

## CHAPTER 1

# Defining Ethics

Perhaps the fundamental distinction in any study of ethics is that between descriptive and normative ethics. Clearly, ethical evaluations always involve norms. However, a study of ethical evaluations – for example, an ethnographic account of the moral principles affirmed by a particular group – need not take any stand on the “true” value of the acts or conditions being evaluated. For example, when the Pope characterizes abortion as sinful, he is engaging in normative ethics. In contrast, when a historian is surveying the development of Catholic views on abortion, he or she is engaged in descriptive ethics (unless, of course, he or she goes beyond a historical analysis to take a stand on the issue).

This distinction is crucial to any treatment of ethics, and it is important that authors be clear about precisely what project they are engaged in. As I have already explained, the first half of the present volume addresses the descriptive task, while the second half takes up normative ethics. In the course of the first four chapters, then, I hope to identify some key features of ethical thought and feeling in general and, perhaps more importantly, to isolate patterns that differentiate distinct, often opposed, ways of thinking about ethics. Specifically, I hope to clarify, systematize, and explain what I take to be two levels of ethical thought and action. The first level, what I call “fundamental orientation,” is very broadly defined by the setting of a few parameters in individual ethical response.<sup>1</sup> That is

<sup>1</sup> In speaking of principles and parameters, I am obviously drawing on some of Noam Chomsky’s work and the mathematical models he took up to think about language variation. This work was brought into the study of ethics by Marc Hauser. My account therefore has some points of contact with those of Chomsky and Hauser. However, it is important to point out right from the start that my use of a principles and parameters framework is very different from that of either author. Most important, I am not positing innate ethical principles with associated sets of variables. Rather, I see ethical attitudes and values as arising from the interaction of a range of nonethical processes, some biologically innate, some developmental, some deriving from group dynamics, and so on. In this way, ethics is not, so to speak, a biological system. It is rather a set of evaluative responses and behaviors that we can categorize together but that are not a “natural kind” in the sense of phenomena that play a coherent causal role in evolution. (In contrast, various nonethical processes that contribute to ethics do form natural kinds

the focus of the current chapter (with further development and literary illustration in the following chapter). In the third and fourth chapters, I will turn to the greater specification of ethical response in prototype assimilation, specifically attending to the role of story prototypes in construing ethically relevant situations and organizing our ethical motivations and actions in those situations.

In both the case of fundamental orientation and the case of prototype assimilation, I am relying on and extending research and theorization in cognitive science, affective science, and narratology, as well as ethical theory. For the broad division between rules and prototypes, I am drawing on work in categorization, as ethical evaluation centrally involves ethical categorization. In particular, I take up research indicating that we organize categories in three ways (Murphy and Hoffman 2012, 166), as follows: (1) by rules with variables (alternatively, principles and parameters), (2) by approximation to weighted averages (prototypes; for an introduction to prototypes, see Rosch 2011), and (3) by particular cases (exemplars). In the sciences, we typically try to define categories strictly in terms of rules with variables. For example, a chemical element has a fixed number of protons, but isotopes of the element vary in the number of neutrons. In ordinary speech, we tend to rely more on average cases. To take the standard example, we tend to think of birds as, roughly, things like robins. Thus, some cases of birds are more “bird-like” or prototypical than others. For instance, we would consider a sparrow to be a “better” case of a bird than a chicken or a penguin. The prototype is a *weighted* average because the psychological processes that produce prototypes enhance the contrast with opposed categories. To take another standard example, lettuce is identified as a highly prototypical diet food (Kahneman and Miller 1986, 143). But lettuce’s zero-calorie value clearly enhances the contrast between diet and non-diet foods. Diet foods have fewer calories than comparable non-diet foods, but they almost always have some calories; their average caloric value is not zero. Similarly, our prototypes for male and female tend to be more “masculine” and more “feminine” than average. Categorization by exemplars – similarity to a prominent, specific case – is very important, including for ethical study. As already noted, I will not discuss it as a theoretical topic here. However, all the literary works I consider are particular cases, thus operating as

in this sense; these nonethical processes would include theory of mind [ToM], affective empathy, particular causal inference, and others.) Moreover, I am skeptical about the ways in which Hauser – and, indeed, most evolutionary and empirical ethicists – gathers relevant data. For example, I do not trust the results of surveys bearing on trolley problems, as will be clear in what follows.

exemplars – or as composed of exemplars – insofar as they, in their specificity, have consequences for ethical response.

The varieties of categorization are crucial to this study because ethical response is, again, inseparable from categorization. I respond to a situation as an ethical dilemma only if I categorize it as such. This does not mean that I have to self-consciously think, “This is an ethical problem” or the like. It merely means that the situation has to activate networks of thought and emotion that are connected via ethical categories, which may be rule-based, prototype-based, or exemplar-based. To take a simple example, suppose Jones is engaged in video piracy. Someone convinces him that he may be caught, so he stops out of fear of unaffordable fines. Smith too is engaged in video piracy. Someone convinces him that this hurts the video artists. Smith feels guilty and stops. It seems clear that Jones’s discontinuation of piracy does not involve categorization of piracy as immoral and is not an ethical response (though we would presumably say that the behavior is the same as would have resulted from an ethical response). In contrast, it seems clear that Smith’s discontinuation is an ethical response. Moreover, I do not believe that this requires Smith to think anything explicitly ethical. The categorization operates as soon as he (tacitly) connects the piracy with its effects on people and feels guilty. In terms of the analysis developed later in this chapter, Smith’s response relies on a fundamental (ethical) orientation toward avoiding harm to others if one has no (nonegocentric) justification for inflicting that harm.

### **Some Complications of Descriptive Ethics**

Again, this and the following three chapters treat descriptive ethics, focusing on some significant, literature-relevant consequences of categorization processes. However, as noted briefly in the introduction, there are difficulties with entirely segregating descriptive ethics from normative ethics. Specifically, I (like everyone) have my own preferences regarding ethical values, and I am necessarily concerned with the normative issues that are addressed by the ethical responses I am describing. These preferences turn up in two ways. First, and most obviously, they turn up in my explicit, self-conscious statements about, and arguments regarding, moral values that I advocate. These statements are part of the book’s attention to normative ethics. These self-conscious attitudes might in principle distort my descriptive and explanatory account of ethics. But it does not seem terribly hard to segregate my explicit preferences from, say, claims that various sorts of ethical preference (including those I advocate and those I reject) may be understood in terms of narrative prototypes and dominant emotion systems. For example, my obvious leaning toward the ethics associated with romantic emplotment does not appear to inhibit my

identification and analysis of other ethical attitudes, such as those associated with heroic employment.

The second way in which my own ethical values bear on the present study is more implicit, and thus potentially more troublesome. It concerns the definition of just what counts as ethical. It might seem at first that what counts descriptively as ethical is easy to ascertain – simply include whatever people use the term “ethical” or “unethical” to refer to. Unfortunately, things do not work out that neatly. The most obvious problem is that different languages use different words, and it is often difficult or impossible to find things that we would consider exact equivalents. The less obvious problem is that even people who use the same labels – such as “ethical” – differ in their understanding of the meaning and extension of those labels. (The meaning of a label – or lexical item – is the complex of concepts connected with it; the extension is the set of physical, mental, abstract, or other objects to which it may be used to refer.) It is, then, generally not possible to simply take what everyone refers to as “ethical” or “unethical” and treat the resulting collection as delimiting a unified social or psychological phenomenon.

One result of the preceding points is that writers commonly rely to some extent on their own moral attitudes in deciding even what is to count as ethics. To take a simple example, Jones might use the word “unethical” to cover cases of behavior that are harmless but violate what he admits are changeable social conventions. In contrast, Smith might use different terms – say, “ethical” and “conventional” – for acts that Jones labels “ethical.” There are at least two ways of treating this contradiction. One might assume that for Jones, ethics is in reality much larger than it is for Smith. However, this risks overextending Jones’s view of ethics, which may in practice be more discriminating than his speech indicates, and indeed it may be more discriminating than he is aware. The other way of treating the contradiction is to assume that Smith’s more fine-grained division is shared by Jones, though obscured by the latter’s imprecise terminology. This risks excessively limiting Jones’s view of ethics, if it is in fact more capacious than that of Smith. My suspicion is that in cases of this sort, we are almost all likely to favor interpretations consistent with our own moral preferences.

