture is the ticking bomb case.<sup>23</sup> Imagine that government agents know the following:

1. There is a terrorist plot to explode a massive bomb, very soon, in a crowded urban center in their nation.
2. The plot has a high chance of succeeding if it is not prevented by their agents.
3. A convicted terrorist in their custody is known to possess information that would permit their agents to prevent detonation of the bomb.
4. Torturing this particular man has a high chance of eliciting the required information in time to prevent the tragedy.
5. There is no feasible alternative means to averting the tragedy.

If government agents really knew statements (1)–(5) to be true, then, we think, the government would be justified in applying torture because the harm involved in torturing this one man would be very small in comparison with the harm to be averted by preventing the explosion of a massive bomb in a crowded urban setting, no better option exists, and the right not to be tortured is not absolute. Even when we factor in additional, secondary harms that are fairly likely to occur – such as possible psychological harms to the agents of torture, societal distress that the government was willing to apply torture, and resentment on the part of persons sympathetic to the terrorist – the opportunity to prevent the disaster would appear

<sup>23</sup> See, e.g., Alan Dershowitz, *Why Terrorism Works* (New Haven, CT: Yale University Press), chap. 4, although Dershowitz argues for the permissibility of torture in real-life circumstances, not just in principle. In response to Dershowitz, Oren Gross defends an absolute legal ban on torture yet defends the very occasional resort to (real-life) torture – in an act of official disobedience – as a necessary means to averting a catastrophe ("Are Torture Warrants Warranted? Pragmatic Absolutism and Official Disobedience," *Minnesota Law Review* 88 [2004]: 1481–1555).

sufficiently valuable to dwarf the harms that might accrue if the individual is tortured.

We accept that torture would be justified in these circumstances. If any reader is inclined to disagree, we could raise the stakes higher and more fantastically by stipulating that it is not just a large population in an urban setting whose lives are at stake, but an entire nation or even all of humanity. We could even stipulate that the prospective torture victim himself will die if the bomb goes off, and that he himself is likely to suffer in dying from nuclear radiation far more than he would suffer if tortured. When the stakes are this high, it seems to us that it would be irrational and fanatical to deny the permissibility of torturing someone to prevent such cataclysmic results.

So we accept that torture can, in principle, be justified in some conceivable circumstances. But this is very different from claiming that torture is ever, in fact, justified in real life. We deny that it is actually justified in real life, which involves the world as we know it and human beings as we know them. Crucial differences between the scenario described above, in which torture would be permissible, and the actual world concern *the limits of human knowledge*<sup>24</sup> and *the frailties of human character*.<sup>25</sup>

Consider first the limits of knowledge. It is highly improbable that actual human beings would ever be in the position of knowing the truth of statements (1)–(5) enumerated above. To be sure, that a particular terrorist plot exists can sometimes be known, where intelligence is reliable and the evidence robust. (Of course, intelligence is often unreliable, as when the United States relied on faulty intelligence about supposed weapons of mass destruction in justifying an invasion of Iraq.) Perhaps, in favorable circumstances, we can also know that a particular plot has a high chance of succeeding if it is not thwarted. Much more doubtful is that agents will really know that a particular person – say, a suspected terrorist – has the crucial information. It is very difficult for individuals to know that someone else, who is not in a cooperative relationship with them, has information that they do not themselves possess. But suppose the captive does have the crucial information. Perhaps the most dubious of the five assumptions is that torture will yield the correct information and do so in time to prevent the tragedy. After all, the torture victim can stop the

<sup>24</sup> See Vittorio Bufacchi and Jean Maria Arrigo, “Torture, Terrorism and the State: A Refutation of the Ticking-Bomb Argument,” *Journal of Applied Philosophy* 23 (2006): 355–373, sect. V.

<sup>25</sup> Henry Shue discusses both types of limits in opposing the use of torture (“Torture in Dreamland: Disposing of the Ticking Bomb,” *Case Western Reserve Journal of International Law* 37 [2006]: 231–239).

torture, at least for the time being, by *appearing* to cooperate and *appearing* to supply the desired information. The final assumption enumerated above – that there is no other way of obtaining the desired information – also seems to us well worth questioning in practice.

In sum, officials cannot know in any given case that torturing a particular individual will enable averting a disaster and will be the only way of doing so. A rebuttal is available, however, to a consequentialist who focuses on probabilities and the magnitude of the disaster to be averted: “It is true that we cannot be absolutely certain that torturing a person of interest will work and be the only way to avert a catastrophe. But we’re talking about a *catastrophe*, or the possibility of one, not the prospect of saving a couple of people. Even if there were only a very small chance, say 5 percent, of being correct in a given case, there must be some magnitude of harm to be averted that would justify taking that chance. So your argument, while successfully showing that we will always lack 100 percent certainty about the relevant facts, does not succeed in showing that torture is always wrong in the real world.”

In response to this argument, we appeal to further consequentialist considerations. Suppose a government sanctioned torture, in the name of averting great harms, despite the impossibility of knowing in any given case both that doing so will be effective and that no less morally problematic alternative would work. What would happen? First, even if the agents of torture and their superiors had morally perfect intentions, there would be “false positives”: cases in which individuals are tortured uselessly (because the torture was ineffective) or needlessly (because torture was not necessary to avert the catastrophe). And because the chances of getting everything right in a given case are unlikely to be very high, false positives will almost certainly outnumber true positives or successes. If there is a 5 percent chance of being correct, then 95 percent of the time there will be unnecessary torture. Second, in the real world we cannot expect agents of torture and their superiors to have morally perfect intentions. If a government gives officials the prerogative to torture, they will surely use this prerogative more often than would be justified. As Henry Shue notes, “history does not present us with a government that used torture selectively and judiciously.”<sup>26</sup> And recent experience in the United States should remind us that human beings simply cannot be trusted with the

<sup>26</sup> *Ibid.*, 234.

power to torture.<sup>27</sup> Third, if a powerful nation such as the United States or United Kingdom permits torture, it sends a signal that even powerful nations cannot adequately defend their citizens without resorting to torture, suggesting a fortiori that less powerful nations cannot be expected to do so. In effect, the result is implicit permission to engage in torture. We believe that this message is more likely to lead to catastrophe than a strict policy of refraining from torture. For these reasons, we submit that even the most resourceful consequentialist defense of torture cannot succeed.

