

Our Duties to Future Generations

By

Molly E. Gardner

A dissertation submitted in partial fulfillment of

the requirements for the degree of

Doctor of Philosophy

(Philosophy)

at the University of Wisconsin-Madison

2013

Date of final oral examination: 05/09/13

The dissertation is approved by the following members of the Final Oral Committee:

Robert Streiffer, Associate Professor, Philosophy, Medical History and Bioethics

Russ Shafer-Landau, Professor, Philosophy

Harry Brighouse, Professor, Philosophy

Peter Vranas, Professor, Philosophy

J. Paul Kelleher, Assistant Professor, Medical History and Bioethics

In memory of my mother, Valerie Gardner

Table of Contents

Introduction.....	p. 1
Chapter 1: Why We Have Duties to Future Humans.....	p. 11
1.0. Introduction.....	p. 11
1.1. The Utilitarian Objection.....	p. 14
1.2. The Relational Factor Objection.....	p. 15
1.3. The Epistemic Indeterminacy Objection.....	p. 17
1.4. The Metaphysical Indeterminacy Objection.....	p. 18
1.5. Conclusion.....	p. 24
Chapter 2: How We Can Violate Our Duties to Future Humans.....	p. 25
2.0. Introduction.....	p. 25
2.1. The Harm of Carbon Dioxide.....	p. 27
2.2. The Harm of Species Extinction.....	p. 30
2.3. The Fungibility Objection.....	p. 32
2.4. Conclusion.....	p. 37
Chapter 3: The Non-Identity Problem.....	p. 38
3.0. Introduction.....	p. 38
3.1. The Non-Identity Problem.....	p. 39
3.2. Some Attempts to Solve the Problem.....	p. 48
3.3. Denying the Counterfactually Worse-Off Condition on Harming.....	p. 55
3.4. Conclusion.....	p. 63
Chapter 4: The Existence Account of Harming.....	p. 64
4.0. Introduction.....	p. 64
4.1. Why H is not a Satisfactory Account of Harming.....	p. 65
4.2. The Existence Account of Harm.....	p. 71
4.3. The Existence Account of Harming.....	p. 73
4.4. Conclusion.....	p. 81
Chapter 5: Causes and Mere Conditions.....	p. 82
5.0. Introduction.....	p. 82
5.1. The Problem of Too Much Harm.....	p. 83
5.2. Causes and Mere Conditions.....	p. 88
5.3. Problems for the Comparative Account.....	p. 94
5.4. Conclusion.....	p. 99

Chapter 6: Benefiting by Causing Existence	p. 101
6.0. Introduction.....	p. 101
6.1. The Argument.....	p. 101
6.2. Benefiting by Causing Existence.....	p. 108
6.3. Objections and Replies.....	p. 111

Acknowledgements

This project would not have been possible without the help of my colleagues, friends, and family. I am immensely grateful to Robert Streiffer for his valuable comments on earlier drafts, and I also thank Paul Kelleher, Justin Weinberg, Russ Shafer-Landau, Harry Brighouse, and Peter Vranas for helpful discussions, comments and feedback. For support throughout the process, I thank my family, Harry, Diane, Bob, Daisy, Lew, Alex, Addie, and Phyllis.

Introduction

My project, in this dissertation, is to explicate some of the moral duties that members of present generations have to members of future generations. I will use the term “future humans” to refer to humans who do not yet exist, but who will exist in the future. I will defend the view that (DV1) individual moral agents who are now alive have pro tanto duties of nonmaleficence and beneficence *to* and *regarding* at least some future humans; (DV2) in the present circumstances, this duty of nonmaleficence grounds reasons for all presently living moral agents to refrain from damaging certain features of the natural environment; and (DV3) in the present circumstances, this duty of beneficence grounds reasons for at least some presently living moral agents to bring future humans into existence. For ease of reference, I will refer to the view consisting of (DV1) – (DV3) as the Duties to Future Humans View (DV).

I will begin by making some clarifications. First, DV is not intended as an exhaustive account of our duties to future humans. In addition to duties of nonmaleficence and beneficence, we might, for example, owe future humans duties of justice that enjoin us to distribute resources to them fairly or to punish those who violate their rights.¹ My hope is that what I say about nonmaleficence and beneficence will at least lay some groundwork for the project of explicating whatever additional moral duties to future humans we might have.

¹ More precisely, there may be duties of justice that enjoin us to support just institutions. Kamm (1999, p. 180) writes, “Notice that a duty to promote basic justice may obligate us only to support just institutions, not to aid directly.” If so, then there might be an interesting difference between duties of nonmaleficence and beneficence, on one hand, and duties of justice on the other: namely, the former might be *to* future people *regarding* future people, whereas the latter might be *to* future people *regarding* institutions.

A second clarification concerns what I mean by both “nonmaleficence” and “beneficence.” I use the term “duty of nonmaleficence” to refer to a duty to refrain from harming, and I use the term “duty of beneficence” to refer to a duty to benefit.² I will later argue that harming is causing a harm and that benefiting is causing a benefit. The importance of causation to my view explains why I will defend substantive accounts, not only of harm, benefit, and duties *to* an individual, but also causation. I intend for DV to be interpreted in light of these substantive accounts.

Here is a third clarification. My claim that we have duties both *to* and *regarding* future humans is intended to be stronger than the claim that we merely have duties *regarding* them. To get an intuitive grasp of the distinction between duties *to* and duties *regarding*, consider some cases in which the party to whom a duty is owed is distinct from the party whom the duty regards: Jones has a duty to Smith to water her plants, but this is a duty *regarding* Smith's plants. Kate has a duty to protect William from one of his obsessive fans, but this is a duty *to* William *regarding* his fan.

Many people suppose that the distinction between duties *to* and duties *regarding* can be adequately characterized at the normative level. It is commonly held, for example, that as normative theories, Kantianism and utilitarianism differ in regard to whom they say we owe duties. Kant is usually interpreted as holding that although we owe a duty of respect to other

2 Some people might want to distinguish between a duty to rescue or assist someone in peril from a duty to benefit someone who is not in peril. For example, Tom Regan (2004) distinguishes between a duty to assist victims of injustice and a duty of beneficence. While I grant that there may be stronger reasons to help someone who is in peril or the victim of an injustice, I contend that there is a single duty of beneficence that generates reasons to help those in peril, to assist victims of injustice, and to improve the well-being of those who are neither in peril nor victims of injustice. This will become clearer when I provide my account of benefiting in Chapter 6.

rational agents, we do not owe such a duty to nonhuman animals. This is not to say that we lack moral reasons to refrain from harming nonhuman animals; rather, our moral reasons to refrain from harming them are grounded in our duties to other humans.³ Thus, the Kantian view is supposed to be that we have duties to other rational humans *regarding* nonhuman animals, but we do not have duties *to* nonhuman animals.

On the other hand, some utilitarians argue that, like most humans, many nonhuman animals have *direct moral standing*, and they define “direct moral standing” (or more simply, “moral standing”) as the status an individual has when we have duties *to* him or her. In the context of a utilitarian theory, this is taken to mean that in our moral deliberations, we have a duty to give nonhuman animals' interests the same weight that we would give to the similar interests of humans: the duty to count all animals' interests equally is supposed to be a duty *to* each animal.⁴

This normative-level characterization of the distinction goes only part of the way towards clarifying my conclusion. On my view, our duties to future people are somewhat like the duties we are supposed to have to rational agents on a Kantian view or to sentient animals on a utilitarian view, and not like the duties we are supposed to have *regarding* nonhuman animals on

3 Kant writes, “If a man shoots a dog because the animal is no longer capable of service, he does not fail in his duty to the dog, for the dog cannot judge, but his act is inhuman and damages in himself that humanity which it is his duty to show towards mankind” (quoted in Regan 2004, p. 178).

4 Some utilitarians distinguish the criterion for the correct decision procedure from the criterion for right action. If utilitarianism is viewed strictly as a criterion for right action, then we have no duty to *literally* count each animal's interests equally; rather, we have a duty to take the action that maximizes utility according to a function that gives each animal's interests equal weight (and fulfilling *this* duty is a way of fulfilling our duties *to* animals). The dilemma I will soon raise for the utilitarian account of “duties to” will work equally well against this more precise account of what our utilitarian “duties to” animals are.

a Kantian view. Just as the Kantians and utilitarians claim that we can *wrong* rational agents and nonhuman animals, respectively, my view implies that we can *wrong* future humans. Thus, the matter of whether our actions regarding future humans are wrong is not settled by whether those actions offend presently living humans.

Nevertheless, I am dissatisfied with the practice of drawing the distinction between duties *to* and duties *regarding* at only the normative level. At that level, there are few resources for explaining why a particular duty qualifies as a duty *to* and not merely a duty *regarding*. Moreover, if we rely solely on our intuitions about the distinction, we can characterize utilitarianism and Kantianism in ways that suggest that our duties on those theories are not, after all, to whom we had thought. For example, utilitarians hold that our only *fundamental* moral duty is to maximize utility, and it seems that *that* duty is not the kind of duty that could be owed to anyone in particular. Now, either the duty to assign similar weight to similar interests must *derive* from the duty to maximize utility, or it must have independent normative force. If the former is true—if the *only* reason to count animals' interests equally is that doing so will maximize utility—then it seems that the duty to count animals' interests equally is really a duty *regarding* animals, and not a duty *to* animals.⁵ On the other hand, if the duty to assign similar

5 Notice that if we have no duties to individuals, then we have no duties to individuals to either benefit them or to refrain from harming them. This could explain why some utilitarians trivialize the moral significance of harm. For example, Alastair Norcross (2005) offers a contextualist account of harm that ultimately reduces harm to a linguistic issue. He writes,

[F]or the purposes of ethical theorizing, harm and benefit do not have the kind of metaphysical grounding required to play fundamental roles in ethical theory, nor do judgments of harm or benefit make any distinctive contributions to reasons for action. If, in considering whether to do A, I correctly judge that A would harm P, I am judging that A would result in P being worse off than s/he would have been if I had performed the appropriate alternative action. This is certainly a relevant consideration from a consequentialist perspective, but it is also one that I would already have taken into account, if I had considered all my available alternatives (pp. 171-72).

weight to similar interests has independent normative force, then utilitarianism is not the monist theory it purports to be.⁶ Likewise, it could be argued that on Kant's view, the duty to treat rational agents as ends in themselves is not a duty to those individual agents, but rather a duty *regarding* such agents, either *to* the humanity in oneself, or to humanity in the abstract.