So, it appears that, in practice, normative ethical preferences are likely to color our selection of what counts as ethics. Thus, such preferences are likely to affect our descriptive and explanatory projects. However, this is not a reason to reject the distinction. The distinction is deeply important. The likely interference of normative inclinations in descriptive programs actually gives us greater reason to try to hold the distinction in mind, making it as clear and explicit as possible at each point in our analysis, in order to limit confusion between the two as much as possible.

## **Some Distinctions That Complicate Both Projects**

There are, of course, other distinctions that are important to keep in mind when treating ethics. One of these is the difference between ethical deliberation and character development. Some ethical theories focus on how one should adjudicate ethical demands in particular cases. For example, how does a physician decide which of two people should receive the one remaining dose of a medication that both need? In connection with cases of this sort, it is usual to distinguish two broad types of moral decision as well – deontological or intrinsic and consequentialist. A deontological evaluation would begin from the premise that the ethically right decision has to do with the nature of the act itself. For example, a deontologist might argue that murder of innocent people is wrong, even if the murder would increase the well-being of many other people, prolonging their lives and increasing their happiness. In contrast, a consequentialist would focus on the likely results of each alternative. (There are also intermediate positions. For example, if I understand her correctly, Martha Nussbaum maintains that there is a core of intrinsic ethical rights, but the majority of ethical issues need to be evaluated in relation to the consequences of actions [see 2016, 173–174].)

But, again, our ethical concerns are not solely a matter of specific deliberative choice. Often, our ethical responses to situations are automatic, a matter of responding spontaneously and unreflectively. Robert Sapolsky refers to cases of people risking their life to save someone in an emergency (e.g., a child who has fallen on the subway tracks). When interviewed later, these rescuers regularly say they did not make a self-conscious decision to help but just acted spontaneously (see Sapolsky 2017, 566). This sort of behavior presumably results in part from innate propensities (e.g., spontaneous empathy). But it also results in part from more individual dispositions that derive from upbringing; in Sapolsky's words, an "act of implicit automaticity" is "the product of a childhood in which doing the right thing was ingrained as an automatic, moral imperative" (2017, 593). We often refer to an individual's interconnected innate and acquired dispositions as "character." A good deal of what we think of as ethics is a matter of the development of character. For example, among other matters, Foucault sets out to examine the history of responses to questions such as "What are the means by which we can change ourselves in order to become ethical subjects?" (1997, 265). In doing this, he is not addressing ethical deliberation but character. We see this also in his stress on care of the self (though that may be more prudential than ethical).

In practical terms, focusing on character makes sense. As Sapolsky points out, "It's great if your frontal cortex lets you avoid temptation,

allowing you to do the harder, better thing. But it's usually more effective if doing that better thing has become so automatic that it isn't hard" (2017, 671). However, I will be focusing on deliberative ethics. This is principally because, in ethical theorization, deliberation is necessarily more foundational than character. Of course, one's character develops in tandem with one's ability to engage in moral deliberation, and it develops prior to theoretical reflection regarding ethics. But that theoretically informed, critical reflection is necessary to determine just what sort of character development is ethical or unethical to begin with.<sup>2</sup> Thus, we need to engage in ethical deliberation, determining what sorts of action are good or bad, before we can engage in a program of training and educating children so that they will be morally good, which is to say, so that they will spontaneously act morally in particular cases. What counts as moral in those cases is defined by, precisely, ethical evaluation.

On the other hand, it is important to recognize that character development is often central to people's conception of ethics. This may be one reason for some divergences in what people consider to be ethical. In certain cases, the difference may not be precisely a matter of the acts themselves. It may be a matter of whether we trust the broader ethical orientation of the people involved. For example, Haidt tells a fictional story of a man who masturbates with a chicken carcass. He finds that at least some groups of US academics tend to say that the act is disgusting but not immoral, whereas (ordinary or "normal") people tend to say it is immoral (Haidt 2012, 111–112). Roughly the same point apparently applies to Haidt's other anecdotes. For example, one concerns a brother and a sister who have consensual sex once. The narrative fiction putatively covers all possible, practical contingencies, stipulating that they took foolproof contraceptive precautions, left no possibility of discovery, and so on. There are many possible reasons for the divergence between academics and nonacademic (normal) people in their responses to these anecdotes. Perhaps academics are more comfortable with the convention of stipulating that particular aspects of the story are certain, even though they would not be certain in real life.<sup>3</sup> But another possible reason is that US academics are confining their judgment to the specific, deliberative

<sup>2</sup> There are other, related possibilities beyond character, such as pride in "moral identity," "a source of moral motivation linking moral reasoning (our judgments about whether certain actions are right or wrong) to behavior" (Hardy and Carlo 2011, 495). These too rely on prior, deliberative judgments, as the preceding quotation indicates.

<sup>3</sup> This is a fatal problem with such famous "trolley" problems as the one in which test subjects are asked whether it would be moral to push a fat man onto the trolley tracks if you knew his body would stop the trolley and thereby save the lives of five people farther down the tracks. In real situations, you could never know such a thing.

choice, whereas nonacademics are implicitly taking into account whether they would trust the moral character of someone who engaged in such an act (e.g., whether they would implicitly assume that these incestuous siblings would more readily violate other, uncontroversial social norms as well). It may be that the responses of the two groups would be more similar if the research focused on our implicit moral trust of the siblings. (I say “implicit,” as I imagine many of us would recoil from people who did this, even if we were to judge deliberatively that it does not in fact reflect negatively on their moral character.)

I suspect that most people would take the most important distinction in types of ethical response and reflection to be that between foundationalism and relativism. The former is the view that norms have an identifiable, objective basis, whereas the latter is the view that they are merely a matter of convention, stipulation, or preference. Within this general division, two further subcategories may appear particularly significant. Among foundationalists, we may distinguish theological from secular approaches. Theological foundationalists view moral norms as defined by God or some other transcendent, religious principle. Commonly, the issue of moral deliberation in this context is a matter of ascertaining divine will (e.g., through the interpretation of divinely revealed scripture).<sup>4</sup> Secular foundationalists appeal to some nonreligious absolute, such as the laws of reason; for example, in Kant’s view, it is immoral to engage in actions that (roughly speaking) would prove contradictory (1981, 32) if their norm were willed as a general law. The most common form of relativism sees moral norms as varying by society or culture. Each of these approaches has some fairly obvious difficulties, which we can leave aside for the present.

For purposes of the current chapter, the key point is that these different methods of ascertaining what is ethical are likely to be consequential for describing ethical deliberation. In connection with this, I should comment on my own normative view here because it bears on the connection between descriptive and normative ethics, which is clearly important for a book that sets out to discuss both. I myself am certainly not a foundationalist. But I also bridle at the label “relativist,” principally as I do not view the ethical customs of a given society as having presumptive validity even within that society. For example, like almost everyone I know, I believe that condemnation of homosexuality is wrong; I do not think that such condemnation is wrong in the United States but right in,

<sup>4</sup> In terms of the definition I will propose in what follows, the “allocentrism” of theological foundationalism is often limited to God or the equivalent (e.g., heaven [tiān, 天], in some ethical traditions of China), which then serves as the single, definitive source of all ethics.

say, Iran (see Walsh 2019). Rather, I see myself as holding a third view, different from foundationalism and relativism, a view that I believe is actually fairly widespread, and even defended in different versions, but that is rarely identified and labeled as such.

Specifically, I hold the same view on ethical norms as I do on aesthetic norms (as discussed in Hogan 2016). I do not believe that such norms have any objective existence making them right “in themselves.” However, I do not believe that the crucial question is whether or not we can establish a foundation for norms. Rather, the key issue is whether we can make rational arguments in favor of one or another ethical position. I believe that we can do that. Specifically, we can do it through dialogue in which we seek to isolate common presuppositions and then reason over the implications of those shared presuppositions. The method for doing this is actually widespread in ethical theory. It involves considering specific cases, finding out whether we agree on those specific cases, seeking to draw out general principles, testing agreement on further cases consistent with those general principles, and so on.

For example, suppose Jones and Smith disagree about abortion. Jones asserts that the fetus is a human life and that the life or death of one person always outweighs lesser interests of others. Smith responds, “But I don’t think you believe that. For example, suppose I needed a kidney to live. You have two kidneys. It would certainly be good of you to give one to me. But do you really think that the state would be justified in mandating that you have to give a kidney to me, since – by your principle – my continued existence would trump any interests you have in keeping your kidney, avoiding surgery, and so on?” Jones may in turn respond, “You are right about the kidney case. But that is because virtually anyone could provide a kidney. The state cannot single me out to make that sacrifice. In contrast, the mother is uniquely able to support the life of the fetus. (By the way, I grant that it should be legal for the mother to have the fetus removed and raised in an artificial uterus, if that were to be developed.)” “But it doesn’t really help your case to say that the mother is unique. Suppose I have a vastly rare blood type and need a blood transfusion, and you are the only known person with that blood type. I suspect you still would not believe that the state could reasonably make it illegal for you to refuse to have your blood drawn.” (A dialogue of this sort might continue until the parties come to an agreement or discover a fundamental difference in their ethical beliefs.)