#### 4.5 Permissible Risks in Pediatric Research

Children, like adults, need validated medicines and medical procedures to be available for treatment or prevention of disease. Such validation requires rigorous biomedical research. Because children's bodies differ from adults' bodies in complex ways, this research typically requires pediatric subjects; research results based on adult subjects cannot simply be extrapolated to children. Most children cannot give valid consent to the intrusive and potentially risky interventions that biomedical research involves. But then when is it permissible to place some children at risk for the sake of benefiting the health of other children in the future? The answer, presumably, will depend on the amount of risks and the extent of the potential benefits. Children have a moral right to adequate protection from harm, but this must be considered together with the value of conducting the research.

What constitutes *adequate* protection in pediatric research? How much risk to pediatric subjects is ethically permissible? The answer depends partly on the circumstances. The most important distinction in this context is that between (1) *therapeutic* pediatric research, which offers the prospect of direct medical benefit to the pediatric subjects themselves, and (2) *nontherapeutic* pediatric research, which does not offer the prospect of such benefit to the subjects.

Consider, first, a ten-year-old boy who suffers from a form of leukemia for which no successful treatment is known. He is likely to die within months if he does not receive effective treatment, the only possibility of receiving which is in a single clinical trial. The risks of the trial drug are not insubstantial, but initial tests have suggested some promise of effectiveness. In this situation, the child's right to adequate protection from harm would

<sup>27</sup> See, e.g., Karen Greenberg and Joshua Dratel (eds.), *The Torture Papers* (New York: Cambridge University Press, 2005).

permit his entering a trial that poses considerable risks to his health so long as the expected benefit – from possibly saving his life – along with the risks is preferable to the risks and benefits of the only alternative: receiving no treatment. In general, if the risk/benefit ratio of entering a trial offering the prospect of direct medical benefit is better than the risk/benefit ratio of any alternative (and the risks of entering the trial are minimized with optimal research design and sound medical practices), then the pediatric subject's right to adequate protection has been respected. In therapeutic contexts where there are few options, adequate protection can be compatible with high risks.

Our conclusion here does not imply that it is permissible to accept high risks whenever some direct medical benefit is in prospect for pediatric subjects. It would be wrong to enroll a child in a trial in which the expected therapeutic benefit was relatively modest while the risks were extremely high. For example, if a promising medication offered a child relief from a case of atopic dermatitis (eczema) that standard treatments failed to ameliorate, but it posed a significant risk of a fatal allergic reaction, the prospect of direct medical benefit would likely not justify the risks. The prospect of benefit, taking into account both the magnitude and likelihood of benefit, must at least roughly offset the risks, taking into account the magnitude of possible harms and their likelihood of occurring.<sup>28</sup>

Consider another case. The parents of a healthy ten-year-old girl are thinking of enrolling her in a trial of a vaccine that is hoped to protect children against a chemical agent that might be used in a terrorist attack. The vaccine has already been found safe and effective in adults. The girl does not have a medical condition that might be treated by enrolling in the trial, and there is no particular reason to expect an attack anytime soon, so this trial would clearly count as nontherapeutic research. Because she does not stand to benefit medically from participating (setting aside the extremely remote possibility that a terrorist attack would occur, she would be exposed to the noxious agent, *and* the vaccine she had taken in the trial would protect her), the trial is unlikely to be in her interests. An adult may accept significant risks to benefit society through participation in research, but children lack the capacity to consent and therefore need special

<sup>28</sup> We say “roughly offset” in response to a suggestion by Dave Wendler (personal communication): that just as minimal (but nonzero) risk is permissible in *nontherapeutic* pediatric research, it would make sense to allow *a bit more* risk than is compensated for by the prospect of direct medical benefit in *therapeutic* pediatric research.

protection. What protection would be adequate for nontherapeutic pediatric research?

An approximate answer is suggested by our intuitive judgments regarding the risks parents can responsibly allow their children to undergo.<sup>29</sup> Because nontherapeutic research does not offer the prospect of direct medical benefit to pediatric subjects, we may think analogously about risks undertaken by children in nonmedical activities whose primary purpose is to benefit others. How much risk could parents responsibly allow their children to take in helping others? The following activities seem within the range of reasonable risk: delivering groceries by bicycle to an elderly neighbor; helping to rebuild a damaged building; and running an eight-kilometer race for charity in the rain.<sup>30</sup> These examples suggest that parents may responsibly allow their children to undergo *some nonnegligible risks* in activities that are primarily altruistic, yet the risks would still have to be *relatively minor*.

This way of thinking about permissible risk in pediatric research can be stated more precisely in terms of the *reasonable subject standard*, which we defend in Chapter 5 as the appropriate standard for medical decision-making when the informed-consent and advance-directive standards are inapplicable.<sup>31</sup> According to the reasonable subject standard, a surrogate decision-maker should decide on behalf of someone lacking decision-making capacity as that individual would decide if she were a rational agent choosing prudently within the constraints of morality. Prudent decision-making in the present context would require not accepting high risks in nontherapeutic trials. Meanwhile, beneficence requires that an agent make some significant contributions to the public good over the course of a lifetime (as discussed in Chapter 6). One way of partly discharging this obligation is to enroll in important nontherapeutic trials. So parents should have the discretion to enroll their children in nontherapeutic pediatric studies to the extent that doing so would be consistent with how a prudent person would discharge their duties to benefit society.

Current US regulations state four risk categories for research involving children.<sup>32</sup> They permit nontherapeutic research that poses “minimal

<sup>29</sup> This approach is further developed in David DeGrazia, Michelle Groman, and Lisa Lee, “Defining the Boundaries of a Right to Adequate Protection: A New Lens on Pediatric Research Ethics,” *Journal of Medicine and Philosophy* 42 (2017): 132–153.

<sup>30</sup> Cf. David Wendler, “A New Justification for Pediatric Research without the Potential for Clinical Benefit,” *American Journal of Bioethics* 12 (2012): 23–31.

<sup>31</sup> See also Joseph Millum, *The Moral Foundations of Parenthood* (New York: Oxford University Press, 2018), chap. 6.