To establish whether utilitarianism and Kantianism really do acknowledge duties to other animals or to other rational agents, respectively, we require a more precise account of the distinction between duties to and duties regarding. Such an account would make reference to the ultimate sources of normativity and would therefore be metaethical, not normative. Here is the account I propose:

Duty to (definition): An agent, A, has a duty, D1, to an individual B if and only if A has D1 *because* B exists and not *merely* because she has a more fundamental duty, D2, which she would have even were B not to exist, such that fulfilling D1 is a way of fulfilling D2.

In Chapter 1, I will defend this account more fully against the utilitarian claim that we can have “duties to” sentient beings in the sense that we ought *merely* to count their interests equally.

When I claim that we have duties to future humans, I do not mean to suggest that we have no other duties or sources of reasons, such as a duty very much like the utilitarian duty to maximize value. Indeed, I have argued elsewhere that our reasons for action may be grounded

6 Of course, not all versions of utilitarianism purport to be monist in the same way. Classical hedonism holds that our only duty is to maximize the good and that the good is happiness. My argument is that if a classical hedonist were to draw a fundamental distinction between the goodness of happiness and the goodness of *counting* happiness in the right way, her theory would collapse into pluralism. G.E. Moore's (1903) version of utilitarianism, on the other hand, is already pluralist at the normative level, if not at the foundational level: while Moore holds that our only duty is to maximize goodness, he claims that a plurality of different things can instantiate goodness. Still, even a Moorean could not consistently claim that a duty to give similar weight to similar interests did not derive from the more fundamental duty to maximize the good. So it still seems that if a Moorean were to remain a monist at the *foundational* level, she would have to hold that our duties to count animal interests equally would be a duty *regarding* animals.

both in duties to individuals and in considerations of what I call “impersonal value,” or value from an impersonal perspective. Sometimes the moral status of a single action is determined by both types of considerations. For example, there may be moral reasons against an action, both because it harms someone *and* because it results in less well-being in the world than some alternative action. Many philosophers have argued that the moral status of any action regarding future humans is determined *only* by impersonal considerations: we act rightly when we bring into existence the particular individuals and the particular living conditions that will maximize the goodness of the world.⁷ However, I hold that the moral status of our actions regarding future humans hinges upon other factors besides impersonal considerations. Specifically, there are moral reasons for or against our actions depending on whether we are benefiting or harming future humans. In this dissertation I set aside impersonal considerations and focus *only* upon the subset of moral considerations that have to do with causing benefits and harms.

For those who prefer the language of rights to the language of duties, the first tenet of DV could be reformulated as the claim that future humans have rights against us that we both benefit them and refrain from harming them. However, the duties posited by the DV are pro tanto duties, which means that countervailing considerations can sometimes override them; any rights that my view establishes are therefore non-absolute. Moreover, the claim that anyone has a right to be benefited is controversial. Since on my view, such a right would not be absolute, and since I find it plausible that the duty of nonmaleficence is more stringent than the duty of beneficence, I

⁷ Although I think it is an incomplete account of morality, I am still sympathetic to this view. For an argument that the duty to maximize impersonal value does not lead to what Derek Parfit calls the “Repugnant Conclusion,” see Gardner and Weinberg (2013).

would not object to the claim that future humans have a right to be benefited, where such a right is more easily overridden than the right to not be harmed. However, many people—even those who acknowledge a general duty of beneficence—might still find it implausible that anyone would have such a right. Those who reject the existence of a right to be benefited are therefore invited to understand my view merely as a view about pro tanto duties and not as a view about rights.

Recall that according to DV2, in the present circumstances, the duty of nonmaleficence we owe to future humans grounds reasons for us to refrain from damaging certain features of the natural environment. The claim that in a particular set of circumstances, pro tanto duties can give rise to reasons is a metaethical claim that I take for granted. It is also controversial. There are, for example, alternative views according to which there *are* no moral duties, or if there are, there are no moral reasons that are less fundamental than such duties. Since an argument against these views is beyond the scope of my dissertation, I ask the reader to suppose, for the sake of argument, a framework on which the grounding relations between duties, particular sets of circumstances, and reasons license inferences of the form,

- (i) A has a pro tanto duty to B to (not to) ϕ .
- (ii) In circumstances C, A can ϕ by ψ -ing.
- (iii) Therefore, in circumstances C, there is a reason grounded in A's duty to B to (not to) ϕ for A to (not to) ψ .

My central argument can then be understood as an argument for the claim that if the above argument form is metaphysically valid, then DV (which includes my substantive accounts of duties *to* an individual, harm, benefit, and causation) is true.

Here is one final point of clarification. DV is a view about our duties to future humans. As I noted above, I use the term “future humans” to refer to humans who do not now exist, but who will exist in the future. I use the term “human” rather than “person” because “human” is a biological term and “person” is a moral term: a person is an individual who has full moral standing. Because part of what is at issue is whether future humans have moral standing, I will call the individuals about whom I am concerned “future humans” and then argue for the claim that at least some of them have moral standing and that their having moral standing implies that we have duties to them. I will not argue that all future humans have moral standing; I leave it open whether future humans who lack cognitive functioning, for example, have moral standing. Nor will I argue that only humans (either present or future) have moral standing. I encourage the reader to remember that the claim that we have various duties to at least some future humans does not imply that we lack such duties to future nonhumans.

Here is how the dissertation will be organized. In Chapter 1, I will present a positive argument for DV1, which is the claim that individual moral agents who are now alive have pro tanto duties of nonmaleficence and beneficence to and regarding at least some future humans. The argument will rely on the premises that future humans have moral standing, and that their having moral standing is sufficient for our having pro tanto duties of nonmaleficence and beneficence to and regarding them. I will then consider and respond to objections.

In Chapter 2 I will introduce my argument for DV2, the claim that in the present circumstances, our duty of nonmaleficence to future humans grounds reasons for us to refrain from damaging certain features of the natural environment. Specifically, I argue that there are

such reasons for us to refrain from causing carbon dioxide to be released into the atmosphere and for us to refrain from causing large numbers of species extinctions. To support this claim I will argue that excessive carbon dioxide in the atmosphere is a harm for future humans, and so is a pronounced loss of biodiversity. Although I will argue for substantive accounts of harm and causation in Chapters 4 and 5, I rely in Chapter 2 on pretheoretical intuitions about these notions. I then consider and respond to a preliminary objection to the claim that there is a non-maleficence reason for us to refrain from damaging the aforementioned features of the environment. The objection is that the very actions that damage the environment in these ways also cause living conditions for future humans that negate any potential harm.

In Chapter 3 I explain the non-identity problem, which is the problem of explaining how an individual can in some way be wronged by an action that is a condition of her own worthwhile existence. I show how the problem gives rise to a second objection to the claim that our pro tanto duties to future people establish moral reasons not to damage the environment. I survey some proposed solutions to the non-identity problem and argue that the most promising of these solutions is to reject the worse-off condition on harming, which holds that an action harms an individual only if it makes her worse off than she would otherwise have been. I consider and respond to the objection that denying the worse-off condition on harming undermines the justification for the view that harming is always or at least usually morally objectionable.

In Chapter 4, I defend a substantive account of harming. The account holds that for any individual, S, and any state of affairs, A, the following are true: (1) A is a harm for S just in case if it were true that both S existed and A did not obtain, then S would have a higher level of well-

being; (2) an action or event harms S just in case it causes a harm for S to obtain; (3) other things being equal, the reason against harming S is stronger, the more similar the world would be if A did not obtain.

In Chapter 5 I motivate the need for a substantive account of causation by arguing that a Elizabeth Harman's (2004, 2009) harm-based solution to the non-identity problem does not work because it relies upon an untenable assumption about causation. I then argue that not only my own account of harming, but other comparative accounts of harming will also get the wrong results if they are paired with an account of causation that does not support a metaphysical distinction between causes and mere conditions. I suggest that fortunately, there are plausible accounts of causation that do support such a distinction, and I endorse Alex Broadbent's (2007a) account.

In Chapter 6 I defend the claim that our pro tanto duties to future people establish reasons for some of us to act in ways that will bring future people into existence. I argue that, just as some procreative actions can harm those whom the actions bring into existence, so can some procreative actions benefit those whom they bring into existence. Since there is a pro tanto reason to benefit future people, there is a pro tanto reason for some of us to cause some future people to exist.

Chapter 1: Why We Have Duties to Future Humans

1.0. Introduction

In this chapter, I will defend DV1, which is the claim that individual moral agents who are now alive have pro tanto duties of nonmaleficence and beneficence *to* and *regarding* at least some future humans. For ease of expression, I will refer to individual moral agents who are now alive as “we” and I will use the term “owe” when I want to say that an agent has a duty *to* someone. The reader should also understand me as claiming that our duties of nonmaleficence and beneficence are duties *regarding* those to whom they are owed unless I state otherwise. (I am not defending the claim that we have duties to future humans not to harm future crocodiles or to benefit Donald Trump, although I will, of course, argue later that there are reasons grounded in our duties to future humans for us to refrain from damaging the environment.) The first tenet of DV can then be reformulated as the claim that we owe at least some future humans pro tanto duties of nonmaleficence and beneficence.

Here is my main argument for that claim:

The Argument from Moral Standing

- (M1) We owe all individuals who have moral standing pro tanto duties of nonmaleficence and beneficence.
- (M2) At least some future humans have moral standing.
- (DV1) Therefore, we owe at least some future humans pro tanto duties of nonmaleficence and beneficence.

Let me first explain M1. Recall from the introduction that some utilitarians purport to use the term “moral standing” to refer to those to whom we have moral duties. Many non-utilitarians also claim to follow this usage. Some of those who use the term “moral standing” appear to take

it as an *analytic* truth that all and only those who have moral standing are owed moral duties; the term “has moral standing” is just a *synonym* for “is owed moral duties.” However, I raised doubts in the introduction about whether utilitarians can consistently hold that we have duties *to* individuals. If it turns out that they cannot, we should not want it to also turn out that, according to utilitarianism, no individual has moral standing. Thus, I think we should take “moral standing” to refer to the kind of status an individual has when her interests matter morally, independently of anyone else's interests, and we should accept the account of *duties to* that I offered in the introduction. That account, again, is the following:

Duties to (definition): An agent, A, has a duty, D1, to an individual B if and only if A has D1 *because* B exists and not *merely* because she has a more fundamental duty, D2, which she would have even were B not to exist, such that fulfilling D1 is a way of fulfilling D2.