One difference between the way I approach these issues and the way many others do is that I take there to be a set of cross-culturally available moral attitudes that are central to our moral thought and action. I am not alone in this view about cross-cultural moral attitudes. It is shared by, among others, some evolutionary ethicists and a number of writers in



cognitive and affective science. For example, Jonathan Haidt (2012) suggests something along these lines in his “moral bases.” In my account, these cross-cultural attitudes – derived ultimately from emotion systems and bound up with narrative prototypes – provide partial, implicit bases for rational dialogue on moral issues. The general operation of empathy, simulation, and other cognitive and affective processes contributes to the possibility of such dialogue as well. Note that we do not all share equal commitment to all cross-cultural moral attitudes. Thus, there are likely to be some issues – sometimes, many issues – on which two people are unable to find a common ground. For example, I will argue that the bonding of romantic love and the associated narrative prototype favor ethical norms supporting free individual choice and opposing social identity categorization. In contrast, heroic narratives – with their celebration of group pride – tend to purvey norms of in-group loyalty and obedience to social authority. People who particularly stress the former are likely to find it difficult to engage in productive dialogue with people who particularly stress the latter, and vice versa.

As to cultures, they clearly disagree to some extent on the precise norms – thus, the emotion systems and narrative prototypes – they emphasize. On the other hand, there is always some distance between the norms affirmed by a society or any individuals in that society and the norms they actually follow in their daily lives. As Justin Barrett (1999) has argued, there is often verbal conformity to “theological correctness,” which is belied by actual practice. The point clearly bears on “moral correctness” as well. In consequence, it is not necessarily true that a difference between two societies’ claims about ethics reflects differences in their real ethical attitudes and behaviors. Moreover, even when two cultures genuinely diverge in which moral attitudes they emphasize, the same sorts of ethical attitudes are likely to appear in both societies, even if they do not do so in the same proportions. In short, these recurring moral attitudes enable rational moral dialogue between some people, even in different cultures, while also inhibiting such dialogue between other people, even in the same society.

### **So, What Is Ethics, Then?**

Given the distinctions just outlined – as well as others that we will be taking up later – the first proposal I would make is that we should not consider ethics to be a concept with strict necessary and sufficient conditions. We should, rather, understand it as a prototype concept. There are, in other words, concerns that are more centrally ethical than others. In defining ethics, we want to capture the most central cases. Where we leave off on more peripheral cases seems less important, though there will

certainly be differences – personal, cultural, and other – in just how much one is inclined to include.

How, then, do we go about identifying core cases of this or any other prototype concept? It is often valuable to consider what sorts of primary contrast there are with a given idea. In the case of ethical decision, there is one clear and consistent contrast: acting out of self-interest. Moral behavior is difficult precisely because it so frequently conflicts with self-interest of some sort. As Schramme puts it, “Moral demands often oppose our own self-concerned desires,” such that “morality and egoism are opposing forces acting on the agent” (2017, 204). Roughley and Schramme write similarly that “moral action . . . is at least primarily a kind of action for other-regarding reasons” (2018, 6). Indeed, the entire point of having a moral code is to establish a norm that can overrule self-interested decisions and behaviors. If ethics converged with self-interest, we would not need ethics. Ethics would simply be what we did when acting selfishly, which we would do automatically. Or, rather, this would be the case if it converged with what we might call “egocentric self-interest.”

I will use the term “self-interest” to refer to any motivation to pursue one’s own well-being. I take it that all our behavior is in fact self-interested in this broad sense – though, of course, the “well-being” at issue may be short term, misguided in various ways, or even simply wrong about one’s preferences. (To take a trivial example, for years I thought I liked jalfrezi and would order it in restaurants, only to recall upon its arrival that jalfrezi is something completely different from what I remembered; nonetheless, on each of these occasions, I was pursuing self-interest with my order.) I would call self-interest “egocentric” insofar as it is not contingent on the well-being of others. I may act out of largely – perhaps even entirely – “allocentric” self-interest if I am distressed by someone else’s suffering and am myself relieved when that person’s suffering is alleviated. Many authors refer to this as “altruism.” That seems to me reasonable only if by “altruism” they mean behavior that provides one with no benefits beyond relief at the ending of another person’s suffering (or, in other cases, satisfaction with his or her joy).<sup>5</sup> However, if one is using “altruism” in the strict, motivational sense – where the agent receives no benefit whatsoever – then this

<sup>5</sup> For example, Rottschaefer explains that he uses “altruism” interchangeably with “prosocial.” He goes on to define “prosocial intentions and actions” as “those that have the benefit of another as their object” (1998, 86). Rottschaefer also explains that “a survey of studies shows that in adults, empathy . . . is positively associated with prosocial behaviors” (1998, 93). For a thoughtful analysis of altruism, though one that bypasses this aspect of the issue, see Kitcher (2006).

characterization seems to me clearly wrong.<sup>6</sup> If I help a suffering person, I do so because I care that he or she is suffering. Thus, I am gratified when his or her pain is allayed. I am not indifferent. It is only if I am indifferent to his or her misery that I could be said not to benefit from the action in any way. (I leave aside cases where I help the suffering person because I want to be rewarded by God or the like, since these are clearly not altruistic.) When we speak of a conflict or an opposition between ethics and self-interest, we are not generally speaking of allocentric self-interest, but of egocentric self-interest.

With this in mind, I give the following preliminary and partial characterization of ethics, or rather of ethical evaluation and related behaviors: A value judgment and associated action are more prototypically ethical to the extent that they oppose egocentric self-interest. If I am correct that all action is in fact self-interested, then an ethical judgment can have motivational force, and therefore can lead to action, only insofar as it is allocentrically self-interested. Thus, judgment and action are more prototypically ethical insofar as one's (self-interested) well-being is contingent on the well-being of others (thus, is allocentric).<sup>7</sup> We have the highest degree of ethical prototypicality when egocentric well-being is fully contradicted by the ethical alternative, as when it leads one to sacrifice one's life.

Though my precise formulation is necessarily specific to my analysis and larger project, this definition is broadly in keeping with the observations of many other writers. For example, Haidt's definition reads as follows: "Moral systems are interlocking sets of values, virtues, norms, practices, identities, institutions, technologies, and evolved psychological mechanisms that work together to suppress or regulate self-interest and make cooperative societies possible" (2012, 314). Durkheim similarly writes, "What is moral is everything . . . that forces man to . . . regulate his actions by something other than . . . his own egoism" (quoted in Haidt 2012, 314). Howard Brinton makes the related claim that "in all religions self-centeredness is the chief sin from which all other sins are derived" (1973, 71). The Golden Rule – the paradigmatic ethical principle in Christianity – adjures reversing egocentrism into allocentrism: "Therefore all things whatsoever ye would that men should do to you,

<sup>6</sup> Another use of "altruism" makes sense here but seems irrelevant to our concerns with ethics, as it is nonmotivational. Specifically, De Waal points out that "'altruism' is defined in biology as behavior costly to the performer and beneficial to the recipient regardless of intentions or motives" (2006, 178).

<sup>7</sup> I should note that this relation may be inverse rather than parallel. Indeed, that is typically the case regarding out-group members, as we will see. However, due to my own ethical preferences, I will most often stress empathic allocentrism.

do ye even so to them" (Matthew 7: 12, King James version).<sup>8</sup> The Confucian stress on shù (恕) or empathy involves the same idea. In a famous passage of the *Analects*, one of Kǒngzǐ's followers asks, "Is there one word which may serve as a rule of practice for all one's life?" Kǒngzǐ replies with the word "shù" (恕), explaining, "What you do not want done to yourself, do not do to others" (15.24, altered from Legge's translation; see also 4.15). The same idea is repeated in "The Great Learning" (*Dà Xué*, 大學, 13).<sup>9</sup> The general idea is consistent with Kant's assertion that one should "act in such a way that you treat humanity, whether in your own person or in the person of another, always at the same time as an end and never simply as a means" (1981, 36; treating someone simply as a means would presumably be a matter of using them to satisfy one's egocentric self-interests). It is also consistent with utilitarian consequentialism, which displaces egocentric self-interest with a calculus of general happiness (in which one's own well-being does count, but only as much as that of any other individual).