<sup>32</sup> The regulations are found in 45 Code of Federal Regulations (CFR), Sections 46.404–46.407 (for the Department of Health and Human Services) and in 21 CFR, Sections 50.51–50.54 (for the Food and Drug Administration). Hereafter we will cite only 45 CFR 46.

risk” – where “minimal” is said to be no greater than the degree of risk that children face in everyday activities and in routine medical examinations.<sup>33</sup> (Follow-up discussions have led to general agreement that here we should imagine healthy children living in safe neighborhoods.) They also permit research that involves greater than minimal risk where the risks associated with participation are compensated “by the anticipated benefit to the subjects.”<sup>34</sup> This corresponds to what we called therapeutic research. In cases in which the research offers no prospect of direct medical benefit to pediatric subjects, but is deemed “likely to yield generalizable knowledge about the subject’s disorder or condition,” the permitted risk level is “a minor increase over minimal risk.”<sup>35</sup> However, we do not think the likelihood of illuminating the subject’s condition is a cogent justification for a higher risk level.<sup>36</sup> The justification we gave for exposing children to risk in nontherapeutic research rests on considerations of beneficence, which apply whether or not the individuals in need share a condition with their potential benefactor. We regard the minor-increase-over-minimal-risk category as justifiable independent of the disease status of the participants.

This third category is not universally acknowledged in regulations for pediatric research. Both Canada and the United Kingdom are more restrictive in terms of the risks that they allow children to undergo in research. Canada limits the involvement of children to research with minimal risks or where the benefits to participants outweigh the risks – that is, to the first two categories we just described.<sup>37</sup> The terminology in the United Kingdom is different than that used in North America. The Medical Research Council’s guidelines allow nontherapeutic research only where the risks are “minimal” or “low,” where “low risk” appears to include similar procedures to those regarded as “minimal risk” in Canada and the United States (e.g., blood draws).<sup>38</sup> Again, higher risks are permissible in the case where the prospect of direct benefits to participants outweighs the risks.

Even if the US category of “a minor increase over minimal risk” is something of an outlier internationally, we believe that – as commonly

<sup>33</sup> 45 CFR 46.404.    <sup>34</sup> 45 CFR 46.405.    <sup>35</sup> 45 CFR 46.406.

<sup>36</sup> Cf. David Wendler, Seema Shah, Amy Whittle, and Benjamin Wilfond, “Non-Beneficial Research with Individuals Who Cannot Consent: Is It Ethically Better to Enroll Healthy or Affected Individuals?,” *IRB* 25 (2003): 1–4.

<sup>37</sup> TCPS 2 (2018): 53.

<sup>38</sup> Medical Research Council, *Medical Research Involving Children* (2004) (available at <https://mrc.ukri.org/documents/pdf/medical-research-involving-children/>).

interpreted in practice – it is justified. This standard is commonly understood as permitting not only such procedures as blood draws (considered minimal risk) but also skin biopsies and chest X-rays, but not procedures much riskier than these.<sup>39</sup> We believe this level of risk is consistent with adequately protecting pediatric subjects from harm.

Nearly all pediatric studies that are approvable under current US regulations will fall under one of the three categories described: (1) therapeutic studies, (2) studies posing minimal risk, and (3) studies posing a minor increase over minimal risk but likely to yield generalizable knowledge about the subjects' medical condition. The regulations also mention a fourth, "wild-card" category of approvable pediatric studies: those posing more than minimal risk with no prospect of direct medical benefit to individual subjects and without the prospect of illuminating a medical condition that affects the individual subjects. Such a study is approvable if it presents a significant prospect of furthering the understanding, prevention, or treatment of a serious problem affecting the health or welfare of children, though approval would require national expert review rather than the usual review by a local institutional review board.<sup>40</sup> Earlier we considered a candidate study for approval under this provision: a pediatric study of a vaccine against a particular chemical agent that could be used in a terrorist attack. Assuming the risks would be more than minimal, and considering the fact that the potential participants do not have a medical condition that the trial is designed to illuminate, the risks they would face in this trial would be considered excessive and therefore unjustified – unless the study is approvable under the wild-card category.

We think that this fourth category identifies a class of permissible research studies, and offer several suggestions with regard to it. First, the regulations would be improved by specifying a limit to the risk levels to which children may be subjected in this category of studies, just as there are limits specified for the other approvable categories. Second, the risk ceiling should still be limited to those that parents could responsibly permit for their children in altruistic activities, consistent with the burdens that someone could reasonably be asked to assume on the basis of beneficence. Third, we anticipate that this limit may be higher than a minor increase over minimal risk, but not much higher. Fourth, as with

<sup>39</sup> For examples from one high-volume pediatric research institutional review board, see Children's Hospital of Philadelphia Research Institute, "Risk of Common Procedures" (2020) (available at <https://irb.research.chop.edu/risk-common-procedures>).

<sup>40</sup> 45 CFR 46.407.



other research studies, pediatric studies that deserve to be approved under the wild-card category are ones in which risks are minimized.

This last category raises an interesting question. It is easy to understand how adequate protection of children in therapeutic research allows greater risks than adequate protection permits in nontherapeutic research: a sufficient prospect of direct medical benefit compensates for higher risks. What may be difficult to understand is why the higher risk standard pertaining to the last category, which we have argued is compatible with children's right to adequate protection, is not the appropriate standard for *all nontherapeutic research*.

Our answer is that this higher risk standard would be the appropriate standard for all nontherapeutic pediatric research in certain ideal circumstances. These would be circumstances in which research institutions and investigators were not biased in favor of conducting research; the government and its officials reliably treated individual children with as much care and protection as good parents treat their children; and prospective subjects and their proxies never felt pressured by biomedical personnel into joining particular studies. Given that these ideal circumstances do not obtain, it is appropriate that in most instances children receive more protection than they would need in ideal circumstances. In this light, the more restrictive standard of a minor increase over minimal risk is the one we endorse as a default level of permissible risk. Occasionally, however, this additional downward (protective) pressure on permitted risk may be appropriately relaxed due to extraordinary circumstances. In these instances, the limits of permissible risk may – if necessary – be raised. But in these instances, as the regulations stipulate, national expert review is required. Thus, while the *substantive* protection is slightly relaxed with an increased level of permissible risk, the *procedural* protection increases.