On my view, the claim that *all of those who have moral standing are owed moral duties* is not analytic, but rather a substantive claim. Nevertheless, such a claim is intuitively plausible. If an individual's interests matter morally, then whatever it is that grounds the fact that they matter morally will also ground a moral duty to honor those interests, where honoring those interests entails benefiting or refraining from harming the individual.

Premise M2 does not specify what the sufficient conditions for having moral standing are. That is because, even without knowing what those conditions are, we can still reasonably infer that some future humans will meet them by way of the following:

The Parity Argument

- (P1) At least some presently living humans have moral standing.
- (P2) At least some future humans have all the same grounds for moral standing that these presently living humans do.
- (M2) Therefore, at least some future humans have moral standing.

I take P1 to be commonsense judgment that does not require much argument. Here is the reasoning behind P2: whatever the features are that make us morally significant, those features probably have to do with our biology and possibly our natural environment. Because of our biological make-up (and possibly the environment), most of us experience pleasure and pain, for example, and have desires, preferences, and life stories. But if biology and our environment grounds our experiences, preferences, and life stories; and if our experiences, preferences, or life stories (or something similar) ground our moral standing, then our biology and our environment ground our moral standing. Even if the human genotype and the environment changes over time, it will be the case that at some future time, at least *some* humans have similar enough genes and live in a similar enough environment that they share with us the same grounds for moral standing. If so, then it is true *now* that these beings who are so similar to us, albeit separated from us in time, have moral standing.

Here is how the rest of the chapter will be organized. Since the plausibility of DV1 hinges almost entirely upon the plausibility of both M1 and P2, I will consider and respond to a series of objections to either M1 or P2. It is not always easy to determine whether a consideration is supposed to undermine M1 or P2, so in some cases, I will introduce an objection as an objection to DV1 and then consider whether it can defeat *either* M1 or P2 before moving on to the next objection.⁸

8 It is difficult to distinguish objections to 1M from objections to 2P precisely because so few people distinguish claims of the form, “x has moral standing” from claims of the form, “we have duties to x.” For example, the following passage illustrates a widespread and mistaken tendency to equate not only an individual's having moral standing with the individual's being owed duties, but also the *strength* of the individual's moral standing with the *strength* of our duties to him or her: “The fact that I stand in such-and-such a relation to animal, A, may itself give rise to an obligation on my part – an obligation that effectively elevates A’s moral status. (To be the

1.1. The Utilitarian Objection

The first objection is clearly directed towards M1, which claims that we owe all individuals who have moral standing pro tanto duties of nonmaleficence and beneficence. The objection specifically takes issue with the way I have defined “duties to.” While a utilitarian might grant that we have “duties to” future humans in the sense that each future human's interests must be given the same weight as the similar interests of any other sentient being, she would deny that our duties to future humans carry normative weight that is independent from the fundamental duty to maximize utility.

The problem with this view is that it has counterintuitive implications. Consider, for example, the following:

Lisa's Choice: Lisa wants to be a mother. Because she has a medical condition, she faces an unusual set of options. Either she can take a fertility drug that will cause her to have sextuplets, or she can choose not to take the drug, in which case she can only have one child, who will not be identical to any of the sextuplets. If she takes the drug she will name her children A, B, C, D, E, and F, and if she does not take the drug, she will name her child G. Each of A, B, C, D and E will have a happy life that measures 100 on the scale of well-being, but F's life, full of nothing but pain and misery, will measure -100 on the scale. On the other hand, G will have a happy life of 100, just like A, B, C, D, and E.

The commonsense judgment about Lisa's Choice is that, other things being equal, Lisa ought not to take the drug. The best explanation for this judgment is that if she takes the drug, then she wrongs F, but if she doesn't take the drug, she does not wrong anyone. Now, on the utilitarian account of what it is to have a *duty to* someone, Lisa does not violate any duty she has to F when

beneficiary of an obligation, I assume, is to have moral status – to count morally, to take up moral space)” (Burgess-Jackson, 1998, p. 175).

she brings him into existence. After all, her action is consistent with her having given equal consideration to F's well-being. We can even suppose that in the scenario in which she takes the drug, she explicitly counts each unit of well-being in F's life and weighs it against each unit of well-being in the lives of A, B, C, D, and E. It simply turns out that taking the drug generates a total of 400 units of well-being, whereas failing to take the drug generates only 100 units. When she takes the drug, Lisa is maximizing utility.

In short, the utilitarian view of what it means to have *duties to* an individual does not account for the intuition that Lisa violates a duty to F when she takes the drug. Thus, the utilitarian's objection that M1 is true only on a utilitarian account of *duties to* is unsuccessful.

1.2. The Relational Factor Objection

The second objection holds that we do not owe future humans pro tanto duties of nonmaleficence and beneficence because future humans are not related to us in the right way. That is, there is some relation that an individual must bear to us, either in order that she have moral standing, or in order that we owe duties of nonmaleficence and beneficence to her, and future humans do not bear that relation to us.

What relation could this be? Some philosophers have argued that the strength of our duties to assist strangers in need varies according to whether those strangers are proximate⁹;

⁹ For an argument that proximity in space affects the strength of our obligations to help strangers, see Kamm (1999).

included in our social circles¹⁰; or able to arouse our feelings of empathy¹¹. Although these factors are cited in the context of the debate about our duties to strangers in developing countries, such strangers have much in common with future humans. Like strangers in developing countries, future humans are not proximate; they are separated from us by a significant distance (if not in space, then in time, or more accurately, in spacetime). Likewise, future humans are not included in our social circles, and it is likely that we would feel less empathy for future humans than we would for presently living humans.

Nevertheless, even those who argue that a relational factor like proximity, social distance, or empathy makes a moral difference claim only that the factor *weakens* our duties of assistance. They do not claim that the factor *precludes* our having duties of assistance, nor do they make any claims about the relation between such a factor and our duty of nonmaleficence.¹² Further, I do not think there *could* be a plausible justification for the claim that a factor like proximity, social distance, or empathy could either preclude future humans from having moral standing or preclude our having any duties of nonmaleficence or beneficence to them.¹³ Thus, the premise

10 For an argument that social distance is relevant to the strength of our obligations to help strangers, see McGinn (1999).

11 For an argument that the strength of our duties of assistance to other individuals is affected by how much empathy the plight of those individuals would elicit in a normally developed person, see Slote (2005).

12 For example, in response to Peter Singer's (1972) argument that we have strong duties to help poor people in developing countries, McGinn writes, "I am with Singer in believing that the richer nations of the world do far too little to aid the poorer nations ..." (1999, p. 160), and Slote writes, "At the very least, then, even if Singer exaggerates what morality demands of us, it may nonetheless be true that many of us should give a good deal more for the relief of famine or disease around the world than we actually do" (2005, p. 153).

13 Moreover, whether we have a reason to refrain from contributing to climate change is also a separate issue from how strong such a reason is, so even if factors like proximity or social closeness weaken our duties to future humans, they do not undermine my argument in Chapter 2.

that relational factors like proximity, social closeness, or empathy make a moral difference simply does not support the conclusion that we do not owe duties of nonmaleficence and beneficence to future humans.

1.3 The Epistemic Indeterminacy Objection

The third objection to the claim that we owe future humans pro tanto duties of nonmaleficence and beneficence appeals to our ignorance of the identities of future humans. Joel Feinberg (1974) articulates the objection as follows: “The real difficulty is ... that we don't know who [our descendants] will be. It is not their temporal remoteness that troubles us so much as their indeterminacy—their present facelessness and namelessness” (p. 181). Feinberg's use of the term, “know ” indicates that he is concerned with epistemic indeterminacy. One could also appeal to metaphysical indeterminacy by arguing that there is no *fact of the matter* who our descendants will be; I will consider the appeal to metaphysical indeterminacy later. The objection here is that epistemic indeterminacy—our lack of knowledge of the identities of future humans—either precludes future humans from having moral standing or precludes our having duties to them.

After articulating the objection, Feinberg offers a response to it. He argues that knowing that individuals are human beings is sufficient for having duties to them, and any lingering epistemic indeterminacy fails to void those duties:

Five centuries from now men and women will be living where we live now. ... The identity of [such individuals] is now necessarily obscure, but the fact of their interest ownership is crystal clear, and that is all that is necessary to certify the coherence of

present talk about their rights. We can tell, sometimes, that shadowy forms in the spatial distance belong to human beings, though we know not who or how many they are; and this imposes a duty on us not to throw bombs, for example, in their direction. In like manner, the vagueness of the human future does not weaken its claim on us in light of the nearly certain knowledge that it will, after all, be human (p. 181-82).

A modified version of Feinberg's bomb-throwing example supports a conclusion that is stronger than the moral he draws from it. If the shadowy forms in the distance are normal humans, there is a moral reason for us not to throw a bomb at them *regardless* of whether we *know* that the shadowy forms are humans. There can be reasons for us to act even if we do not know about these reasons, and these reasons can be grounded in duties we do not know we have. In any case, either Feinberg's weak conclusion or my stronger one is sufficient to undermine the claim that our knowing the identities of particular future people is necessary, either for their having moral status or for our having duties to them; the epistemic indeterminacy objection fails.

1.4. The Metaphysical Indeterminacy Objection

A fourth objection to DV1 claims that we do not have duties to future humans because their identities are *metaphysically* indeterminate. To say that the identities of future humans is metaphysically indeterminate is to say that there is no fact of the matter who future humans will be. We can perhaps make true claims about future humans using definite descriptions (“future humans,” “the humans who will live in 2153,” “the humans who will live in Australia in 3024”), but these claims will be true in virtue of something other than the existence of concrete individuals who meet the descriptions.¹⁴

¹⁴ If we can have *de dicto* duties to groups of indeterminate individuals the way that the President, for example, might have *de dicto* duties to “the taxpayers,” whoever the taxpayers might be, then perhaps we could still have

Notice that this objection is not merely the denial of physical determinism. According to physical determinism, the laws of physics and the conjunction of all the facts about the world at any given point in time are jointly sufficient to determine the truth values of any proposition about the world at any future point in time.¹⁵ If physical determinism is false—i.e., if physical *indeterminism* is true—then there are some propositions about the future whose truth values are not determined by the facts about the past plus the laws of physics. This does not imply, however, that there are some propositions about the future whose truth values are *indeterminate*. Physical indeterminism leaves open the possibility that something besides the previous state of the world and the laws of physics determines the truth values of propositions about the future. For example, facts about the future could determine the truth values of propositions about the future. A physical indeterminist is not committed to the claim that any propositions about the future are indeterminate; *a fortiori*, she is not committed to the claim that the identities of future humans are indeterminate.