Note that this definition or characterization does not bias ethics toward any particular outcome. It is consistent with supporting the free choice of lovers, including choice that violates racial or other forms of segregation. But it is also consistent with supporting the self-sacrifice of soldiers, or the killing of unknown enemies, or the decision of lovers to separate in conformity with dominant racist prejudice. Personally, I favor universalism in ethics and find in-group preference highly immoral. As Patricia Greenspan writes, "The universal element of morality lies . . . in its application to everyone, regardless of culture or some narrower set of personal affiliations" (2010, 244). But clearly many other people disagree. As Marilyn Brewer explains, "Research indicates that laypersons . . . view ingroup favoritism as normative . . . and that they reveal implicit preference for an ingroup member who discriminated in favor of fellow ingroup members over one who behaved in a fair, egalitarian way"

<sup>8</sup> See [www.kingjamesbibleonline.org](http://www.kingjamesbibleonline.org) (accessed September 2, 2020).

<sup>9</sup> As I discuss in connection with Goldie's views, we usually need to alter our simulation of others' experiences and desires, taking account of differences between us. However, a risk of such "correction" is that it will rely on stereotypes or other false beliefs, thereby increasing rather than decreasing error. As Nickerson, Butler, and Carlin point out, "under certain conditions, people who project their own opinions to others may predict other people's opinions more accurately than those who do not" (2009). On the other hand, it also remains true that "models of others' knowledge that are based solely on one's own knowledge are likely to need some degree of adjustment on the basis of individuating information if they are to avoid being inaccurate in many particulars" (2009, 48).

(2017, 102).<sup>10</sup> Both the in-group favoritism and the universalism count as descriptively ethical by this definition, as neither is simply egocentric.

A further benefit of this formulation is that it both distinguishes and relates ethics and prudence, which is, I believe, appropriate. There is something almost moral about being prudent. Doing what is best for me in the future is clearly not the same sort of thing as doing what is best for someone else. However, “me in the future” is different from me now. To choose the prudent option, I must in a sense choose the well-being of someone who is, so to speak, “less me” than I am right now. Indeed, when I imagine myself in the future, I have to take up empathic processes to try to determine what I will feel in the future, much as I have to do with others (see Soutschek and colleagues 2016; Yong 2016). Of course, the difference is that if something harms me in the future, I will directly feel the harm; it will not merely be empathic at that future time. Thus, prudent decisions remain egocentric, but they are egocentric in a somewhat attenuated sense.

A possible objection to this characterization is that it to some extent excludes forms of behavior that many people would consider ethical, specifically conformity to precepts due to fear of punishment or desire for reward. It makes such behavior prudential, rather than ethical as such. But that seems right to me. At the very least, it seems relatively clear that it is more ethically praiseworthy to help one’s neighbor out of fellow feeling than out of fear of postmortem torment. In the latter case, the judgment and action involve a form of egocentric self-interest, even if one that is (again) attenuated.

### **Specifying a Fundamental Ethical Orientation**

Up to this point, I have only set out a broad characterization of ethics and pointed to a few general parameters (e.g., deontological versus consequentialist) that, when set differently, may lead to different ethical orientations (which in turn will need to be further specified into more particular ethical ideas and attitudes, which I refer to as “prototypes”). Rephrasing the preceding conclusions, I might say that ethical norms are principles that one acts on for nonegocentric reasons, thus reasons other than direct benefit to oneself; more precisely, the actions benefit oneself only as a function of benefiting someone else. This definition is intended to mark a broad area constituting ethical thought in general. Thus, it encompasses a range of contradictory systems, including the entire

<sup>10</sup> This is consistent with de Waal’s plausible suggestion that “morality likely evolved as a within-group phenomenon” (2006, 53).

spectrum of political attitudes. For example, it includes the Fascist morality of absolute devotion to the state, articulated by Mussolini and Gentile as a “conception of life” that is “ethical” (2012, 6; “etica” [1935, 2]) “in a world sustained by moral forces” (2012, 6; “forze morali” [1935, 2]); more exactly, “anti-individualistic, the Fascist conception of life stresses the importance of the State and accepts the individual only in so far as his interests coincide with those of the State” (2012, 7).

Clearly, then, this broad definition of ethics is too broad to lead to any particular ethical decisions, either spontaneously and intuitively (when we feel a particular action is moral or immoral) or through self-conscious deliberation (as in ethical theory). To get any sort of concrete ethical response, we need to specify the values of some key ethical concerns that vary across moral agents. I have referred to this as establishing a fundamental ethical orientation. It is probably the case that, for each of us, these variables are initially set through developmental experiences in childhood. Those developmental settings are probably enduring and are likely to be crucial to our spontaneous ethical response. However, we may also come to self-conscious conclusions about which settings are valid, as when we engage in ethical philosophy. The unselfconscious and self-conscious settings may be consistent or inconsistent with one another. In addition, both types of setting may vary with context. For example, we might spontaneously evaluate acts in terms of intrinsic morality but adopt an ethical theory that is consequentialist; in addition, our spontaneous evaluation might change in contexts where the harmful results of an action are unusually salient. Thus, there is considerable complexity in the actual operation of these principles and parameters. In this section, I will consider some of the more consequential variables that differentiate kinds of ethical orientation, conceptualizing these in terms of the parametric specification of ethical principles. However, to make clear the basic operation of such parameter setting, I will largely leave aside the complexities just mentioned.

I take three ethical parameters to be particularly consequential for defining types of ethical orientation. Again, ethical issues enter when there is a conflict between our individual well-being and that of others. However, an ethical preference for others’ well-being does not tell us how to resolve several difficulties. The first is that other people’s conditions of well-being will come into conflict with one another. Our ethical orientation must be able to determine which people should be the object of our ethical concern. Put differently, we need to set a parameter defining who receives our ethical consideration. I refer to this as the *scope parameter*. For example, nationalists confine ethical preference to the nation (and, within that group, Fascists confine it to the state). (There may be scope subvariables that hierarchize nonnational groups as well.) In contrast, ethical

universalists set the scope parameter at humanity (or, in some cases, all sentient beings).

The scope parameter alone is insufficient to specify ethical response. Another concern is just what constitutes the “well-being” at issue. We tend to evaluate well-being relatively, not absolutely. This is indicated by the association between well-being and “downward comparison” (see Wills 1981), a feeling of (in this case) in-group dominance in status, wealth, or some other property or condition. It is also indicated by the preference of individuals for establishing in-group superiority over out-groups, even when the absolute level of concrete benefit to the in-group would be higher if the groups were equal (see Duckitt 1992, 68–69). Moreover, except in the case of universalism, the scope parameter identifies preferred and dispreferred groups – for example, the home nation and other nations, respectively, for a nationalistic scope. Once we have a particular set of targets for ethical preference, along with a set of targets defined as falling outside this preferred group, we may be supportive of the well-being of the preferred group or critical of the well-being of the dispreferred group. Specifically, our particular ethical judgments, and the encompassing systems of interlocking principles that they form, tend to stress either pleasure or pain, commonly seeking to confine one or the other to putatively “merited” cases. I refer to this as the *valence parameter*.

As just indicated, ethical preference is defined by reference to group membership – for example, national or racial group membership in the case of fascism. But that is not the same as merit. Whatever scope governs one’s ethical orientation, there are some cases of unmerited pleasure and some cases of unmerited pain for members of all groups (though one is likely to be concerned with only some of these cases, depending on how one’s valence parameter is set). One needs to determine precisely what constitutes merit. Merit is a function of conforming to a moral precept. Unmerited pain occurs when the target did conform to the relevant moral precept or when his or her violation is considered to be less severe than the punishment. Unmerited pleasure occurs when the target did not conform to the relevant moral precept (e.g., did not get married before having sexual relations) or did not engage in the more morally elevated, benevolent activities (e.g., religious practice or charitable giving) that would ethically justify reward. Thus, one needs a means of isolating what constitutes merit in these cases. I refer to this as the *method parameter*. As already suggested, methods are generally seen as falling into two broad groups: consequentialist (e.g., utilitarian), which evaluates ethical alternatives on the basis of (predicted) outcomes, and intrinsic or deontological (e.g., Kantian), which grades actions on (inferred) intrinsic worth.