Might there be rare cases in which it would be ethically permissible to conduct even riskier research on children where the expected social benefits are very high? Here our response is the same as it was to the question of whether there were exceptions to the right not to be tortured: in principle, yes, but not in practice. We can imagine circumstances in which (1) enrolling children in nontherapeutic research with net risks substantially higher than a minor increase over minimal risk would offer the prospect of the best overall consequences *by far*, while (2) decision-makers would have all the information needed to make this judgment reliably, and (3) they would possess the intellectual and moral character needed to resist institutional pressures biased in favor of research. But, in

the world as we know it with people as we know them, we believe it would be better to maintain a strict ceiling on the net risks to which children may be exposed.

Let us take stock. In this chapter, so far, we have articulated the ethical principle of nonmaleficence, investigated the concept of harm at the heart of this principle, and identified several general moral rules that capture different aspects of nonmaleficence. In addition, we have noted that nonmaleficence supports a variety of moral rights in the form of negative constraints, including a right not to be tortured. We have also found that nonmaleficence supports negative constraints and moral rights, regarding the imposition of risk, exploring in some detail the appropriate risk-related constraints in pediatric research. In our final application of the principle of nonmaleficence, we explore its implications for medical assistance-in-dying. We find, perhaps surprisingly, that whether death constitutes a harm depends not only on one's circumstances but also on one's values.

## 4.6 Medical Assistance-in-Dying

### *Conceptual and Legal Background*

Consider two types of harm. *Experiential harm* includes pain, distress, suffering, and other unpleasant experiences. Accordingly, one of the general moral rules that we derived from the principle of nonmaleficence was "Do not cause pain, suffering, or other experiential harm." And one of the commonly acknowledged goals of medicine is to alleviate pain and suffering. A second type of harm is *death*. In ordinary cases, death harms a human being by depriving them of the goods that their life would have contained had they lived. So one of the general rules we mentioned earlier was "Do not kill." And one of the goals of medicine, of course, is to preserve life.

In some cases, the objectives of avoiding or ameliorating experiential harm and of preventing death conflict because the continuation of life entails suffering due to painful or debilitating medical conditions. In many cases, the conflict is tolerable because the suffering is not so great as to call into question the assumption that continuing to live is in someone's best interests. This might be the case, for example, where imperfect recovery from a surgery has caused a low level of recurring pain that is not fully relievable with medications. But there are also cases in which the prospect of unrelenting, severe suffering causes the patient or their proxies to doubt that their remaining alive is worthwhile.

Such a patient might have an end-stage metastatic cancer that has defeated the best treatments and is causing pain and discomfort that pain medications cannot satisfactorily alleviate. Or they might be a patient with locked-in syndrome whose paralysis permits them to move only their eyelids, and who has determined after trying to adjust to their disability that they would prefer dying sooner than later. Cases could be multiplied indefinitely. What such patients have in common is that their situation, in their opinion, reverses the usual presumption that death constitutes a harm. In their view, on balance, the continuation of life constitutes a harm for them. Such cases have motivated the movement away from the traditional medical imperative of preserving human life in all cases toward assertions in law and medical ethics of “a right to die.” Here, we explore the ethics of pursuing various ways to bring about a patient’s death, with special attention to the concept of harm. We assume, unless otherwise specified, that the patient in question is a competent adult.

In a medical setting, there are various ways of bringing it about that a patient dies sooner rather than later. We may distinguish five modalities:

1. Forgoing of life support
2. Unintentional killing through escalating use of pain medications with the intention of relieving pain
3. Terminal sedation
4. Physician-assisted suicide
5. Euthanasia.

Each merits a brief explanation.

Medical life-support measures such as respirators, feeding tubes, and antibiotics for a potentially lethal infection can be forgone either through stopping their use or withholding them (not starting). Because it is well-established in law and ethics that a competent adult patient has a right to refuse medical treatment, it is widely accepted that such patients have the right to forgo life support.<sup>41</sup>

A distinct modality, widely accepted in practice yet less often discussed in the literature, involves unintentional killing or shortening of life. In these cases, a physician does not intend to end a patient’s life but is willing to increase the dosage of an opioid to reduce or eliminate pain *even if there is a significant chance that the medicine will induce respiratory failure, causing*

<sup>41</sup> See Ruth Faden and Tom Beauchamp, *A History and Theory of Informed Consent* (New York: Oxford University Press, 1986), Part II; and Gregory Pence, *Medical Ethics*, 6th ed. (New York: McGraw-Hill, 2011), Part I.

*the patient to die.* So long as the physician does not use more medication than she judges necessary to reduce suffering and does not intend to cause the patient's death, then if the patient dies as a result of opioid-induced respiratory failure, the killing is deemed unintentional and consistent with appropriate medical practice. In the legal jurisdictions of which we are aware, killing in these circumstances is not considered homicide under the law. Properly used, however, this method cannot guarantee a patient's death because pain relief might prove successful while the patient is alive.

In *terminal sedation*, a patient agrees to be sedated to the point of entering a coma, remaining in the coma due to continuing sedation, and to withholding of life-support measures. By this method, a patient can be assured of dying painlessly from an underlying illness or from dehydration.<sup>42</sup> Although the legal status of terminal sedation tends to be somewhat ambiguous, at least in certain cases it would appear to be legal by virtue of the legality of two separate acts: providing sedatives to a suffering patient and forgoing life support, including nutrition and hydration.<sup>43</sup>

A fourth modality is *physician-assisted suicide (PAS)*, in which a physician provides crucial support for a patient who commits suicide.<sup>44</sup> In order for an act to be an instance of suicide and not euthanasia, it must be the patient who performs the act that causes death. One way a physician may provide assistance is by writing a prescription for a lethal dose of medicine and providing instructions for how to take the medicine. A doctor might also supervise the process or even give a bedridden patient the pills and water to swallow. At the time of this writing PAS is legal in Belgium, Canada, Luxembourg, the Netherlands, Switzerland, Columbia, the Australian province of Victoria, and eight US states and the District of Columbia. PAS remains illegal in, for example, New Zealand, the United Kingdom, and most US states.

The most controversial mode of medically bringing about a patient's death is *euthanasia*, in which a physician intentionally causes a patient's

<sup>42</sup> Any patients who prefer to remain conscious during their remaining days can refuse hydration and be sure of dying from dehydration within two weeks. Palliative measures can reduce pain and discomfort without inducing coma.