However, a group of philosophers whom I will refer to as “dynamists,” *are* committed to the claim that the identities of future humans are indeterminate. These dynamists, so-called because they endorse a dynamic theory of time, fall into two main camps. *Growing block*

de dicto duties to future humans. However, this would not help my present project, which is to establish that we have duties to future humans *de re*. For more on whether we can have *de dicto* duties to classes of individuals, whoever the particular individuals who comprise those classes might be, see Hare (2007).

15 Some people use “causal determinism” as a synonym for “physical determinism.” According to causal determinism, the complete state of the world at a time and the laws of nature *causally* determine the truth value of any proposition about the world at a later time. It is not clear what “causally determine” is supposed to mean, but if it means anything other than “determine,” then causal determinism is not synonymous with physical determinism. Further, if “causally determine” means *cause*, then according to the theory of causation I discuss in Chapter 5, causal determinism is false even if physical determinism is true.

theorists hold that while the past and present exist, the future does not.¹⁶ *Presentists* hold that only the present exists.¹⁷ Since concrete objects must exist at a time, and since nothing can exist at a non-existent time, growing block theory entails that the only concrete objects that exist are those that exist in the past or the present, and presentism denies that there are either past *or* future concrete objects. Thus, growing block theorists reject the claim that there are future space stations on Mars, and presentists reject the existence of such space stations as well as dinosaurs, Socrates, and New Coke. More pertinently, all dynamists are committed to the claim that future humans (where “future humans” is *not* used as a description) do not exist, and thus, that the identities of future humans are indeterminate.

As far as I can see, the *only* plausible motivation for the metaphysical indeterminacy objection to DV1 is a view like dynamism. But how, exactly, does the objection undermine my argument? One way to see the connection between dynamism, metaphysical indeterminacy, and DV1 is to consider the following argument, defended by Dennis Earl (2011)¹⁸:

- (E1) Future humans do not exist.
- (E2) In order to have moral duties to X, X must exist.
- (E3) Therefore, we do not have moral duties to future humans.

The argument is valid, and I accept premise E2. My account of what it is to have a *duty to* someone implies that the individual to whom the duty is owed must exist. However, Earl's conclusion contradicts DV1: if we do not have moral duties to future humans, then we do not

16 Defenders of the growing block theory include Michael Tooley (1997) and Peter Forrest (2004).

17 Defenders of presentism include C. D. Broad (1923) and Craig Boorne (2006).

18 Earl offers two formulations of the argument, one in terms of *obligations* to future generations, and one in terms of the *rights* of future generations. I have included just one of the formulations and substituted “humans” for “generations” and “duties” for “obligations,” but I do not think these substitutions distort Earl's argument.

owe pro tanto duties of nonmaleficence and beneficence to future humans. Either E1 is false, or else DV1 is.

Which claim should we reject? I believe that the best way to resolve the conflict between dynamism and DV1 is to determine which view coheres better with our other moral and metaphysical beliefs. In the remainder of this section I will argue that the view that *best* coheres with our other beliefs is DV1.

Let me start by giving credit where credit is due. Dynamists are motivated by a number of platitudes that do seem to be compelling features of commonsense metaphysics. One such intuition is that “the past is no more, while the future is yet to be.”¹⁹ Another is that our knowledge that we are in the present is as privileged as our knowledge that we exist.²⁰ Growing block theorists, on the other hand, are motivated by the view that while past and present are determinate, the future is open.

Many presentists and growing block theorists also hold that their view can vindicate the truth of claims that appear to refer to past and future concrete objects.²¹ Different versions of dynamism offer different accounts of the truthmakers for these claims. For example, on Prior's presentist view, the claim “Socrates taught Plato” is true in virtue of a primitive, concrete, present fact that does not contain either Socrates or Plato as its constituents. On the other hand, presentist Craig Bourne (2006) argues that the truthmaker for the claim “Socrates taught Plato,”

¹⁹ Sider 2003, p. 11.

²⁰ See Bourne 2006, p. 21.

²¹ However, not all dynamists worry about making claims about future concrete objects true. For example, Broad argues that the statement, “tomorrow it will be wet” is simply not true; see Broad (1923), p. 73.

is an abstract object that consists of ordered sets of propositions about present facts.²²

What is problematic about these views, however, is that in their attempt to vindicate commonsense platitudes such as “the future is yet to be,” presentists and growing block theorists commit themselves to claims that are even less plausible than the denial of the original platitudes. For example, Bourne's view implies that claims about the present and claims about the past have radically different truthmakers. Although “I am eating a cheese,” and “my great-grandmother was eating a cheese,” seem to assert structurally similar facts, Bourne's view implies that only the former is directly grounded by facts about me and a cheese. The latter is directly grounded by an abstract object and only indirectly grounded by various concrete objects, and those concrete objects might not even include a cheese.

Dynamists must also deny the truth of many of our moral judgments. Consider, for example, the following case:

Toxic Delicacy: Jones has been attempting to get pregnant. As a matter of fact, she has just engaged in an act of procreation that will result in a successful pregnancy, but the sperm that has made its way into her uterus has not yet borrowed into the egg. Jones is now dining at a restaurant, where she is offered a delicacy that contains a slightly toxic substance. She is told that the substance will not harm her at all, but if she should become pregnant in the next few weeks, it will cause her child to develop a painful and permanent health condition when the child turns eight. Jones shrugs and ingests the delicacy.

It seems clear that at the time she ingests the delicacy, Jones is doing something wrong, and that what grounds the wrongness of her action is that she is violating her duty to her future child.

²² Presumably, if dynamists can offer semantic accounts on which “Socrates taught Plato” or “Abraham Lincoln was taller than Napoleon Bonaparte” come out true, they can also offer an account that implies that “we have duties to future people” is true. However, vindicating the truth of “we have duties to future people” is not the same as vindicating DV. That is because DV includes a substantive account of the “duties to” relation, and this account implies that an individual who is owed a duty must exist. No presentist can hold both that humans who will exist in the future *exist* and that we have moral duties to them.

However, if a dynamist wants to hold that Jones has a duty to her future child, then she must give an account of what “her future child” refers to. The referent of “her future child” might, for example, be a concrete object consisting of the egg and sperm that will later become Jones's daughter, or it might be an abstract object such as a set. Whatever the dynamist decides Jones's child *is*, the claim that Jones has duties to *that thing* will be implausible unless *that thing* is an existing child, and the dynamist *cannot* say that Jones's future child is an existing child. On the other hand, the dynamist could deny that Jones has a duty to her future child, but this would also conflict with commonsense morality. In either case, then, the dynamist is committed to an implausible moral claim.

I think, then, that we should reject both presentism and the growing block theory. But if both theories are false, then the original metaphysical indeterminacy objection lacks motivation; if we do not deny that the future exists, then there seems to be little reason to deny that the identities of future humans are determinate.²³ Instead, we should conclude that, although it is false that future humans exist *now*, it is nonetheless true that they exist.

It might seem as though my response to the metaphysical indeterminacy objection begs the question. The objection, after all, is that future humans do not exist, and *therefore* we do not

²³ One who is sympathetic to my view might try to argue that even *if* the identities of future humans are indeterminate we can still have duties to them, but I do not think the strategy will succeed. Alastair Norcross (1990) proposes a thought experiment in which a canister of poison gas is positioned near the air vents of two rooms. In one room is Smith, and in the other is Jones. An indeterministic process will close the air vent to either Smith's room or Jones's room just before the canister releases its contents. If a third party, Fowl, intervenes to prevent the process from closing either of the vents, then there is no fact of the matter about who has been deprived of a future that was otherwise in prospect. Nevertheless, Norcross thinks it is clear that Fowl is responsible for “one more death than would otherwise have occurred” (p. 270). However, this case does not show that violating a direct duty to someone is compatible with the metaphysical indeterminacy of the individual's identity. According to the account of harm I provide in Chapter 4, Fowl harmed both Smith and Jones; thus, the identity of each of those to whom he violated a moral duty was determinate, after all. I will discuss this case further in Chapter 3.

have duties to them. I *appear* to have responded that we have duties to future humans, and *therefore* they exist. However, my response invokes more than just the premise that we have duties to future humans. I have appealed, for example, to the *particular* judgment that Jones has a duty to her future child. I have also argued that dynamism is in tension with too many of our other beliefs, some of which are not moral judgments, but metaphysical ones. For all these reasons, I conclude that the metaphysical indeterminacy objection is unsuccessful.

1.5. Conclusion

In this chapter, I have argued for DV1, which is the claim that individual moral agents who are now alive have pro tanto duties of nonmaleficence and beneficence to and regarding at least some future humans. In support of that claim, I argued that at least some future humans have moral standing and that their having moral standing is sufficient for our having pro tanto duties of nonmaleficence and beneficence to and regarding them. I then considered an number of objections and argued that they were unsuccessful.

Chapter 2: How We Can Violate Our Duties to Future Humans

2.0. Introduction

In the previous chapter, I argued for DV1, which is the claim that we owe at least some future humans *pro tanto* duties of nonmaleficence and beneficence. In this and the next four chapters, I will argue for DV2, which is the claim that in the present circumstances, the duty of nonmaleficence that we owe to future humans grounds reasons for all presently living moral agents to refrain from damaging certain features of the natural environment. My positive argument for DV2 depends upon the accounts of harming and causation I offer in Chapters 4 and 5, so my full defense of DV2 will not be complete until the end of Chapter 5. My specific aims in this chapter are to clarify DV2, to provide a brief sketch of how the argument for it will go, to motivate some of the premises, and to defend the argument against one preliminary objection.

The first clarification is that the phrase, “certain features of the natural environment” should be understood in the *de re* and not the *de dicto* sense. When I say that we have a reason not to damage “certain features of the natural environment,” I have two particular features in mind: the atmosphere and the planet's biodiversity. More precisely stated, DV2 is the claim that in the present circumstances, our duty of nonmaleficence to future humans grounds a reason for us to refrain from causing the release of carbon dioxide into the atmosphere and to refrain from causing a high number of species extinctions. I refer to these specific actions as “damaging certain features of the natural environment” only for ease of reference.