In short, I will be stressing three parameters of allocentric judgment and behavior: first, the scope of the targets; second, the valence of the

targets' experience; and third, the method for adjudicating among evaluative and behavioral alternatives. There may well be other important parameters. However, it seems clear that these three are highly consequential, even if they turn out to be three items on a longer list. I will consider them in turn.

Before going on to these, however, I should say something about the next level of ethical categorization – prototypes. Though we will not turn to prototypes until Chapter 3, it is important to briefly consider their place in ethical response. As work by Eleanor Rosch and others has demonstrated (see Rosch 2011 and references therein), human beings tend to categorize things and events, concrete and abstract, not by applying strict, rule-based, necessary, and sufficient conditions. Rather, in ordinary life, we tend to categorize by evaluating a target's similarity to some (roughly) average case of a category. Thus, scientists judge something to be a bird by reference to strict criteria, but in ordinary life, we basically judge things to be birds based on their similarity to robins. Likewise, when doing ethical theory, we may evaluate the morality of an act by reference to strict criteria, beginning from the self-conscious version of our ethical orientation. However, when reacting ethically to situations in real life, we are likely to rely on their assimilation to standard cases. In ethical response, the situations and responses at issue are not static things. They are, rather, dynamic sequences of events and (human) actions. The trajectories of such sequences are given in stories, and thus the prototypes that guide our ethical understanding, evaluation, and response in such cases are, most importantly, narrative prototypes.

This is not to say that fundamental orientations do not play a key role in our ethical response. Rather, they play precisely an orienting role, guiding our attention to certain aspects of a situation, tilting our emotions one way or the other, affecting just what narrative prototypes are activated (thus likely to further organize and specify our response), and so on. Indeed, the use of prototypes in ordinary categorization has been overstated in the psychological literature. Prototypes are crucial, but they interact with necessary and sufficient conditions. For example, in ordinary activity, being organic is a necessary condition for being a bird, as opposed to a bird toy or a painting of a bird. (We may speak loosely of these as being "birds," but obviously, if asked, we would all say that a painting of a bird is a painting, not an actual bird.) This sort of interaction between strict conditions and prototypes occurs in ethics as well.

### *The Scope Parameter*

Again, the scope parameter sets the extent to which a non-self preference extends ethically. By the definition given above, simply in order to count



as ethical, a motivation must be nonegocentric. As a general principle, we could say that acts become more paradigmatically ethical to the extent that they generalize beyond the self up to some ideal level, at which point ethicality begins to decline. A minimal extension of ethical obligation may be to one's immediate family. A broader extension might encompass one's nation, ethnic or racial category, or religion. A still broader category would include all humans. All sentient beings would go beyond even that. In the *Xiàojīng* (*Classic of Filial Piety*), this parameter is set to the family, or more precisely one's parents (see section 9 in de Bary and Bloom 1999, 328; see also the *Analects* 13.18). In Mussolini and Gentile's *The Doctrine of Fascism* (2012), the nation – more exactly, the state – is the definitive level for ethical decision. This means that in cases of conflict, one should base one's ethical decisions on the condition of the nation, not of one's individual family or of humanity as a whole. In contrast, setting the parameter to humanity as a whole would subordinate concerns of the nation.

It is probably obvious that in setting the scope parameter, I favor a very broad scope as morally definitive.<sup>11</sup> In the context of the present book, I will stress humanity as the morally definitive level of generalization. My argument in favor of this is based primarily on the fundamental characterization of ethics. Again, to count as ethical, a judgment and an action must derive their motivational force from the condition of some other person (or other appropriate target). Since the definition itself takes the ethical norm to be opposed to egocentric action, it seems clear that the most paradigmatic target of our ethical concern would be the one that is most fully removed from the egocentric interests of the ethical agent. Families involve more egocentric concerns than nations, which involve more egocentric concerns than humanity as a whole. It would thus seem that the logic of ethical nonegocentrism pushes toward more and more inclusive groups. This logic actually leads beyond species and therefore arguably extends to sentient beings as the proper level for setting the scope parameter. However, there are complications here (e.g., bearing on the similarity between the empathizer's experiences and those of the target, which presumably vary more radically

<sup>11</sup> Though this part of the book is devoted primarily to descriptive ethics, I believe it is valuable to inform the reader of my normative views on the parameters and to give a basic idea of my reasons for these views. I say this for two reasons. First, I serve as an illustration of the description, which is of course relevant to the descriptive part. Second, I will not be focusing on fundamental orientation – or even prototypes – in the normative part, so it is more fitting to sketch my normative views on these topics here.

across species than within species). Consequently, I will largely leave aside the issue of species preference.

Before going on, I should briefly respond to an objection that sometimes comes up in this context (see, e.g., Churchland 2019, 160). If I say that the scope of ethics is this broad, does that not commit me to the view that one should let one's child starve to death if it means providing food to two people in Africa? No, in fact, it does not mean this. First, the scope parameter does not mean that there are no social arrangements that bear on ethics. Social organization entails that I have some responsibilities and not others. For example, it may be that, on a certain day, I would do more good by visiting some lonely patients in the hospital rather than teaching my class on empathy, ethics, and narrative. It does not follow that I should abandon my class simply for that reason. Our ethical obligations are affected by social arrangements that allocate obligations, even in cases where that allocation is arbitrary. Clearly, given the organization of society, one has obligations to one's own children that one does not have to others' children. Moreover, socially defined obligations of some sort would bear on the ideal society. If we understand ethical action as the type of action everyone would undertake in an ethically ideal society, that gives further moral weight to fulfilling one's socially defined obligations (though of course in both real and ideal societies, there would be exceptions).

But what about foregoing nonobligatory benefits for one's children in order to benefit unknown people in another continent? In fact, this seems to me a plausible ethical choice, assuming one has already sacrificed one's own superfluous pleasures. (It hardly seems ethical to deny one's children, say, a vacation if one takes vacations oneself.) Of course, socially allocated obligations are not entirely clear, nor are they all equally obligatory. Providing food and shelter are strongly obligatory; providing a college education is intermediate; providing fashionable, prestigious clothing has arguably no moral weight. Note that moral weight here needs to incorporate considerations of the child's sense of being cherished by his or her parents; it is not merely a matter of some sort of objective need. It does seem to me that giving money to keep someone from starving is indeed ethically preferable to spending money on a vacation. It is even arguable that it is ethically preferable to paying for one's child's college education (again, given the qualification that one has deprived oneself as far as ethically reasonable already), though I personally see educational support as having the stronger argument, given current social circumstances. However, given common social practices, we are likely to count such acts of charity to strangers as morally benevolent rather than as morally obligatory. In other words, even in cases where favoring others over one's own family seems ethically preferable, we are

unlikely to judge the kin preference to be sinful; we are more likely to see the alternative as saintly. In any event, it is a preference that we can most often discuss rationally. Moreover, in my view, we generally overvalue kinship relations in this regard and should typically count more actions benefiting non-kin as obligatory rather than as merely benevolent.

These points may be captured by a slight complication of the theoretical apparatus I have been invoking. In addition to the larger parameters, which may be thought of as defining ethical systems, there are also more restricted variables that affect the particular ways in which the systems operate. So the scope parameter fixes the paradigmatic reference level for ethical adjudication generally. However, there are different ways in which this general adjudication could be worked out. For example, societies are typically set up in such a way as to assign different sorts of obligation to different people. In traditional Hindu ethical theory (see O'Flaherty 1978; Parekh 1989), for example, there is the range of "swadharmas," the particular dharmas or duties that apply to us individually, as well as the principles of *sādhāraṇadharmā* or universal dharma, which apply to everyone. The swadharmas include one's particular obligations within the family. In almost all societies, there is some sort of special ethical obligation one has toward one's family. There is also often a special ethical obligation that derives from one's profession (e.g., doctors and lawyers have ethical obligations that bear on their professions and that do not bear on the profession of, say, a literature professor). Just how we set these subparametric variables will determine how we believe someone should act in particular cases of ethical conflict. For example, one's special, socially defined responsibility to one's family may lead one to having certain obligations to them over others even when one's fundamental ethical relation is to humanity as a whole. The scope parameter, then, does not reduce all ethical preference to a single level. However, it provides the broader context in which subparametric variables operate. For example, one's primary moral obligation may be to humanity, but in particular circumstances (e.g., saving people from some natural disaster), one's first obligation may be toward the segment of humanity that constitutes one's family, or some other group for which one has some socially defined responsibility.