<sup>43</sup> In 1997 the US Supreme Court provided at least some support for terminal sedation: "A patient who is suffering from a terminal illness and who is experiencing great pain has no legal barriers to obtaining medication from qualified physicians, even to the point of causing unconsciousness and hastening death" (Justice Sandra Day O'Connor, Concurring Opinion in *Glucksberg v. Washington* [521 U.S. 702, 737 (1997)]). In 2016 France legalized terminal sedation in cases of refractory pain and imminent death within two weeks (Ruth Horn, "The 'French Exception': The Right to Continuous Deep Sedation at the End of Life," *Journal of Medical Ethics* 44 [2018]: 204–205).

<sup>44</sup> Our use of the term "physician" in "physician-assisted suicide" refers – as it does throughout the book – to any medical doctor.

death. The standard method of euthanasia involves injection of a lethal agent, such as potassium chloride.<sup>45</sup> Of the five modalities under consideration, it is the only one (at least on standard interpretations) in which a physician intentionally kills a patient.<sup>46</sup> Because of the moral rule against intentionally killing human beings and corresponding laws against homicide, acceptance of euthanasia is considered somewhat radical. The practice is currently illegal in the United Kingdom and the United States, for example, but legal in Belgium, Canada, Columbia, Luxembourg, and the Netherlands. In the 1990s euthanasia was also briefly legal in the Northern Territory of Australia.

### *Conventional Thinking about This Issue and Its Limitations*

All five modalities are ways in which a patient's death can be made to occur earlier than it would with full application of life-support measures. But conventional thinking about the ethics of end-of-life decision-making does not regard all of the modalities as morally equivalent. It is widely believed – among scholars in medical ethics and among the general public – that forgoing life support, unintentional killing through appropriate use of opioids, and terminal sedation can be morally permissible, assuming certain conditions are met. It is also widely – though by no means universally – thought that euthanasia is morally impermissible under any circumstances and should be illegal. Attitudes about physician-assisted suicide are more ambivalent. PAS is legal in some, but not all US states, legal in Canada, and illegal in the United Kingdom.”

Reflections on nonmaleficence and the nature of harm put the coherence of this conventional thinking in doubt. Consider this view's prohibition of euthanasia. At first glance, the prohibition seems reasonable insofar as nonmaleficence supports a rule forbidding the killing of human beings. Reflection on unintentional killing in pursuit of the legitimate goal of pain relief motivates specifying the rule so that it forbids only *intentional* killing. Reflection on killing as a last resort in self-defense against an aggressor

<sup>45</sup> This practice is sometimes called “active euthanasia” and contrasted with “passive euthanasia,” the forgoing of life support. Probably due to widespread acceptance of the latter practice and (for some people) negative connotations of “euthanasia,” the term “passive euthanasia” is rarely used today. So we reserve the term “euthanasia” for intentional medical killing of a patient by a physician.

<sup>46</sup> Some commentators argue that terminating life support, in some circumstances, intentionally causes the patient's death and therefore constitutes intentional killing. It has also been argued that a patient's decision to forgo life support can, depending on details, constitute a decision to commit suicide, in which case a physician's complying with the request constitutes assisted suicide. We set aside these conceptual complexities.

motivates a further specification, so that the rule forbids the intentional killing of *innocent* human beings. Even with these refinements, prohibiting euthanasia seems quite reasonable, because euthanasia standardly involves the *intentional* killing of *innocent* human beings. And that, according to common morality, is wrong.<sup>47</sup>

But we must ask *why* common morality considers such killing wrong. The most cogent answer is that intentionally killing innocent people is wrong because (1) death is a terrible harm, (2) killing violates a victim's rights, (3) the victim's innocence removes special justifications of killing (e.g., self-defense), and (4) the intentional nature of the act removes excuses that may apply to accidental killing (e.g., in pursuit of pain relief). Focus now on reasons (1) and (2), which underlie the general pro tanto wrongness of killing human beings. We are considering situations in which patients give *valid consent* to administration of a lethal injection. So the second reason does not apply. Meanwhile, we have been considering cases in which a patient does not regard staying alive, given her condition and prognosis, as worthwhile in view of the suffering involved. Given our subjective view of well-being, which we defend in Chapter 8, the counterfactual account of harm has a momentous implication. Provided that the patient correctly judges that her life will involve more suffering than enjoyment and her desire to die is suitably informed and responsive to her values, *death is not a harm*. Indeed, in these cases, it is the continuation of life that is harmful. A "good death" – the literal meaning of "euthanasia" – would put an end to the harm.

The recognition that death is not a harm in the relevant set of cases, while euthanasia would be respectful of patients' autonomy, suggests that the basic reasons for opposing the intentional killing of innocent human beings simply do not apply. Why, then, do so many people and so many jurisdictions favor the prohibition of euthanasia? Is there any cogent basis for doing so?

One possibility is that euthanasia is to be condemned on the basis of a moral principle of *avoiding killing* that is independent of nonmaleficence. The idea would be that intentionally killing a human being is at least pro

<sup>47</sup> In this discussion, we bracket the morality of suicide. Some thinkers consider it wrongful to end one's own life, and this would, of course, have implications for the ethics of others' helping one to end one's life. Interested readers might turn to Margaret Battin, *Ethical Issues in Suicide* (Englewood Cliffs, NJ: Prentice Hall, 1995); Tom Beauchamp and Robert Veatch (eds.), *Ethical Issues in Death and Dying*, 2nd ed. (Upper Saddle River, NJ: Prentice Hall, 1996), Part 3; and Margaret Battin (ed.), *The Ethics of Suicide* (New York: Oxford University Press, 2015).

tanto wrong, independently of whether it would harm or violate the rights of the individual in question.<sup>48</sup>

We find this an unpromising strategy for defending a prohibition on euthanasia. It seems that, when killing is wrong, it should be possible to say why it is wrong. Consider that killing human beings is usually wrong, killing nonhuman animals is often wrong, and killing some living things (e.g., bacteria, ivy) doesn't seem at all morally problematic. These points are easily understood on the basis of differences in the harms or rights violations typically involved in each case. Human death is ordinarily a very substantial harm, and killing an innocent human person against their will violates their rights. The death of sentient nonhuman animals is ordinarily somewhat harmful to them. The strength of our reasons not to kill nonhuman animals plausibly depends on both the magnitude of this harm and our views about whether they have rights. The death of insentient creatures is not plausibly regarded as harmful to them, and their lack of sentience means there is no subject who could have rights. The harm of death and the violation of rights involved in killing some creatures seem sufficient to explain what is wrong about killing. But this explanation would not support a prohibition of euthanasia – in which death is not a harm and the patient's rights against bodily trespass have been waived.