Here is a second point of clarification. Much has been written about the political dimensions of environmental degradation. There is an important question, for example, about

whether countries such as the United States or China should do more to prevent climate change. That question is largely about the principles of distributive justice: the atmosphere's capacity to absorb greenhouse gases is a limited resource, and the question is how our institutions ought to have distributed that resource among presently living and future generations and how they ought to rectify any injustices in the way it has so far been distributed. Since this dissertation is not about the principles of distributive justice, I will not engage in the debate about what the United States or China ought to do. However, I think that the important moral questions about the environment are not exhausted by questions about what countries should do or how natural resources should be distributed. There are also questions about what individuals owe each other, and one way of getting at these questions is to ask whether individuals can harm other individuals by damaging the environment. It may well be that in order to avoid harming future individuals, present individuals will need to establish political institutions that protect the environment; a failure to do so might result in both harm to future individuals *and* injustice. In this dissertation, however, my focus is only on the harm.

One final point of clarification has to do with the relationship between harming and benefiting. I believe that in *addition* to grounding reasons for us to refrain from damaging certain features of the environment, our duties to future humans *also* ground reasons for us to take various actions to protect the environment. More precisely, whereas the duty of nonmaleficence grounds reasons to refrain from damaging the environment in certain ways, the duty of beneficence grounds reasons to protect the environment in certain ways. However, I will not provide a full argument for the latter claim; my hope is that what I say about the reasons to

refrain from damaging the environment will show how the argument for reasons to protect the environment would go.

Here is how my argument for DV2 will go. For ease of reference, I will use the term “fh-reason” as shorthand for *a reason that is grounded in a pro tanto duty of nonmaleficence to and regarding future humans*. (I use “fh” as a mnemonic for “future harming.”) My argument will rely on two inferences of the following form:²⁴

- (R1) We have a pro tanto duty to some future humans not to harm them.
- (R2) In the present circumstances, we can harm those humans by ψ -ing.
- (R3) Therefore, in the present circumstances, there is an fh-reason for us not to ψ .

For each inference, I will substitute one of the following specific actions for ψ -ing: (1) causing carbon dioxide to be released into the atmosphere; and (2) causing a high number of species extinctions. In the next two sections, I will present some considerations about the particular circumstances we are in that provide some intuitive support for each version of R2. I will then consider and respond to one preliminary objection to R2. In the next chapter, I introduce a second objection to R2, and I will use that objection to motivate my accounts of harming and causation.

2.1. The Harm of Carbon Dioxide

In the time it takes me to type this sentence, human activities will have caused approximately 2,000 metric tons of carbon dioxide to be released into air.²⁵ Since the Industrial Revolution,

²⁴ Recall from the introduction that I am assuming that the form of this argument is metaphysically valid.

²⁵ I used the figure of 31.6 gigatonnes a year divided by 31,622,400 seconds in a year, multiplied by 2 seconds. The figure of 31.6 is from International Energy Agency (2012).

humans have been by burning fossil fuels on a massive scale, and the gases released when the fuels are burned—most notably, carbon dioxide—have been changing the composition of the planet's atmosphere. By some estimates, the concentration of carbon dioxide in the atmosphere is now higher than it has been since the Middle Miocene era, about 15 million years ago.²⁶

This ongoing, rapid rise in atmospheric carbon dioxide is affecting the planet's weather. According to the IPCC, an intergovernmental panel appointed to assess the relevant science, high levels of atmospheric carbon dioxide intensify the greenhouse effect, a process in which the heat reflected by the Earth's surface is absorbed by the atmosphere and then redirected towards the Earth's surface. Because the Earth's surface temperature is rising, the planet's weather systems are gathering energy. This extra energy is fueling *climate change*, or a change in the statistical distribution of weather patterns.²⁷

The carbon dioxide and other greenhouse gases that we released into the atmosphere over the past 200 years have already caused some changes in weather patterns. For example, the IPCC claims that the late 20th century saw an increased frequency of heat waves, an increase in the areas affected by drought, and an increase in intense tropical cyclone activity in the North Atlantic. However, if we continue to add carbon dioxide to the atmosphere, we can expect to see even further increases in the frequency and intensity of heat waves, droughts, floods, fires, and tropical cyclones. We can also expect a decrease in sea ice that will lead to a higher sea level. Many islands and densely populated coastal areas such as Bangladesh, Tokyo, and New York

26 Wolpert (2009).

27 IPCC (2007).

City could be flooded and eventually made uninhabitable.²⁸

Changes in the climate also have secondary effects. For example, water shortages and extreme weather events reduce agricultural productivity and contribute to starvation and malnutrition. Changes in weather also affect the latitudes at which mosquitoes, ticks, and other vectors of infectious disease can survive. Already, the World Health Organization estimates that disease and malnutrition due to climate change were responsible for 0.2 percent of the deaths that occurred in 2004.²⁹

Commonsense morality holds that forced migrations, property damage, disease, starvation, and death are all harms to human beings. What should we say, then, about a state of affairs that raises the likelihood of such harms? I believe we should say that, because it raises the likelihood of serious harms, an excessive amount of carbon dioxide in the atmosphere is, itself, a harm. More precisely, given that there is already too much carbon dioxide in the atmosphere, any further rise in carbon dioxide levels is a harm. When we cause carbon dioxide levels to rise, we are causing that harm. That is why, in the present circumstances, causing carbon dioxide to be released into the atmosphere is a way of harming future humans.

Nevertheless, this description of our behavior faces some objections. One of these is related to what is known as the “non-identity problem,” and I will address that objection in the next chapter. The other objection is based upon the assumption that we can negate any potential harm to future humans by improving their standard of living at the same time that we damage the

28 IPCC (2007).

29 WHO (2009).

environment. I will respond to this other objection soon, after first considering a second example of a way in which we can harm future humans.

2.2. The Harm of Species Extinction

We are living in an era that scientists have termed, the “Holocene Extinction.” The era began about 13,000 years ago when humans expanded around the globe and woolly mammoths and other large mammals died off. Since that time, more and more species have vanished. According to the International Union for the Conservation of Nature, of the 59,507 existing species that scientists have studied, 19,265 are now threatened with extinction.³⁰

Much of this loss of biodiversity is due to human activity. Humans have destroyed the habitats of numerous species in order to create farms, pastures, and cities. We have also caused species to go extinct by fishing, hunting, introducing non-native species to new regions, spreading diseases, and polluting the air and waterways. Our contributions to climate change have caused and will continue to cause extinctions.

Many environmentalists hold that this destruction of biodiversity is wrong because it violates duties that we have to whole species or ecosystems, or perhaps to various nonhuman individuals within ecosystems. I do not take a stand on whether we can have a duty to a species or to particular plants or animals. However, the loss of biodiversity also has consequences for future humans. When high numbers of species are destroyed, the ecosystems they were part of are weakened or destroyed. With fewer ecosystems in existence to purify the air and water,

³⁰ International Union for the Conservation of Nature (2013).

renew soil fertility, replenish genetic diversity in agriculture, and stabilize the climate, human health suffers. Further, biodiversity contributes to the development and production of new foods, fibers, petroleum substitutes, and medicines.³¹ It has become a standard practice for medical researchers to gather samples of unfamiliar, wild species and test them systematically for their potential to treat cancer and other diseases. Edward Wilson writes,

Revolutionary new drugs have rarely been developed by the pure insights of molecular and cellular biology Rather, the pathway of discover has usually been the reverse: the presence of the drug is first detected in whole organisms, and the nature of its activity is subsequently tracked down to the molecular and cellular levels. Then the basic research begins.³²

Many species are going extinct before researchers can test them for their medicinal potential. It is likely that because of these extinctions, at least some future people will not have medicines and other valuable compounds that they could have otherwise had.

Finally, there is some reason to think that reduced access to ecosystems directly contributes to increased incidences of anxiety, depression, and physical illness. In a study that supports this claim, people who lived more than a kilometer from a park or wooded area experienced significantly greater incidences of 15 out of 24 diseases including heart disease, lung disease, neurological disease, anxiety, and depression.³³ Other studies have shown that people recover from illnesses more quickly when they are in hospital rooms with windows than without, and that in areas with lots of green space, there are fewer health inequalities between rich and

31 World Wildlife Fund (2013).

32 Wilson (2002), p. 120.

33 Maas et al. (2009).

poor people.³⁴ If it turns out that access to thriving ecosystems is important for our mental and physical health, then causing high numbers of species to go extinct may result in poorer health among future generations.

Does it follow that, by causing causing high numbers of living species to become extinct, we are harming future people? Commonsense morality holds that anxiety, depression, poor health, and the destruction of valuable resources are harms to people. Since a significant lack of biodiversity is a lack of a valuable resource and might be a source of anxiety, depression, and poor health, I believe we should say that a significant lack of biodiversity is, itself, a harm. Since new species are not evolving quickly enough to replace our destruction of them, our causing high numbers of species to go extinct is, in the present circumstances, a way of causing harm to future people. Nevertheless, just as there are objections to my previous claim about carbon dioxide, so there are objections my claim that causing high numbers of species to go extinct is a way of harming future humans. I will address the first of these objections in the next section.

2.3. The Fungibility Objection

The first objection to the claim that damaging the environment in the ways described above is a way of harming future humans appeals to the fungibility of natural resources. The objection grants that both the diversity of species on the planet and the climate's limited capacity for absorbing carbon dioxide while maintaining equilibrium are valuable resources.

Nevertheless, there are many other resources that are also valuable. Such resources include

³⁴ Gardner (2009).

money, knowledge, technology, political stability, and large populations of human beings. The very actions that cause carbon dioxide to be released into the atmosphere also generate the above sorts of resources, and so do the actions that lead to species extinctions. Indeed, future humans will be *as well off or better off* living in a world with a high GDP, political stability, advanced technology, and so on than they would have been living in a world whose climate had not changed and whose ecosystems were more intact. The objection concludes that higher levels of atmospheric carbon dioxide and high numbers of species extinctions are not harms to future humans, at all; thus, in causing these things, we do not harm future humans.