### *The Valence Parameter*

Our response to other people's emotions is focused principally on valence. Thus, it is primarily a matter of responding to their hedonic or aversive experiences. The point appears to hold both in general and for ethical response in particular. As indicated by the basic characterization of ethics as allocentric, what makes a given act ethical or unethical is

necessarily its relation to other people or to other sentient beings. For example, stealing a necklace from a storefront display is immoral not because the vandals have been mean to windows or jewelry, but because their actions on windows and other items have consequences for people. Those consequences are presumably painful in this case. If we are considering a case of generosity, rather than vandalism, then the consequences would presumably be pleasurable for some other people. In principle, we could all have ethical concerns for both the pleasures and the pains of other people. Indeed, this strikes me as preeminently sensible. However, it seems that we tend to focus our ethical interests on one or the other, though the point does not appear to have been widely recognized. Specifically, our ethical response to other people tends to focus on the degree to which those people may be considered to merit some quality of experience. To put the point simply, some of us appear to be most concerned to restrict unmerited pleasure; others appear to be most concerned to restrict unmerited pain. In some cases, these orientations produce the same or nearly the same result. For example, moral judges of both persuasions would condemn the theft of jewelry. The anti-pleasure group would grudge the thief the unmerited enjoyment of the necklace. The anti-pain group would pity the shopkeeper for his or her undeserved loss of the investment. In other cases, however – for example, in moral assessments of consensual sexual indulgence – the orientations will diverge.

I imagine it will not come as a surprise to any reader that I am in the second group. My argument for this is twofold, concerning both the ethics of hedonic restriction (against unearned pleasure) and the ethics of relief (against undeserved pain), as we might call them. The ethics of relief is somewhat simpler, so I will begin with it. The desire to lessen or end someone else's pain is generally an empathic response. In some cases, it may be an egocentric response to one's own personal distress (e.g., when one cannot simply leave the situation). But most often it is empathic. As such, it is highly prototypical for ethical allocentrism.

The ethics of hedonic restriction, on the other hand, seems to be routinely connected with egocentric motives, most obviously envy (for unmerited wealth or social status), jealousy (for unmerited love), or disgust (for unmerited physical enjoyments, especially sexual, but also gustatory). Indeed, in the case of moral condemnations of pleasure, it seems to me often very difficult to separate a nonegocentric component from a person's own feelings of deprivation and resentment. For example, I am often unhappy when literary "theorists" receive great adulation for analyses that I believe are muddled and lead to badly mistaken conclusions. This is in part an ethical attitude, and I do not wish to deny that there are ethical issues bearing on elevation in social prestige, which typically

goes along with increased wealth and power as well. However, it does not take a psychoanalyst to figure out that my animus against such celebrated figures is inseparable from my feeling that my own (lucid and insightful?) analyses are (woefully?) underappreciated. Moreover, insofar as these concerns can reasonably be judged to have genuine ethical force, it seems that such force is largely a matter of harm anyway – the harm done to theorists who do not receive the advancement needed to continue their work, the harm to other academics who are deprived of the possible insights, and so on.

But, of course, as with the other parameters, this is not even close to a definitive argument. It is merely an indication of why we might reasonably prefer an ethics that seeks to limit harm over an ethics that focuses on limiting pleasure. I believe this is also consistent with the general asymmetry between pleasure (to which we readily habituate, so that the initial experience stops being pleasurable) and pain (to which it is much more difficult to habituate), as well as that between opportunities and threats.

### *The Method Parameter*

The most complicated of the three parameters, at least in my own reflections and attitudes, is that of method. Indeed, in this case, the necessity of subparametric variables becomes abundantly clear. The first parameter setting is one in which one defers to a moral authority or, alternatively, assumes the individual's ability to make moral judgments and infer moral precepts himself or herself – roughly, ethical heteronomy versus ethical autonomy. For example, someone with the former ethical orientation may take morality to be defined by the prescriptions of a sacred text. Of course, things do not stop there. Subparametric variables include new versions of heteronomy versus autonomy. For example, in interpreting the sacred text and in applying it in particular situations, one may defer to canonical commentaries or engage in interpretive effort oneself. I will largely ignore this parameter, as I am interested in autonomous moral responses. This is in part due to my own ethical preferences. But it is principally due to what I see as a general human propensity to assume moral autonomy in intuitive moral responses. We may assert adherence to an exogenous moral code, and indeed that moral code may make some difference, particularly in our self-conscious judgments. But it seems that our (tacitly autonomist) spontaneous judgments are primary and that we interpret our external moral laws in such a way as to make them fit our intuitions, rather than abandoning the latter when they do not fit the external system. Notorious cases include the convoluted justifications of aggression by Christians (i.e., followers of Jesus, who advised turning the

other cheek [see Matthew 5:39], not, say, dropping cluster bombs). (See Haidt 2012, chap. 2 for empirical research supporting this view.)

Setting aside autonomy/heteronomy, then, the main issue in method (as I am using the term here) is the degree to which one conceives of acts as intrinsically ethical or unethical and the degree to which one conceives of ethicality as a function, not of the acts themselves, but of the outcome of the acts. To complicate matters, not everything falls neatly into one of these categories. There are also mixed cases, such as Nussbaum's advocacy of a core of rights (which define intrinsic moral values) integrated with a wide-ranging consequentialism (2016, 173–174). A particularly important mixed case is the ethics of care. In that case, the parameter is set to "intrinsic" for the core value of care about particular persons in a given situation, but the response following the cultivation of such care is consequentialist (at least in the cases I am familiar with).<sup>12</sup> Alternatively, one could say that almost all approaches are mixed, assuming some intrinsic value or values and supplementing those with attention to consequences. For example, utilitarians mandate utility maximization as their intrinsic good, leaving all the rest to consequences. Indeed, even Kant would seem to rely on some sense of consequences. For example, how can one possibly attend to treating someone as an end or as a means without considering the consequences of one's acts? Moreover, Kantian universalization explicitly relies on hypothetical consequences.

One reason this is a difficult parameter to discuss is that, I suspect, people's preferences in this case are less clear and less stable than in the cases of scope and valence. In other words, I imagine that we are all inclined to accept some intrinsic values and some consequentialist concerns, but not always consistently. Nonetheless, for ease of exposition, I will treat these alternatives as if they were as clearly separated in practice as they are in concept.

There are many arguments for and against both approaches. For example, as to deontology, few people are willing to agree with Kant that one is morally obligated to tell the truth about the hiding place of an innocent man who is being sought by a murderer (see "On a Supposed

<sup>12</sup> Care ethicists are often opposed to being characterized in terms of these two categories. I believe that this is a mistake and results in explanations of care ethics sometimes having a paradoxical quality. For example, in his valuable treatment of care ethics and empathy, Hamington asserts that "care does not take a preemptive formulaic approach to right action." But, in the immediately following sentence, he also asserts that "the care giver must be attentive and responsive to the other"; the "must" here is clearly moral and defines a preemptive formula for right action. He goes on to explain that "the moral concern is for the other" (2017, 265), which clearly implies an intrinsic value.

Right to Lie" in *Grounding* 1981, 63–67).<sup>13</sup> Note that this view is not a unique result of Kant's specific deontology. Someone committed to divine determination of moral value might find textual support for the view that lying (or bearing false witness) is morally wrong without exception. It is arguably the sort of problem that is likely to arise in relation to a range of intrinsic ethical orientations.

On the other hand, cases like the murderer seeking a victim – where the (evil) consequences are apparently quite clear – are not common, and that poses problems for consequentialism. Most often, the ramifications of an act are too complex and diffuse to make sense of beforehand. For example, the Nazi regime in Germany was obviously morally repulsive. It was also engaged in invading other nations. Going to war against Nazi Germany seems to be the one case of war that even the strictest pacifists tend to agree was apparently necessary. In this case, the standard argument would be that at the time when the United States decided to enter the war, it seemed likely that this entry into the war was necessary on consequentialist grounds. Moreover, despite the incredible complexity of war, it appears retrospectively that such a consequentialist assessment was correct. However, we do not in fact know what would have happened if the United States had not entered the war but had undertaken some nonviolent (or at least less violent) alternative. For example, Robert Paxton (2005) notes that the Nazi "policy of expulsion" of Jews gave way to "a policy of total extermination," and that change needs to be accounted for. He goes on to explain that "a recent plausible theory locates" the change in December 1941, in Hitler's "reaction to the entry of the United States into the war and its transformation into a truly worldwide conflict. Hitler would thus be fulfilling the threat he made in a speech on January 30, 1939—that if the war became worldwide, the Jews were to blame and would pay (Hitler believed Jews controlled American policy)." I do not know enough to have a worthwhile opinion on this theory. The point is that something like this could turn out to be true, thereby potentially undermining even the most apparently uncontroversial consequentialist inferences.