### *Defense of a More Progressive View*

The foregoing arguments suggest that euthanasia is morally justified, in principle, in some circumstances. We mean two things in saying that euthanasia is sometimes justified *in principle*: (1) negatively, that nothing about the act of euthanasia justifies an automatic prohibition, and (2), positively, that strong reasons – concerning avoiding the harm of continuing life and promotion of patient autonomy – count in favor of allowing euthanasia in some circumstances. But a type of action might be justified in principle without being justified in practice, after accounting for factors such as the possibility of well-intentioned error and some people's willingness to act wrongly. We argued that this is true of torture.

So let us consider practical concerns that might justify prohibiting euthanasia. Some of these practical concerns are especially acute in the United States, which has a highly inequitable distribution of income,

<sup>48</sup> For a defense of this view, see Robert Veatch, *A Theory of Medical Ethics* (New York: Basic Books, 1981), chap. 10. Cf. William Frankena, *Ethics*, 2nd ed. (Englewood Cliffs, NJ: Prentice-Hall, 1973), 55–56.

wealth, and access to health care (see Chapter 6). The major concerns about the legalization and practice of euthanasia, among those who acknowledge that this act may sometimes be justified in principle, are concerns about five possibilities: (1) error, (2) coercion or undue pressure, (3) abuse of the discretion afforded to physicians, (4) a slippery slope toward unacceptable practices or policies, or (5) inconsistency with the role morality of clinicians. We address these possibilities in turn.

The most important type of error in this context is incorrectly thinking that a particular person is an appropriate candidate for euthanasia. Suppose, for example, that a patient is suffering terribly, and professionals involved in the patient's care believe there is no realistic hope of relief – although in fact there is a method of pain alleviation available that caretakers have not attempted or considered. With the best palliative care, let us suppose, this person would not want to die. It would be tragic for them to die unnecessarily as a result of this error. Of course, such an error would be equally tragic if it led to a patient's death via terminal sedation or simply the withdrawal of life support.

The second concern is that, once euthanasia is an option, individuals may be coerced or otherwise pressured into agreeing to undergo it. Suppose a nursing home patient lacks insurance that covers nursing home care and has multiple medical problems requiring costly services to address. The patient's care providers now face the choice of either not providing needed services or providing them without compensation. Such a patient might feel pressure to consent to euthanasia, for example, as a result of cues from staff members that the patient's care is very costly. This concern will be particularly pressing in countries like the United States where many people are uninsured or underinsured. But they will also arise in countries with universal health care, like Canada and the United Kingdom, where public provision of some services is limited. Even in a country with a strong health care system, individuals may feel pressure to consent to euthanasia on the basis of feeling that they pose a disagreeable burden on family members and hospital staff. Pressure to agree to euthanasia might undermine the voluntariness of the decision, if it involves a threat (see Chapter 5). But even pressure that does not have this effect – simply being reminded of the sacrifices others have to make and the option of euthanasia to relieve them – could lead someone to make a decision that they would rather not make.<sup>49</sup> Death would then be contrary to the

<sup>49</sup> J. David Velleman develops an argument against giving patients the right to request euthanasia on the grounds that providing that option may make people worse off, even though it may be rational



patient's wishes and potentially seriously harmful insofar as his death, given his values and priorities, would deprive him of a life worth continuing. Once again, this concern also applies to terminal sedation and withdrawal of life support, both widely accepted ways to hasten death.

In addition to concerns about patients being pressured into accepting euthanasia, there are concerns about possible abuse. Some health professionals, knowing that euthanasia is an accepted practice, may commit involuntary euthanasia – killing a patient *against* the patient's wishes – and claim that the patient had validly consented. This is a truly disturbing possibility. Like the other concerns, though, it applies to even the most widely accepted modalities for hastening death. For example, a health professional might also discontinue life support without the patient's permission and then claim that the patient had consented to the decision.

A fourth concern is the possibility of a slippery slope. The idea here is that institutionalization of morally acceptable forms of euthanasia might lead to the widespread practice and perhaps institutionalization of forms of euthanasia that are morally wrong. An early critic of euthanasia, Yale Kamisar, referred to “the danger that legal machinery initially designed to kill those who are a nuisance to themselves may someday engulf those who are a nuisance to others.”<sup>50</sup>

In order to be persuasive, a slippery slope argument against euthanasia must combine a reasonable projection about where acceptable euthanasia practices are likely to lead and a considered moral judgment that the end point is morally unacceptable. In other words, a strong slippery slope argument against euthanasia must be both empirically and morally plausible. Let us consider three specific “slopes” on which acceptable euthanasia practices might be thought to slide toward unacceptable ones.

First, in terms of *medical condition*, a policy that requires terminal illness as a condition for euthanasia might slide to a practice or policy in which terminal illness is not required, greatly expanding the pool of candidates for euthanasia. Second, regarding the *basis for decision-making*, a euthanasia policy that requires consent might slide, over time, to permitting cases in which consent is impossible but an advance directive had earlier been completed and now applies; then later, to permitting cases in which consent is impossible and no advance directive had been left, but there

for them to exercise the option once it is available (“Against the Right to Die,” *Journal of Medicine and Philosophy* 17 [1992]: 665–681).

<sup>50</sup> “Some Non-Religious Views against Proposed ‘Mercy-Killing’ Legislation,” *Minnesota Law Review* 42 (1958): 969–1042. Cited in L. W. Sumner, *Assisted Death* (Oxford: Clarendon, 2011), 175.

are thought to be adequate grounds for euthanasia based on the patient's previously expressed preferences; and, finally, to permitting cases in which no such judgment can be substantiated, but death is judged to be in the patient's interests. Third, regarding the *patient's age*, a policy that initially includes only adults as eligible for euthanasia might slide, over time, to include adolescents and then younger children and finally infants.

