

Hare, Caspar. *The Limits of Kindness*.
Oxford: Oxford University Press, 2013. Pp. 229. \$40.00 (cloth).

The Limits of Kindness is an exciting book. Caspar Hare seeks to make progress in what he sees as the argumentative stalemate between deontology and consequentialism. He confesses to feeling pulled by both deontological and consequentialist intuitions, but unlike many others who feel such pulls, he reports that upon reflection and in the face of the usual arguments, he remains torn. Hare seeks another route to answering some central moral questions, and in this book he offers one.

Hare proposes to draw moral conclusions by reflecting on what a rational, morally decent person would prefer. He develops a new style of moral argument called a “morphing argument.” These arguments are ingenious. They start with modest, intuitively plausible premises and yet draw strong and surprising moral conclusions. I will illustrate by giving one example of a morphing argument.

q1

q2

THE HUMAN FUEL MORPHING ARGUMENT

You are captaining your boat alone, when you learn that Amy is in trouble on a nearby island. You save her life, but then you learn that Brian and Celia are in trouble on a far-off island. You are the only one who can save them, but you don't have enough fuel to reach them. You could save them by knocking Amy on the head and using her body as fuel. Amy's death would be quick and painless, while Brian's and Celia's deaths would be slow and painful. You must choose between killing Amy and letting Brian and Celia die. Hare argues that a rational, morally decent person would kill Amy.

Hare asks us to consider a series of possible worlds. The first is the world in which you don't kill Amy, and Brian and Celia die. The second world is very similar, but Amy's life is changed a bit to be like Brian's life, Brian's life is changed a bit to be like Amy's, and Celia lives a bit longer. (That is, Amy's personality and biographical details are changed to be more like Brian's, but her life is still just as long, and Brian's personality and biographical details are changed to be more like Amy's, but his life is still just as short.) The third world is one in which Amy's life is a bit more like Brian's (but still just as long), Brian's life is a bit more like Amy's (but still just as short), and Celia lives yet a bit longer. Jumping to the end, the last world is the world in which you do kill Amy, and Brian and Celia live long lives. The second-to-last world is very similar to the last world, but Brian's life is changed a bit to be like Amy's (but still just as long), Amy's life is changed a bit to be like Brian's life (but still just as short), and Celia's life is a bit shorter. Hare claims, and we should agree, that there is a series of possible worlds which (a) gradually morph the person in the Amy role in the first world (a person who lives a long life) from being just like Amy to being just like Brian, (b) gradually morph the person in the Brian role in the first world (a person who dies at the

For permission to reuse, please contact journalpermissions@press.uchicago.edu.

time of your decision) from being just like Brian to being just like Amy, (c) gradually morph Celia's life from a short life (one that ends at the time of your decision) to a long life, and (d) morph each of these three chains of persons in such a way that each person along the chain is better off than (or at least as well off as) the person immediately before her.

This is not a series of worlds in which Amy, Brian, and Celia all exist. Many of the worlds in the middle of the series do not have Amy or Brian in them (nor their counterparts, if one prefers to think in terms of counterparts)—rather, these worlds contain people who are roughly half like Amy and roughly half like Brian. Some of these people may be of indeterminate gender. No matter, Hare says: we suppose that in these worlds, there are no prejudices that systematically disadvantage the indeterminately gendered.

Hare claims that every decent person will prefer each world to its immediate predecessor in the series. Any two adjacent worlds in the series do contain people who are counterparts of each other: the persons in the two worlds differ only very slightly from each other. And each world is a *pareto improvement* over the predecessor world in the series: some people are better off, and everyone is at least equally well off.

Hare argues that every rational person will have transitive preferences and will act in accord with her preferences. Hare concludes that if you are a rational, decent person, you will prefer the last world in the series to the first world, and so you will kill Amy to save Brian and Celia. His conclusion is that you should kill Amy.