Here, we might recall Gandhi's advocacy of nonviolence in resisting Nazism (see Kling 1991). Gandhi is often criticized or simply dismissed as naïve for this view. But, even on the face of it, one has to admit that the military response to Hitler, though successful in ending Nazism eventually, hardly prevented atrocities. Put crudely, the military approach succeeded only after millions upon millions had been killed.

<sup>13</sup> Varden has argued that Kant's claim has been seriously misunderstood. I adopt the common interpretation of Kant's view here because it poses an apparently greater challenge to the universalism I am advocating.



It is not necessarily obvious that a Gandhian approach would have led to more deaths. Of course, the idea is that the Nazis would never have stopped exterminating people and simply would have added the non-violent protesters to the (ever-lengthening) list of victims. But, as Kling points out, there was a famous case of nonviolent resistance in Berlin in 1943. Jewish men married to non-Jewish women were arrested and imprisoned for deportation. But 6,000 women went to Rosenstrasse, where the men were housed, and refused to leave. They did not have grenades and machine guns. They did not throw bombs. But the Nazis eventually backed down and released the men. (Margarethe von Trotta made a fine film of the incident in 2003.) This at least appears to suggest that the outcomes of a nonviolent response to Nazism may have had a reasonable chance of success, possibly even at much less cost than resulted from the military option (a point suggested by Paxton 2005 as well; see Sinclair 2017).

Moreover, in such political cases as the decision of a government to enter a war (as opposed to an individual's personal decision about a private matter), our understanding of consequences is rendered even more uncertain and distorted by the interests of the political leaders and conflicting groups involved. Consider, for example, a standard case of "humanitarian [military] intervention" – the NATO attack on Serbia. Extensive Serbian atrocities are widely cited as evidence that "humanitarian intervention" was needed in that case. ("Intervention" in this context means military action, not diplomacy or other nonviolent forms of response.) However, as Chomsky (2002) has shown, for the most part, the Serbian atrocities cited retrospectively in justification of NATO bombing did not precede that bombing. Rather, the atrocities in question were provoked by the NATO attacks, as NATO leaders actually predicted beforehand. (Unsurprisingly, all sides had committed atrocities earlier, but that hardly supported bombing Serbia specifically.)

The preceding examples suggest that there are problems with both deontological and consequentialist approaches. My preference is to set the parameter to "deontology." As a default position, I believe that the ethical status of the act should be decisive. Or, rather, I believe that acting according to the deontological default is always ethically "safe" and thus not blameworthy. However, there is a subparametric variable here that allows one to override the default for consequentialist reasons. The strength of a case for a consequentialist override is a function of three variables: (1) the degree to which the act in question is unethical; (2) the degree of certainty of the relevant consequences of that act; and (3) the degree to which these consequences would violate ethical desiderata. The consequences in the last case need to be understood broadly, encompassing not only the direct consequences for the people involved, but also



such concerns as the impact of having one's action stand as a model for other people (e.g., encouraging others to commit similar [unethical] acts in perhaps very different circumstances). The case for overriding the deontological default (e.g., telling the truth) becomes stronger to the degree to which the outcomes of all possible actions (e.g., lying and telling the truth) seem particularly clear; the outcomes of the default deviate greatly from ethical desiderata; and the intrinsically immoral act (e.g., lying to the murderer) deviates only minimally from the ethical desiderata.

For example, consider Kant's case in relation to these variables, and to my preferred settings of parameters, with intrinsic moral value defined by treating others as ends and not simply as means. Lying in this case would involve not treating the murderer as an end in himself on this occasion, which is unethical, but less unethical than causing someone's death (thereby ending the possibility of anyone ever treating him or her as an end). The likely consequences of both the truth and the lie are fairly clear. For example, it seems unlikely that such a lie would serve to encourage lying more generally. In contrast, the nearly certain outcome of telling the truth would be the murder of an innocent person.

Before going on, it is worth turning briefly to a type of case sometimes brought up in this context. A group of people is hiding from the Nazis. A baby in the group starts to cry. If he or she continues, it is virtually certain that the group will be discovered and then sent to a concentration camp where they will eventually die, after the excruciating experience of the camp itself. The baby will presumably be killed as soon as the group is discovered. For some writers, the ethical point of the anecdote is that there is an ethical obligation to suffocate the child to prevent the group's discovery. For example, this is the conclusion suggested by Paul Bloom's contention that the advocacy of empathy is ethically misguided because empathy is (among other things) innumerate (see 2016, 32–36).

The account suggested by Bloom's arguments, I believe, oversimplifies ethical issues such as this by treating them as if they are a matter of ethical criteria being or not being sensitive to numbers. By the preceding parameters and variables, this case should involve the following considerations. First, the act of killing the baby is an extreme violation of the principle of treating others as ends and not simply as means. Indeed, it is as extreme a single case as we can imagine since it involves the ultimate harm (death) doled out to a completely innocent person (a baby). Second, despite the instructions given with the example, it is never entirely certain in practice that the crying of the baby would lead to the other innocent people being found and killed; moreover, it is similarly unclear that killing the baby would prevent them from being found and killed. Third, any time we justify committing heinously immoral acts in some relatively public

context, we risk establishing a precedent for the commission of other, similar acts in the future. It is true that the likely fate of the others (and even of the baby), were they discovered, could weigh in favor of overriding the default of nonharming. But every other morally germane aspect of the situation weighs in favor of maintaining the deontological default, thus against murdering the child. (Of course, this might change in specific circumstances.)

Part of the problem with this and other cases of moral reflection in ethical deliberation concerns the nature of thought experiments in ethics, which tend to stipulate certainties inappropriately. Ethical thought experiments typically follow the general experimental principle of seeking to control variables. This includes controlling for outcomes. For instance, in a famous trolley dilemma, we are told that a trolley is rushing toward four people trapped on the tracks. We are also told that if we push a fat man onto the trolley tracks, he will be killed, but the trolley will be stopped, thus preventing the deaths of the four other people farther down the tracks. To my mind, the key point about this ethical (pseudo-) dilemma is that the whole scenario is wildly implausible. The man would presumably resist being pushed; he might fall only partially on the tracks; and even if he falls fully on the tracks, how can we possibly know with adequate certainty that his body will stop the trolley?

This problem with thought experiments is particularly crucial for evaluating how one should set the method parameter. The nature of thought experiments tends to bias them toward consequentialism. I fear that this may be a problem with literature as well. For example, literature (including films and television programs) about torture often shows us that the suspect really did plant a bomb, really does know where it is and how to defuse it, and so on. It gives us certainty about these issues that no police force is likely to have (in addition to ignoring the tendency of such harmful acts as torture to proliferate once allowed). In contrast, it is usually a central feature of our moral decisions that they are made in conditions of uncertainty.

But, of course, to adopt a deontological orientation, we need to determine what makes an act intrinsically ethical or unethical. Here, I to a great extent agree with Kant, if Kant revised his view in light of the preceding discussion. First, I would say that an excellent rule of thumb is – as stipulated in the preceding examples – to “act in such a way that you treat humanity, whether in your own person or in the person of another, always at the same time as an end and never simply as a means” (1981, 36). This tells us that we should always try to empathize with others and constrain our behavior toward them by reference to their interests, particularly their physical or emotional pain, which we should always seek to minimize (given the “ethics of relief” setting of the valence parameter).

(In addition, given our uncertainties and fallibility, we need to add to this empathy a degree of self-criticism, as I will discuss in Chapter 7.)

Kant's categorical imperative is also relevant here, primarily in indicating that we should not exempt ourselves from general moral principles, as we are very much inclined to do. But it is not relevant in the criterial way that Kant imagined. Kant believed that we could infer moral precepts by universalizing the maxim that covers our behavior in any given situation. But, as a way of deciding what is ethical or unethical, this is notoriously faulty. Most crucially, it sets aside "the problem of relevant act descriptions" (see Nell 1975, chap. 2), which is exactly what the method parameter should address. The case of the murderer leads to one ethical decision when construed as "lying versus telling the truth" but another when construed as "assisting versus impeding the murder of an innocent person." I actually agree with Kant that lying is morally wrong even in this case. Where I disagree with Kant is that I see this as a case of moral conflict. In cases of moral conflict, I give greater weight to the less harmful option. In this case, I favor "impeding the murder of an innocent person" over "telling the truth," principally because the death of innocents is more important ethically than the communication of misinformation.