There is some intuitive plausibility to this objection. Suppose that, as some scientists believe, it was humans who caused high numbers of the planet's megafauna, including woolly mammoths, mastodons, giant sloths, sabertooth tigers, and so on to go extinct about 13,000 years ago. It is likely that the human activities that would have caused these extinctions (namely, over hunting and altering ecosystems in a way that caused a mini episode of climate change) also enabled humans to expand in population size and develop new technologies. Indeed, the technologies these humans developed led directly to the high standard of living that many people enjoy today. Possibly, the average person today is *better off*, not worse off, as a result of what early humans may have done, both to cause a miniature episode of climate change and to cause the extinction of megafauna. Thus, even if humans caused an episode of climate change and numerous species extinctions, they did not harm us. Likewise, if we cause another episode of climate change and numerous more species extinctions, we will not harm future humans.

Or so it might seem. However, this objection relies upon two questionable assumptions.

One is a conceptual assumption about the nature of harm. The other is an empirical assumption about the net effects of the present processes of climate change and species extinction. The conceptual assumption is that an individual is not harmed unless she is worse off *on the whole*, compared to some alternative. The empirical assumption is that at most, carbon dioxide in the atmosphere and loss of biodiversity will make future humans worse off (compared to some unspecified alternative) in *some respect*, but not worse off *on the whole*.

Let me respond to the empirical assumption first. It may well be that the actions that damage the environment also generate money, new technologies, and other things we find valuable. Perhaps the money and technology can mitigate some of the bad effects of climate change and species extinctions. For example, the money and technology might lead to an easy cure for malaria. The fact that climate change caused an increase in cases of malaria would then not be much of a problem at all. People who caught malaria would be worse off in one respect—namely, insofar as they caught malaria—but better off, on the whole, for living in a wealthy world, especially since their malaria would be easily cured. Or perhaps we might develop new treatments for depression and anxiety. Humans who were depressed or made anxious by a lack of exposure to healthy ecosystems could then undergo such treatments and be as mentally healthy or healthier than they would have been if ecosystems had remained intact.

However, these empirical speculations are problematic for at least two reasons. First, there may be a level of biodiversity that is necessary for human flourishing, such that any drop below that level cannot be made up for by any increase in wealth or technology. As an analogy, consider the various nutrients that are essential in our diet. If we lack a certain amount of, say,

Vitamin C in our diets, then our health will suffer, no matter how much Vitamin D we get, no matter how many antidepressants we take, and so on. Likewise, if we do not have enough different animals and plants in the world, then we may only be able to reach a certain level of well-being, no matter how many new antidepressants we develop. It might be false, then, that the basic goods that are necessary for human flourishing are fungible.

Second, there is no guarantee that the money and technology generated by our destruction of the environment will get into the hands of those future people who will be affected by the destruction. For example, in just seven years, the population of Bangladesh is expected to increase by 14 million, while, as a result of climate change, agricultural yields in Bangladesh are expected to drop by 50 percent. We can be virtually certain that the wealth generated from burning fossil fuels will not compensate these 14 million future humans for the shortage of food they will face. Even if other people in the world are better off as a result of the actions that caused climate change, these future individuals in Bangladesh will not be; instead, they will have been harmed.

While the first of my responses to the fungibility objection may not be decisive, I believe that the second one is. My claim, after all, is that certain environmentally destructive actions can harm at least *some* future humans; I do not need to argue that our activities will harm *all* future humans. Undoubtedly, causing carbon dioxide to be released into the atmosphere will harm *at least* a few million Bangladeshis who will be born in the next few decades, for they will face a serious risk of starvation. The fungibility objection therefore relies on the false empirical claim that *no one* affected by our environmentally destructive behavior will be worse off *on the whole*.

The fungibility objection also relies on a questionable conceptual assumption. The assumption, again, is that an individual is not harmed unless she is worse off *on the whole*, compared to some alternative. According to the accounts of harming and causation that I will advance in the next few chapters, this assumption actually turns out to be *true*, but only when we are appealing to the right alternative. I will reject the claim that an action harms an individual only if it makes her worse off on the whole than she would have been, had the action not been taken. I will also reject the claim that an action harms an individual only if it makes her worse off on the whole than she was before the action. Instead, an individual may be harmed if she is merely worse off *in some respect* than she was or would have been, either before the action or had the action not been performed.

A full defense of this claim must wait until I provide my account of harming. However, as some preliminary motivation for the claim, consider Seana Shiffrin's (1999) example of a wealthy character who decides to share some of his wealth with other people. His method of sharing the wealth is to create cubes of gold bullion, each worth five million dollars. The wealthy character climbs into his plane and flies over people, dropping the gold bullion cubes on them. A cube falls on one individual whose arm is consequently broken. Shiffrin argues that, although the five million dollars more than covers the cost of getting the arm repaired, the person whose arm was broken was nevertheless harmed by the wealthy character's action. Notice that the victim of the broken arm is better off *on the whole* than he was before the wealthy character dropped the brick on him. He is also better off *on the whole* than he would have been, had the wealthy character not dropped the brick on him. Nevertheless, he is worse off in *one respect*. Namely, he

got a broken arm.

Now compare the victim of the falling brick to the victims of the loss of biodiversity and climate change. Even if the victims of the latter are better off *on the whole* than they would have been, had we not destroyed certain features of the environment, this does not mean that they are not harmed by our destruction of the environment. There is some respect in which they are worse off: namely, they live in a more unstable climate, and they live in a world depleted of biodiversity. Likewise, we can perhaps claim that *we*, ourselves, are worse off in one respect, as a result of what the humans who lived 13,000 years ago may have done. We might be worse off insofar as there are no woolly mammoths around for us today to study and enjoy.

2.4. Conclusion.

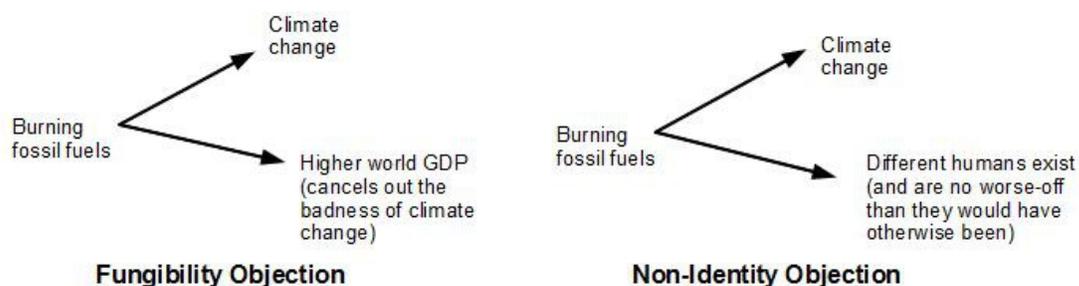
In this chapter I introduced my argument for the second tenet of DV, which is the claim that in the present circumstances, our duty of nonmaleficence to some future humans grounds reasons for all presently living moral agents to refrain causing carbon dioxide to be released into the atmosphere and to refrain from causing high numbers of species extinctions. I provided some empirical information that suggests that high levels of carbon dioxide in the atmosphere and a profound loss of biodiversity are both harms to at least some future humans. If harming is causing harm, then causing these things is a way of harming future humans. Since we have a duty to refrain from harming some future humans, there is an fh-reason for us, in the present circumstances, to refrain from damaging certain features of the environment in the aforementioned ways

Chapter 3: The Non-Identity Problem

3.0. Introduction

Recall that I have defined an “fh-reason” as a reason that is grounded in a pro tanto duty of nonmaleficence to and regarding future humans. In the previous chapter I considered what I called the “fungibility objection” to my claim that in the present circumstances, there is an fh-reason for us to refrain from damaging certain features of the environment. Although I did not put it in precisely these words, one way of thinking about the fungibility objection is in terms of *double effects*: according to the objection, the actions that cause carbon dioxide emissions or species extinctions have *other effects* as well. Since some of those other effects are *good* for future humans, the good cancels out the bad, and the aggregated effects of the actions do not qualify as harms.

In this chapter, I will focus on another objection that can also be formulated in terms of double effects. I will refer to this second objection as the *non-identity objection*. Like the fungibility objection, the non-identity objection appeals to consequences that are distinct from either an increase in atmospheric carbon dioxide levels or high numbers of species extinctions. However, whereas the additional effects on which the fungibility objection focused were the *creation of wealth and technology*, the additional effects at issue in the non-identity objection are the *creation of future humans*. The following figure illustrates the objections' relationship:



The non-identity objection is motivated by one family of solutions to what Derek Parfit calls the “non-identity problem.” In the next section, I will provide an overview of the non-identity problem and I will explain how a particular approach to the problem—one that endorses what I will call a *comparative account of harming*—gives rise to the non-identity objection. In section 3.2 I will consider some specific proposals that endorse an implication of the comparative account of harming that I call the *counterfactually worse-off condition on harming*, and I will argue that none of the proposals support successful solutions to the non-identity problem. In section 3.3 I will endorse a solution to the non-identity problem that rejects the counterfactually worse-off condition on harming, and I then will consider and respond to an objection to rejecting the counterfactually worse-off condition.

3.1. The Non-Identity Problem

The term, “non-identity problem” was coined by Derek Parfit (1987), but the fundamental tensions that constitute the problem were identified, somewhat independently, by Parfit (1976), Robert Adams (1979), Thomas Schwartz (1978), and Gregory Kavka (1982). Parfit sets up his description of the problem with the following case:

The 14-Year-Old Girl. This girl chooses to have a child. Because she is so young, she gives her child a bad start in life. Though this will have bad effects throughout this child's life, his life will, predictably, be worth living. If this girl had waited for several years, she would have had a different child, to whom she would have given a better start in life (Parfit, 1987, p. 358).

In order to support his contention that, if the girl had waited, she would have had a different child, Parfit argues that one's identity is dependent upon the particular sperm and egg pair from

which one originated. The matter of which particular sperm will fertilize which particular egg usually depends, in turn, upon the time at which a couple engages in intercourse. An individual who was conceived at a some particular time will not be identical to the individual who would have been conceived, if intercourse had taken place sufficiently earlier or later.

How large is the window of time in which an individual's conception must occur in order for *that* individual, and not some other individual, to come into existence? Parfit imagines that on some views, switching out one spermatozoon for another but holding the identity of the ovum constant may not result in a determinate change in who is conceived. Nevertheless, he thinks that, at the very least, a change in the identity of both the spermatozoon *and* the ovum will mean that a different individual comes into existence. Parfit therefore builds a month-long window into his formulation of what he calls the “time dependence claim”:

If any particular person had not been conceived within a month of the time when he was in fact conceived, he would in fact never have existed (p. 352).