Before talking about how we might react to the morphing argument as a whole, let me mention an earlier argument Hare makes on which his morphing arguments rely. One thing that will strike us immediately about the morphing argument is that it does not make any mention of the different *causal relations* you stand in to the different possible deaths. In the first world in the series, you fail to save Brian and Celia, but in the last world, you kill Amy. So there must be some change in the causal relation between you and the three relevant people, as we proceed along the series of worlds. Whether this change is abrupt or gradual, it may be that considerations regarding your causal relations to the events that occur in each world warrant different preferences from what Hare says a decent person must have. Before presenting his morphing arguments, Hare argues that considerations merely regarding one's causal relations to events cannot affect a decent person's preferences. He offers a version of the "dirty hands" objection to deontology, the complaint that according to deontology, agents should care about whether they are getting their hands dirty rather than caring about people.

A deontologist might point out that in preferring not to kill Amy, you are caring about Amy, not just about yourself. Hare has a clever variant of the dirty hands objection which resists this point. He asks us to imagine a version of the above case in which you know that Amy, Brian, and Celia are all trapped on islands, but you do not know who is where. The one person you have just saved is taped up in a gorilla suit—and you cannot manage to undo the tape. All you know is that one of the three is here on your boat, and you must decide whether to kill her or him to save the other two. Hare argues that if you really care about Amy, Brian, and Celia, you will kill the person in the gorilla suit. As I understand

the argument Hare gives, we can gloss it as follows (though he doesn't put it this way): concern for Amy would have you choose the option in which she has a 1/3 chance of a short, painless death and a 2/3 chance at a long life rather than the option in which she has a 2/3 chance of a long, painful death and a 1/3 chance at a long life; similarly regarding concern for Brian and Celia; so, if you care about those three people, you should quickly, painlessly kill the person in the gorilla suit and use her or him as fuel to save the other two from slow, painful deaths. Hare says that if you prefer to refrain from killing in this case, it must simply be out of a self-centered desire that you yourself not do any killing. Hare uses this case to conclude that if an agent contemplates two scenarios, in the second of which the prospects for every person involved are better than in the first (or better for some and at least as good for others), then if the agent is morally decent, she will prefer the second scenario, regardless of how her causal contributions to the events in those scenarios may differ.

So, we can see that even before he gets to the morphing arguments, Hare is already arguing for a surprising conclusion: he argues that in the gorilla suit version of the rescue case, you should kill the person in the gorilla suit in order to save the other two people. There are a number of different ways a deontologist might try to resist Hare's argument here. Let me mention one. A deontologist might point out that we can name the person in the gorilla suit. Call that person "Sam." We can then think about what you would be doing to Sam and whether it can be justified. It seems that Hare's argument does not address this question. Killing Sam does not give Sam a 2/3 chance of survival.

Let's return now to thinking about the morphing argument for the conclusion that you should kill Amy in order to save Brian and Celia. One thing we should understand about the morphing argument is that it offers a series of worlds in which you act in the face of *different predicaments*. In the first and last worlds, you face the same predicament: you face a choice between the first and last worlds. But in each of the other worlds, you act in the face of different predicaments—different from the choice available in the first and last worlds and different from the choice available at each of the other worlds. The preferences Hare is talking about are not preferences between options a person faces in one situation but are preferences between actions a person takes in different situations. Now, consider the following query:

Which would you prefer?

- A. You cut off Joe's arm to save his life, when you're the only person who can save him and that's the only way to do it.
- B. You lie to Joe to cheat him out of a job he really wants so that you can have it for yourself.

Scenario A is one in which you act morally rightly, while scenario B is one in which you act morally wrongly. But Joe is better off in B than in A.

How would a morally decent person answer this query? Surely this is an odd question to ask. But Hare is asking this kind of question. There are two ways we might answer this question. Following them through can help us to explore how we might try to resist Hare's morphing argument.

First, let's suppose that a morally decent person would always prefer to act morally rather than immorally, even preferring to be in worse situations in which she acts well rather than better situations in which she acts poorly. On this supposition, a morally decent person would answer my query with option A. As we've already noted, the worlds in the morphing series differ in your causal relations to the events in question. If we answer the query in this way, we may resist Hare's claim that just because the fates of the people in the three roles improve along the series, a morally decent person would prefer the latter worlds in the series to the earlier worlds. Rather, a morally decent person would stop preferring the next world in the series as soon as we reach a world that involves killing one person and using her body to save two. (Or perhaps she would stop sooner, when we reach a world in which what the agent does is indeterminate between killing and letting die, if there are such possibilities.) As Hare would see it, this way of resisting his arguments would involve taking morally decent people to be immorally preoccupied with themselves and what they make happen.