#### *A Brief Note on Blame: Character versus Action*

I am of two minds as to whether to identify a fourth parameter, which identifies the primary target of ethical evaluation. Personally, I am a firm believer that acts are good or bad, not people. Moreover, I have also argued that our ethical theory must be based on the evaluation of at least some set of behaviors, as otherwise we have no way of distinguishing good from bad characters. But it does appear to be the case that many people think of ethical response and judgment as applying primarily to persons.<sup>14</sup> Indeed, as already noted, I suspect that a distinction along these lines is one reason for differences Haidt found in responses to his anecdote about a man masturbating with a chicken carcass (2012, 111–112) or the one about consensual incest. I suspect that many of the people who saw these behaviors as morally neutral were considering the act (in addition to confining ethics to issues of undeserved pain).

<sup>14</sup> This does not contradict the argument just mentioned because even agent-focused ethical response would need to be grounded in a sense that a (good or bad) person is, precisely, a person inclined to engage in particular sorts of (moral or immoral) acts. The difference between agent-focused and action-focused is not a matter of which is theoretically definitive. It is, rather, a matter of what one cares about ethically – what the target of one's emotion is and what one sees oneself as responding to.

In contrast, many of the people who judged the behaviors immoral might have been responding to a sense that the people involved had a bad character.

I do believe that these alternatives (agent-focused evaluation versus action-focused evaluation) are among the most important and fundamental ethical divisions. In that sense, they may be said to define possible parametric values for a fundamental ethical orientation. However, I have not included this as a fundamental parameter per se because I am inclined to view differences on this axis as deriving from other ideas and attitudes. Specifically, I am inclined to infer that we may differ in the degree to which we explain behaviors generally by reference to free choice along with dispositions (e.g., personality traits) versus the degree to which we explain such behaviors by reference to circumstances. It seems likely that an explanatory preference here would orient us toward an evaluative target. Thus, I am inclined to explain an agent-focused evaluative orientation by reference to a preference for dispositional (thus agent-focused) explanations of behavior. Similarly, a preference for situational explanations would seem more consistent with an action-focused evaluative orientation. This is not a specifically ethical parameter as it bears on explanatory preferences that go well beyond ethically relevant actions.

## **Conclusion**

There is a fundamental distinction between a descriptive study of ethics and a normative study of ethics. Descriptive ethical study investigates what ethical beliefs and attitudes may be found in or across different societies, historical periods, or individual people. It may stress the particularity of some local system, as would be the case with ethnography. However, it may also stress recurring patterns (as in Chapter 4 of the present volume). In contrast, normative ethical study treats what genuinely is ethical or unethical, rather than what people think is such. Ethical differences – for example, that between a focus on moral character and a focus on ethical deliberation – complicate both forms of study. Nonetheless, it is possible to isolate a coherent basic characterization of ethics that will allow discussion of both descriptive and normative topics.

Descriptive ethics proceeds by way of the usual methods of empirical study, conceptual analysis, and explanation through cognitive, affective, historical, or other relevant structures and processes. (The following chapters will stress cognitive and affective science.) Normative ethics cannot be established by empirical study. However, it is possible to rationally debate normative ethics, thus avoiding the apparent arbitrariness of relativism, without assuming some absolute foundation for morality. Such rational debate involves isolating shared ethical responses and

then identifying and testing principles relative to those responses. There are, however, problems with this method. They include, among others, the artificial and unrealistic nature of the moral cases that ethicists typically consider, as well as the limited overlap in people's fundamental moral orientations even in a single society.

But just how might we define ethics as an area of study? Our actions require motivation and thus emotion. These motivations are always self-interested in that they bear on our achievement of goals defined by emotion systems. Those goals, however, may be either egocentric or allocentric. In other words, they may be contingent upon one's own condition or someone else's condition. Ethics is fundamentally an attempt to limit egocentric motivation and action. Thus, ethical evaluation and action are prototypically allocentric. In other words, they are responses to the emotion-defined goals of other people and thus, broadly speaking, their happiness. This leads to my views, developed in subsequent chapters, concerning the ethical centrality of empathy.<sup>15</sup> This basic characterization of ethics has the salutary consequence that it relates prudence to ethics. Specifically, prudence is closely related to ethics as it involves a form of empathy with a person who is between self and Other – oneself in the future (as opposed to oneself right now).

In considering people's ethical attitudes, three parameters appear to be particularly important. The first is the scope parameter. This governs just what set of targets (e.g., people) are viewed as the proper objects of ethical deference. There seem to be two widespread tendencies here. Some people are inclined to see ethical obligation as strongest regarding groups that are closer to oneself. For example, in this view, one's obligations to one's family are greater than those to one's nation, which are in turn greater than those to humanity. The other tendency goes in the opposite direction, with one's obligations to humanity being greatest. However, in practice these groups often differ less than one might expect since social arrangements assign obligations in ways that do not necessarily track general obligations (e.g., even the most ardent universalist is likely to recognize that social practices commit him or her to providing special care for his or her own children). Theoretically, we may refer to these more specific qualifications on the general ethical orientations as subparametric variables.

A second parameter is valence. We tend to focus our ethical attention on either pleasure or pain and on the degree to which people merit or do

<sup>15</sup> It is worth noting that the idea is suggested by several lines of research. For example, as Zaki and Ochsner emphasize, one important evolutionary function of empathy is enabling cooperative activity (see 2012, 207). In connection with this, Tomasello's (2016) linking of ethics with cooperation suggests the centrality of empathy to ethics.

not merit their experiences of pleasure or pain. While it is in principle possible to attend to both, as a matter of fact it appears that most people are particularly concerned with one or the other. Thus, we find two broad types of ethical orientation – ethics of hedonic restriction (in which one condemns unmerited pleasure) and ethics of relief (in which one condemns unmerited pain).

The third parameter is that of method.<sup>16</sup> This concerns the manner in which one determines what is ethical or unethical. The two main alternatives here are intrinsic (or deontological) and consequentialist. The former evaluates acts as intrinsically ethical or unethical, whereas the latter relies on a calculus of consequences. In fact, it seems that both are in some ways intrinsic and in some degree consequentialist. On one hand, the deontological view that Jones does not treat Doe as an end if he plunges a knife into Doe relies on the fact that doing so will cause Doe pain and very likely lead to his death. On the other hand, a consequentialist calculation of an act's consequences for producing or curtailing happiness clearly presumes that happiness has intrinsic value (as it is not a means to something else). The difference between the two methods is in part a matter of which consequences one takes into account. In the deontological case, the consequences are not confined to the case at hand but are generalized to the behavior of all moral agents in an ethically ideal society. In a pure deontological ethics, the important consequences are hypothetical, bearing on what would happen if everyone behaved ethically. In other words, they concern an ethical ideal. Pure consequentialism, in contrast, concerns what is likely to occur in actual, often very nonideal circumstances. A reasonable alternative for this parameter might be that we should generally set it to the "intrinsic" value of an act. But (as a subparametric variable) we may override that intrinsic value in cases where the intrinsic immorality of the act is low, the relevant consequences of all alternatives appear clear, and the likely consequences of the intrinsically ethical act are highly morally objectionable (e.g., when telling the truth would lead to the death of an innocent person.)

A further, important orientational difference is that between agent focus and action focus. An agent-focused ethical orientation concentrates moral evaluation on the person. An action-focused orientation, in contrast, concentrates on the behavior. This is as fundamental and

<sup>16</sup> Since this is not a book devoted to articulating a complete system of ethics, there are necessarily many aspects of ethical deliberation that I cannot discuss. However, it is important to note again that I am not claiming these are the only parameters that operate in defining ethical systems. Even more clearly, my brief references to subparametric variables are nothing more than a gesture in the direction of an actual account of such variables.

consequential a parametric division as the others. However, it seems likely to derive from explanatory preferences, favoring personal dispositions or circumstances. Thus, I have not counted it as an ethical parameter *per se*.

Finally, it is important to add that the setting of parameters is not fixed immutably. It certainly appears as if we all have inclinations toward one or another valence, scope, method, and evaluative target. However, it also appears to be the case that we may change these orientations in different circumstances, with some of us being more or less likely to change than others. One important context for change is self-conscious ethical deliberation. When we reflect on our ethical orientation, we may change it significantly.