Recall that the 14-year-old girl's action—namely, having a baby at the age of 14—sets her child up for a bad start in life. Parfit claims that intuitively, her action *wrongs* her child. However, if wronging an individual is a sufficient condition for acting wrongly, I believe that this description of the relevant intuition is too strong. We might not think that the girl acted wrongly, all things considered, but there is still a sense that one of the moral reasons against what she did had something to do with her child. To use my own terminology, I think the intuition most people would have about this case is that there is, at least, an fh-reason against the 14-year-old girl's action. I will assume that the claim that there is an fh-reason against such an action is

equivalent to the claim that the girl's action *in some way* wrongs the child, where the qualifier “in some way” indicates that the action might or might not be wrong, all things considered.

Now, the time dependence claim implies that if the 14-year-old girl had not acted as she did, then the particular child to whom she gave birth would not have existed at all. Since that child has a life worth living, he appears to be no worse off in any respect than he would have been, had his mother not acted as she did; if she had not done so, he would not have *been* at all. Since the child is no worse off in any respect than he would otherwise have been, it seems that he has not been harmed by his mother's action. But if he has not been harmed, then it is difficult to explain how he could have been in some way wronged.

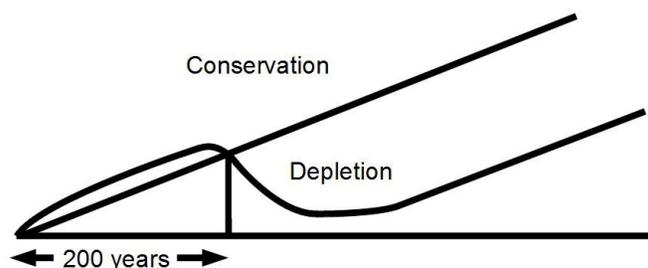
The 14-Year-Old Girl is thus a case in which we find it difficult to explain why an individual *appears* to be in some way wronged by an action that is the condition of his or her own worthwhile existence. I shall refer to this kind of case as a *non-identity case*. The problem of explaining why the individual in such a case appears to have been in some way wronged is the *non-identity problem*.³⁵

35 I do not claim that this is the *standard* definition of the non-identity problem. My definition has two features that are at odds with some other accounts in the literature. First, my definition implies that the non-identity problem arises in *any* case in which there appears to be an *fh*-reason against bringing an individual into existence. This distinguishes my definition from those that take the non-identity problem to arise *only* in cases in which it appears to be wrong to bring one individual into existence instead of *some other individual*. Second, I hold that the intuition that gives rise to the non-identity problem is the intuition that there is a *reason* against the procreative action, and not the intuition that the action is *wrong*, all things considered. This distinguishes my definition from Harman's (2009) definition; she characterizes the non-identity problem as one of determining why people agree that the procreative action is *wrong* in some identity cases but disagree about whether the procreative action is wrong in other cases. She writes, “A ‘non-identity case’ is a case in which an action affects whether some people exist, and it appears (either to everyone, or to some people) that the action is wrong in virtue of harming these people, although they have lives that are worth living” (2009, p. 140). Parfit, himself, provides two different definitions of “non-identity problem” that are somewhat imprecise. He first claims that the non-identity problem is the problem of determining what is the “objection” to the 14-year-old girl's decision (p. 359). He then invokes the case he calls “Depletion” and states that the non-identity problem is the problem of answering the question, “What is the moral reason not to choose Depletion?” (p. 363).

The action of *having a particular baby* is not the only kind of action that can give rise to the non-identity problem. Because of the time-dependence claim, acting in a way that results in *other* people's having babies at different times can also give rise to the problem. So can changing the conditions in which people live and work; our identities depend not only on *when* our parents conceived us, but on *who* our parents were. For example, if my father had worked in a different industry, he may not have met my mother, and if he hadn't met my mother, then I would not have been born, no matter *when* my father procreated. To illustrate the sensitivity of our identities to the conditions in which we live and work, Parfit envisions a scenario in which a community enacts a policy that changes people's living and working conditions in such a way that some individuals meet different individuals, procreate at different times, and conceive different children than they would have conceived, had the community not enacted the policy. The children of these individuals have children of their own, and eventually the last of those who would have existed regardless of whether the policy had been enacted die off. Two hundred years after the policy has been enacted, *no one* exists who would have existed, had the policy not been enacted. Parfit then adds a twist to scenario; instead of just determining their identities, the policy also results in a lower quality of life. Here is how Parfit describes the case:

Depletion. As a community, we must choose whether to deplete or conserve certain kinds of resources. If we choose Depletion, the quality of life over the next two centuries would be slightly higher than it would have been if we had chosen Conservation. But it would later, for many centuries, be much lower than it would have been if we had chosen Conservation. This would be because, at the start of this period, people would have to find alternatives for the resources we had depleted (p. 361 – 62).

Parfit illustrates the choice of policies with the following graph:³⁶



Let us call any individual who comes into existence at least 200 years after we have enacted the Depletion policy a “depleted person.” Parfit stipulates that, due to the factors I cited above, no depleted person would have existed if we had decided not to deplete the resources. Depletion is thus a non-identity case: it *appears* as though each depleted person is in some way wronged, but none of the depleted persons is worse off in any respect than he or she would have been, had we not depleted the resources. Had we not done so, *those* depleted persons would not have existed. Just as we had trouble explaining why there appears to be an fh-reason against the 14-year-old girl's action, so we have trouble explaining why there appears to be an fh-reason against our choice in Depletion.

Notice the similarities between Depletion and the situation we are actually in. I argued in Chapter 2 that by depleting the climate's capacity to tolerate excess carbon dioxide, we are harming future humans, and also that by depleting biodiversity, we are harming future humans.

³⁶ Parfit uses two separate graphs that correspond with choices between what he calls “Greater Depletion” and “Lesser Depletion,” but the difference between the two types of depletion do not matter for my present purposes, so I have pictured only the graph that corresponds to “Greater Depletion.”

However, if we failed to deplete either of these resources, then, as Parfit's reasoning suggests, a different set of humans would come into existence. Technically, of course, Parfit supposes that it takes at least 200 years for the set of humans to be *completely* different from what it otherwise would have been. If this is true, then for all that Parfit has said, there might be a period of 200 years when the environmentally destructive actions we take today harm people who would have existed *regardless* of whether we had taken those actions. Nevertheless, I intend for the claims I asserted in Chapter 2 to be stronger than this; I believe that our environmentally destructive actions harm those who will live, for example, 400 years from now. Moreover, 200 years might be a generous estimate of how long it takes for the repercussions of certain actions to determine the identities of everyone in a society. For these reasons, I will simply assume from this point on, unless I indicate otherwise, that *every* action that adds carbon dioxide to the atmosphere or causes high numbers of species extinctions determines the identity of *every* future human who will be affected by that action. Such an assumption will clarify the explanation I will give below about why the non-identity problem is a problem for DV2.

Before I present that explanation, one more difference between Depletion and the situation we are actually in is worth mentioning. In Parfit's graph,³⁷ quality of life is at its *lowest* in the present. Certainly, there is a drastic downturn in the quality of life for those who live about 400 years in the future, but even at the lowest point of the curve, people have a better quality of life than those who are presently living. Parfit's assumption all future people will be better off than we are (at least for the next thousand years or so) is not an essential feature of the non-

37 Although I included only one graph, this is true in both of Parfit's graphs.

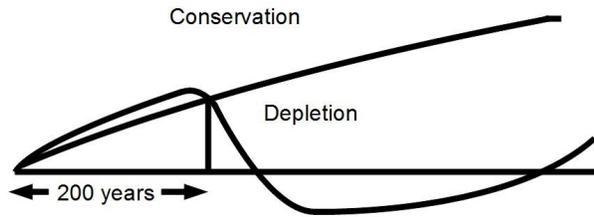
identity problem, but it is relevant to other questions about what we owe future humans. For example, the assumption that the average future human will be better off than the average presently living human may have implications for how we ought to distribute resources to future humans.³⁸ Although I do not take up questions of distributive justice in this dissertation, I still wish raise some doubts about Parfit's assumption. First, rates of depression have been on the rise for the past 50 years. In many countries around the world, a person born today is about three times more likely to suffer from major depressive disorder than a person born before 1955. If we use mental health as a proxy for well-being and rates of depression as a proxy for mental health, then we have some reason to challenge Parfit's assumption that well-being will continue to improve over time. Alternatively, if we use wealth as a proxy for well-being, there are other reasons to challenge Parfit's assumption. Although the world economy has been growing by tens to hundreds of billions of dollars each year for almost two centuries, that growth has been fueled by institutions and practices whose very existence is threatened by climate change.³⁹ It may well be, then, that the following graph better represents our current situation:

38 For example, Gollier (2007) writes, “By investing in technologies to reduce the impact of climate change in the distant future, we redistribute wealth from the poor current generations to the wealthy future ones” (p. 5). Lomborg (2001) makes a similar claim, writing,

We know that future generations will have more money to spend. Because of growth we are actually the poor generation, and future generations will be richer than us. We expect that in 2035 the average American will be twice as rich as she is now. For this reason it is perhaps not entirely unreasonable that society expects richer, later generations to pay more towards the costs of global warming—in exactly the same way as high income groups in our society pay higher taxes (p. 314, cited in Gardiner 2011).

If we think that justice does not require the poor to distribute resources to the rich, these kinds of consideration would undermine justice-based reasons for us to take costly actions to reduce the impact of climate change.

39 See Gardiner (2011), argues that climate change poses a threat to our economic institutions, and that we are therefore not justified in believing that future generations will be wealthier than our own.



As I mentioned, however, although the discrepancy between graphs might matter to questions of distributive justice, it does not matter to either the non-identity problem or the question of whether we can harm future humans. Just as the question of whether I have harmed you is not settled by whether I have made you worse off than I am, the question of whether we can harm future people is not settled by whether we can make them worse off than *we* are.

Now that I have explained why the situation we are currently in is very much like Parfit's Depletion case, I will explain how the non-identity problem motivates what I above referred to as the “non-identity objection.” First, one approach to solving the non-identity problem is to accept what I call the *comparative account* of harming. Proponents of the comparative account endorse the following:

(C1) An action or event, A, harms an individual, S, if and only if A causes a state of affairs, W, to obtain, such that W is a harm for S with respect to A.