Second, let's suppose that a morally decent person would always prefer better situations to worse situations, even preferring to be in a better situation in which she acts morally wrongly to a worse situation in which she acts morally well. On this supposition, a morally decent person would answer my query with option B. If we are understanding preference in this way, then we may reject Hare's claim that a morally decent person would be guided by her preferences. Because a morally decent person will always act morally permissibly, but her preferences may commit her to sometimes favoring worlds in which she acts impermissibly rather than permissibly, a morally decent person may sometimes act contrary to the ways she would prefer. This way of looking at things may seem to make a lot of sense from a deontological perspective. "I wish I could help him, but it would be morally wrong" is something that a moral agent might well find herself thinking, just as she might think "It would be better if things went that way, but it would be wrong to make them go that way." But Hare sees a deep irrationality in this cleavage between preference and action. And he is right that resisting his arguments in this way commits one to seeing preference as irrelevant to action in a surprising way.

OTHER MORPHING ARGUMENTS

Hare offers morphing arguments for several different moral conclusions, not all as radical as the claim that you should kill Amy. For example, Hare offers a morphing argument as a solution to a same-number nonidentity problem. Suppose you are trying to decide whether to conceive now or to wait a few months. You have German measles, and if you conceive now, your conceived child will suffer from a serious health problem that will cause pain throughout his life, but he will still have a life worth living. If you wait to conceive, your condition will clear up and you will be able to conceive a healthy child. Intuitively, it is morally wrong to conceive now, but conceiving now would not be worse for your child. If you conceive now, then it will be true that had you waited, your child would not have existed; it will not be true that he would have been better off if you had waited. Hare offers an argument that involves a series of worlds that morph the child you would have if you conceive now into the child you would have if you

wait. If you are rational and decent, you will prefer the last world in the series. So you should wait to conceive.

Hare's conclusion here is the sensible one: you should wait to conceive. One might worry, however, that the argument is too strong. The same conclusion would follow, via the same kind of argument, if you were now in a position to conceive a healthy child but by waiting you could conceive a child who would experience some special kind of positive benefit. In such a case, it would not be morally required to wait, yet Hare's arguments would imply that it would be morally indecent to conceive now.

Hare also argues that morphing arguments can establish, for certain kinds of cases, that one should save more people rather than fewer people. Another sensible conclusion.

And Hare also offers a morphing argument for the conclusion that a morally decent person would make a small sacrifice to save the distant needy. His conclusion that one should make such a sacrifice is another controversial claim.

Hare offers a new moral picture, and, as his title indicates, he also explores its limits. One such limit is that morphing arguments cannot show that one should benefit one person a lot rather than another person slightly less, if the agent knows a lot about the people in question. That's because, from the agent's perspective, the quality of their two lives may be *on a par* (neither better than the other and yet not equally good). This inhibits the formation of a strictly improving morph from one person to the other. Thus, it turns out that on Hare's view, whether you know a lot about the people involved in certain situations affects what is morally required. In fact, Amy is safe from having to be killed, if you know enough about her. (Note that it's mere knowledge that matters here, as opposed to loving relationships.) Another limit to kindness arises if there are infinitely many people: if so, then one cannot take a benevolent attitude toward them all, Hare argues. One provocative result of this is that whether you should prefer that a particular person is saved from suffering may depend on whether you believe there are infinitely many other uninvolved people in existence.

The Limits of Kindness offers a big new idea. It is delightfully fun to read, and yet it is extremely clear and careful. If, like me, you find its conclusions somewhat maddening, you will read the book as a challenge to what you believe; it is a challenge worth answering.

ELIZABETH HARMAN
Princeton University

QUERIES TO THE AUTHOR

q1. AU: Italics omitted here and throughout when used for emphasis yet meaning seems clear without them or when not at first use of key terms.

q2. AU: Sentence beginning “I will illustrate” ran into prior paragraph to avoid single-sentence paragraph, to conform to style.