The three slippery slope arguments just presented are empirically fairly plausible. Indeed, in at least Belgium and the Netherlands, all three slides have already occurred: terminal illness is not required; decision-making standards less demanding than informed consent have been accepted in certain cases; and there is euthanasia involving minors.<sup>51</sup> Even in more "conservative" jurisdictions in which there were greater emphasis on safeguards it would seem reasonable to judge that each slide down a slippery slope is *quite possible and not at all far-fetched*. But there remains the question of whether these arguments are morally plausible – that is, whether the bottom of each projected slide is clearly morally wrong. We answer negatively for all three.

In terms of the patient's medical condition, we favor not requiring terminal illness. Many patients who are appropriate candidates for euthanasia, such as individuals with cancer or a progressive neurological disorder, will be terminally ill. But, in our view, some patients who are appropriate candidates for euthanasia lack a terminal illness but are suffering terribly, with no reasonable prospect of relief, as a result of a medical condition they cannot bear. This might be the situation of certain patients who have medical conditions such as locked-in syndrome or other severe forms of paralysis and have grown weary of struggling with their physical limitations. So we do not regard the slide in question as morally unacceptable. With sufficient safeguards, it is appropriate not to require terminal illness.<sup>52</sup>

The other possible slides may give greater pause. Yet we believe it would be ethically permissible to permit euthanasia in some cases in which consent cannot be obtained. Euthanasia would be morally permissible,

<sup>51</sup> See, e.g., Andrew Siegel, Dominic Sisti, and Arthur Caplan, "Pediatric Euthanasia in Belgium: Disturbing Developments," *JAMA* 311 (2014): 1963–1964; Barron Lerner and Arthur Caplan, "Euthanasia in Belgium and the Netherlands: On a Slippery Slope?," *JAMA Internal Medicine* 175 (2015): 1640–1641; and Scott Kim, Raymond De Vries, and John Peteet, "Euthanasia and Assisted Suicide of Patients with Psychiatric Disorders in the Netherlands 2011 to 2014," *JAMA Psychiatry* 73 (2016): 362–368.

<sup>52</sup> For arguments in support of requiring terminal illness as a condition for any form of physician-assisted death, see Lynn Jansen, Steven Wall, and Franklin Miller, "Drawing the Line of Physician-Assisted Death," *Journal of Medical Ethics* 45 (2019): 190–197.

in our view, where a *valid* advance directive was completed, was never put in doubt by an apparent change of mind, and now *clearly applies* to the patient's situation and calls for euthanasia. For example, a patient's directive might call for euthanasia in the event that her inoperable brain tumor has deprived her of medical decision-making capacity, apparently irreversibly, and she appears to be suffering greatly with no realistic prospect of relief. In these cases, euthanasia is justified on the basis of an autonomous decision that the patient made at an earlier time – rather than on the basis of a judgment about what she would have wanted or what is best for her. So we are confident about the permissibility of euthanasia in some cases in which a patient currently lacks decision-making capacity.

We are less confident about permitting euthanasia in cases that lack any autonomous authorization (via consent or an advance directive) by the patient. We believe that euthanasia *might* occasionally be permissible, given the right sorts of safeguards, on the basis of the reasonable subject standard. In any such cases, prudential considerations on behalf of an incapacitated adult or a minor would favor a quick, minimally painful exit from life, which could be achieved with either euthanasia or terminal sedation. Chapter 8 discusses a small number of cases involving severely impaired newborns in which euthanasia might be appropriate from the standpoint of our theory.

At this point, even those who agree with us that euthanasia is not inherently wrong may challenge our position on the following grounds: *With each step away from the decision-making authority that accompanies an individual informed consent process, the possibilities of error or abuse increase.* The power of decision-making moves farther away from the patient's current will and toward other people's interpretations of what he wanted or, where there was no advance directive, what a reasonable choice for him would be.

Yet these possibilities of error and abuse already attend widely accepted modalities for hastening death, including withdrawal of life support. In addition, more restrictive policies also have a terrible downside: preventing the death of individuals for whom continuing life is worse than ending it. The best approach is not to prohibit euthanasia but to institutionalize safeguards that deal sensibly with these areas of concern. Suppose, for example, that a child suffering from an intractable type of leukemia with no hope of relief (while alive) satisfied all reasonable criteria for euthanasia in principle. If euthanasia were precluded for minors, then what options would he have? He might commit suicide with the help of a physician (if this were permitted). But that expectation might be cruel for a child.

He could be subject to pain relief that may or may not be effective and may or may not end his life, which seems less than optimal. He could undergo terminal sedation, in which case it is unclear how this has any advantage over euthanasia. Or life support could be withdrawn without sedating him to unconsciousness, in which case he would suffer more than necessary. Once we understand that euthanasia is not wrong in principle and that a practice of euthanasia is generally no more vulnerable to error, undue pressure, or abuse than clearly acceptable modalities of hastening death, we must accept that euthanasia is acceptable in practice. The key is to implement safeguards that best reconcile the importance of protecting vulnerable patients while permitting a quick exit from a life whose continuation is harmful to the patient.

A somewhat different concern remains: that euthanasia might be inconsistent with the role morality of clinicians – that is, with the professional ethics of the doctors, nurses, and others who would be involved in its administration. Some opponents of euthanasia claim that the role morality of physicians, for example, includes not only certain traditional goals, as discussed at the outset of this section, but also a negative constraint: “the inviolable rule that physicians heal and palliate suffering but *never intentionally inflict death*.”<sup>53</sup> Another group of authors argues, “The essence of medicine is healing, managing pain, and alleviating suffering. Doctors [participating in PAS or euthanasia] jeopardize the moral integrity of the medical profession.”<sup>54</sup> However, these claims beg the question of the goals of medicine in assuming that euthanasia is inconsistent with those goals. An international group of scholars convened by the Hastings Center, a respected bioethics think tank, listed among the goals of medicine the relief of suffering caused by medical maladies, the care of patients who cannot be cured of their maladies, and the pursuit of a peaceful death.<sup>55</sup> Euthanasia can serve all three of these goals. In our estimation, appeals to the role morality of clinicians do not undermine our arguments in favor of euthanasia in some circumstances.

### *Policy Suggestions*

Although we cannot offer a full characterization of an adequate policy here, we offer several suggestions for a policy that permits euthanasia.

<sup>53</sup> C. L. Sprung et al., “Physician-Assisted Suicide and Euthanasia: Emerging Issues from a Global Perspective,” *Journal of Palliative Care* 33 (2018): 197–203, at 197 (emphasis added).

<sup>54</sup> *Ibid.*, 200.

<sup>55</sup> See Daniel Callahan, “The Goals of Medicine: Setting New Priorities,” *Hastings Center Report* 25 (6) (1996): S1–S26, at S10–S13.

- *Physician-assisted suicide would be legal (under constraints similar to those for euthanasia) and, in most cases involving adults, encouraged as an alternative to euthanasia. Euthanasia would be permitted when PAS was either infeasible due to the patient's condition, inappropriate due to the patient's young age (if euthanasia is permitted for minors), or unpalatable to the patient for some other reason.* The rationale for a presumption favoring PAS for competent adults is that it affords the patient slightly more control over the final act that causes death. An example of a condition making PAS infeasible is one in which the patient cannot swallow. The reason PAS may be inappropriate for children is that the prospect of committing suicide may be emotionally overwhelming to them. An example of a reason that PAS may be unpalatable to a given patient is that she finds the prospect of swallowing an enormous number of pills (if that is required in her jurisdiction) excessively arduous.
- *Euthanasia would generally be permitted only if the patient is suffering badly enough that her life is reasonably judged not worth living, the suffering is due to one or more medical conditions, and there is no reasonable prospect of recovery or of relief from the suffering while alive.* What makes it the case that continued life constitutes a harm rather than a benefit for the patient is the presence of great suffering without the prospect of relief, not the length of life remaining. So we do not endorse a requirement of terminal illness. Nor would we restrict the causes of suffering to "physical illness." It is possible that psychiatric conditions could cause sufficient suffering with no realistic prospect of relief (see below). The reason to require that the suffering be due to medical conditions is that other reasons for preferring a quick death are not the business of practicing physicians. If medical centers became places in which individuals could receive a quick exit from life for reasons unrelated to medicine, this development might badly damage public trust in the medical profession.<sup>56</sup>
- *In rare instances, euthanasia may be permissible in a patient who is not currently suffering badly enough but is expected – on the basis of highly compelling evidence – to experience sufficiently bad suffering in the future.* There is necessarily some risk of error in permitting euthanasia on the basis of a *prediction* of the degree to which someone will otherwise suffer. But we believe that, in some rare cases, the risk of someone's

<sup>56</sup> See Franklin Miller, "Should a Legal Option of Physician-Assisted Death Include Those Who Are 'Tired of Life'?", *Perspectives in Biology and Medicine* 59 (2016): 351–363.

suffering terribly and unnecessarily is even graver. This may be true of some impaired infants whose conditions worsen relentlessly over time. As we discuss in Chapter 8, death as a self-aware child may be far more traumatic than death as an unaware infant.

- *The patient must either (1) provide valid consent for euthanasia in a nonpressured environment with a thorough discussion of alternatives, including all available palliative care and terminal sedation, and reaffirm consent after a waiting period (perhaps a week); (2) have completed a valid advance directive that now unambiguously applies; or (3) be determined by a legally authorized proxy to be in a state in which the reasonable subject standard recommends euthanasia.* The third condition is one for which particularly stringent safeguards would be needed, including independent third parties who could confirm that the proxy decision-maker was deciding according to the appropriate standard and taking into account all the relevant facts.
- *An independent physician must confirm the patient's condition and prognosis and, if the patient retains decision-making capacity or presents a reasonable hope of regaining capacity, must rule out reversible depression.* Confirmation by an independent physician helps to reduce the possibilities of error or abuse. The relevance of reversible depression is that a patient, while depressed, may see less value in remaining alive than a patient who recovers from depression and is closer to his psychological and cognitive baseline. It is worth noting here that euthanasia in response to suffering due to psychiatric conditions – not only depression but also, for example, schizophrenia, personality disorders, and posttraumatic stress disorder – is a highly complex and hotly contested issue.<sup>57</sup> Some authors argue that euthanasia (and perhaps also physician-assisted suicide) should not be available on the basis of suffering due to psychiatric disorders.<sup>58</sup> We have reached the opposite judgment – where the other conditions detailed above are met. Our position rests significantly on the thesis that the value basis for determining whether a competent adult's life is worth continuing is that individual's well-informed judgment. It also depends on a fair measure

<sup>57</sup> Marie Nicolini et al., "Should Euthanasia and Assisted Suicide for Psychiatric Disorders Be Permitted? A Systematic Review of Reasons," *Psychological Medicine* 50 (2020): 1241–1256.

<sup>58</sup> See, e.g., Scott Kim and Trudo Lemmens, "Should Assisted Dying for Psychiatric Disorders Be Legalized in Canada?," *Canadian Medical Association Journal* 188 (October 4, 2016): E337–E339; and, for empirical data that can motivate this view, Kim, De Vries, and Peteet, "Euthanasia and Assisted Suicide of Patients with Psychiatric Disorders in the Netherlands 2011 to 2014."

of confidence in the possibility of adequate safeguards to protect adequately against errors and abuse.

This is a very rough sketch of the policy we recommend. At this point, a reader might wonder why we have defended a policy permitting euthanasia when terminal sedation, which offers equally fast relief from suffering, is already available in some jurisdictions. Our answer has three parts. First, although terminal sedation has the practical advantage of already being legal, *it has no ethical advantage over euthanasia*. The cases in which either modality would be appropriate are cases in which it is continued life, rather than death, that would be harmful to the patient. And terminal sedation does not afford more control to the patient than euthanasia does. By contrast, PAS does afford the patient a bit more control at the time of causing death and so is appropriately encouraged as an alternative to euthanasia. The second reason for permitting euthanasia despite the availability of terminal sedation is that many patients, and their families, may prefer that the patient undergo a “clean” death rather than dying more slowly from dehydration or an underlying medical condition. Even though the dying process would be unconscious, sparing the patient from suffering and awareness of his deterioration, the patient and his family may find the prospect of dying this way repugnant. We believe this sensibility ought to be accommodated. Third, it is wasteful to devote costly medical resources to keeping someone alive as she gradually dehydrates to the point of death if that individual meets appropriate criteria for euthanasia.