They also endorse at least one of the following:

(C2) A state of affairs, W, is a harm for an individual, S, with respect to an action or event A if and only if W is worse for S in some respect than some other state, X, such that X obtained immediately before A.

(C3) A state of affairs, W, is a harm for an individual, S, with respect to an event or action A if and only if W is worse for S in some respect than some other state, X, such that X would have obtained if A hadn't happened or been performed.

The combination of C1 and C3 entails the following *counterfactually worse-off condition on harming*: An event (or action) harms an individual, S, only if it makes S worse-off in some respect than S would be, had the event not occurred (or action not been performed). Some philosophers who write about the non-identity problem endorse the counterfactually worse-off condition on harming but do not invoke the complete comparative account of harming. I nevertheless take the comparative account to be their underlying motivation. I will discuss some problems with the comparative account in Chapter 5. For now, I mean to simply introduce the account so that its relationship to the counterfactually worse-off condition on harming and its contrast with the non-comparative account of harming, which I will later endorse, will be more apparent.

The non-identity objection appeals to the counterfactually worse-off condition on harming and can be formulated as follows:

- (N1) The counterfactually worse-off condition on harming is true.
- (N2) If the counterfactually worse-off condition on harming is true, then our damaging certain features of the natural environment does not harm future people.
- (N3) If our damaging certain features of the natural environment does not harm future people, then DV2 is false.
- (N4) Therefore, DV2 is false.

Here is the reasoning behind the premises. N3 is true because DV2 is the claim that we have an fh-reason to refrain from damaging certain features of the natural environment. There would be such a reason only if our damaging those features harmed future people. N2 is also true on the

assumption that in every case in which we damage certain features of the natural environment, we determine the identities of all future people who will be affected by that damage. Given this assumption, the future people are no worse off than they would have been, had we not damaged those features of the natural environment.

I hold that the false premise in the non-identity objection is N1. In the next section, I will consider three specific proposals that attempt to solve the non-identity problem by appealing to or at least implicitly endorsing something like the counterfactually worse-off condition on harming, and I will argue that none of the three offers a successful solution to the problem.

3.2. Some Attempts to Solve the Problem

The first unsuccessful approach to the non-identity problem is the “biting-the-bullet” approach taken by David Heyd (2009) and David Boonin (2008). Both authors accept the inference that begins with the claim that individuals in non-identity cases are no worse off than they would have been, had the agents who brought them into existence acted differently, and ends with the claim that such individuals have not been in any way wronged.

Any biting-the-bullet approach that attempts to defend what initially strikes one's opponents as a counterintuitive conclusion will face two difficult challenges: the first is to make a convincing case that the counterintuitive conclusion is *more* plausible than any other solution to the problem. The second challenge is to “explain away” the intuitions that appeared to support some other solution. Heyd attempts to argue both that his solution to the problem is the “lesser theoretical evil” and that the intuitions that appear to go against his view can be accounted for.

Nevertheless, the argument that Heyd's solution is the “lesser theoretical evil” becomes much less convincing as soon as an alternative theory—one that is even *less* evil, so to speak—is available. I think that the solution I will defend presently is better, primarily because it accounts for the intuition that there is an *fh*-reason against having a child in non-identity circumstances. Moreover, Heyd's attempt to “explain away” the intuitions that support the view that there is such a moral reason is unpersuasive. Heyd writes,

What counts in our moral judgment [of actions in non-identity cases] is the parents' moral profile, the kind of people they are, rather than the objective condition of the child (or world). And it is of course also true that if parents are completely indifferent to the welfare of their planned future child, they are liable to become bad parents and to violate their parental duties *to* the child once she is born (p. 16).

In other words, to account for the intuition that, for example, the 14-year-old girl's action is in some way wrong, Heyd attempts to locate the wrongness of the action in the girl's attitude toward the welfare of her future child. But if this account is intended to *fully* account for our intuitions about non-identity cases, it is unsatisfactory. After all, if the 14-year-old girl's character or attitude towards her child is in some way defective, then there must be an explanation of *why* her attitude or character is defective. But the only plausible explanation is that the girl's *action* fails to treat her future child with due care and respect. Therefore, Heyd's attempt to account for the intuition that non-identity cases are morally suspect will ultimately have to rely upon a claim that the actors in those cases are in some way *wronging* their children. But if Heyd accepts the latter claim, then he will be abandoning the biting-the-bullet approach.

David Boonin (2008) attempts a different strategy for “explaining away” the intuition that the individual in a non-identity case is wronged. Appealing to his own version of a non-identity

case in which a character named Wilma brings her son into existence and also appears to be responsible for the child's blindness, he writes,

At a purely intellectual level, we understand that Wilma's act harms no one, but our moral intuitions are not sufficiently fine tuned to respond at every instance to what, at a purely intellectual level, we understand to be the case. As a result, because the Wilma case so closely resembles other cases that do involve causing significant harm to others, and since the features of the case that account for the fact that Wilma does not cause any harm at all in this instance are difficult to keep prominent in our mind's eye as we reflect on the case, we still feel the force of the intuition that her act is wrong because we feel the force of the general moral belief that causing significant harm to other people is wrong (p. 146).

To explain why we have a the non-veridical intuition that Wilma is causing significant harm, Boonin argues that we are simply confusing the technically true claim that Wilma makes it the case that “her son,” *de dicto* is blind with the false claim that Wilma is causing blindness in the particular individual who is her son. Boonin then offers two cases that he thinks are less likely to lead our intuitions astray because they do not involve a similar *de re/de dicto* confusion. He holds that these two cases are still analogous to non-identity cases in all morally relevant ways. Since both of Boonin's cases are structurally similar, I will consider only the second case:

Barney: Barney is walking past a lake in which two young women are drowning. One is near him, the other is further away. Barney has one life preserver. He can toss the life preserver to the woman who is near him, but by the time she is safely ashore, it will be too late to save the further away woman. He can run down the shore and toss the life preserver to the further away woman, but by the time that she is safely ashore, it will be too late to save the first woman. So Barney can only save one of the two women. . . . Barney also knows the following facts about the women: each of them has just had sexual intercourse and conception will occur in each of them in approximately an hour, if they do not die first. One of the women will conceive an incurably blind child. The other will conceive a child who is not blind. The woman who will shortly conceive a blind child is the one who is closer to Barney. Barney decides to save the woman who will shortly conceive a blind child because this is more convenient for him (p. 151).

Boonin argues that in this case, we are not likely to be confused about the *de re/de dicto*

distinction. Moreover, in this case it is clear that Barney's action is not morally wrong. Since Barney's action is not wrong, then Wilma's action is not wrong, either.

I agree with Boonin that Barney's action is not wrong. Or to put things in my own terminology, I agree that Barney's action does not wrong the blind child *in some way*. Nevertheless, I think there is a morally relevant difference between Barney and a standard non-identity case. In Barney, Barney does not *cause* the child who exists as a result of his action to be blind, even though his saving the child's mother is a condition of the child's being blind. However, in Boonin's own example involving Wilma, Wilma really *does* cause her son to be blind. Likewise, in 14-year-old girl, the girl *causes* her child to have a bad start in life. The account of causation that I endorse in Chapter 5 will support the intuitive appeal of these causal judgments. In any case, I take the causal differences between the cases to be morally relevant, since causing harm seems to be more serious a matter than allowing harm to occur. If so, then Boonin's cases do not support his claim that our intuitions about non-identity cases are the product of confusion. I think a better explanation for our intuitions about non-identity cases is that they are *true*.

Advancing an alternative approach to the non-identity problem, Caspar Hare (2007) agrees with Boonin and Heyd that no determinate individual is harmed in non-identity cases. Nevertheless, he argues that in such cases, the agent's procreative action still harms her child in a *de dicto*, rather than a *de re* sense. To harm an individual in a *de re* sense is to harm *that particular* individual, regardless of how the individual is described; to harm an individual in a *de dicto* sense is to harm *whatever individual it may be* who happens to fulfill a certain description.

Hare draws the inspiration for this approach from a joke about Za Za Gabor. In the joke, Gabor claims that she has found a way to keep her husband “young and healthy, almost forever.” Her interviewer asks for the secret, and Gabor replies, “I get a new one every five years!” (Hare, p. 514).

Hare argues that in the sense that Gabor succeeded in keeping “her husband” healthy, a prospective parent has a duty to ensure that “her child” is also healthy. Sometimes, fulfilling this duty requires that the parent refrain from conceiving one child, at one time, in order to have *a different child* later, who will be healthier than the child who would have been born earlier. If a parental duty is fulfilled this way, the parent ensures that “her child” (*whoever* that might be) is healthy, even if the parent does little to improve the health of any particular child who happens to fit the description, “her child.” If the parental duty is *not* fulfilled—and if, in fact, a parent conceives an unhealthy child when she could have conceived a healthy one—then she has *de dicto* harmed “her child” by making “her child” worse off than “her child” would otherwise have been.

The problem with this account is that when we consider cases like The 14-Year-Old Girl, the intuition we have that she has in some way wronged her child does not seem to be explained by the intuition that she has in some way wronged “her child,” *de dicto*. To see this, suppose that we name the girl's child Billy and then explain the case again: “The 14-year-old girl wanted to have a baby, so she got pregnant and gave birth to Billy, and Billy had a bad start in life.” When the case is explained this way, we continue to have an intuition there was an *fh*-reason against the 14-year-old girl's action, such that this reason was grounded in a duty to *Billy*. The intuition is not

about a *phrase*; it is about an individual. If this is so, then like Boonin and Heyd, Hare has failed to “explain away” the intuition that tells against his preferred solution to the non-identity problem.

A third approach to the non-identity problem is to question the assumptions that are built into non-identity cases. Recall that in Parfit's 14-year-old girl case, the child was stipulated to have a life “worth living.” However, David Benatar (2006) rejects the claim that *anyone* who has ever lived has had a life worth living. His proposal for solving the non-identity problem is to say that the 14-year-old girl harms her child in the same way that *every* couple impermissibly harms their child when they bring that child into existence. Benatar reasons that a person is harmed when she is made to suffer a disadvantage. Coming into existence is always a disadvantage because the presence of the pain of existence is bad, in comparison to the absence of pain that comes with non-existence. On the other hand, although presence of the pleasure of existence is good, it does not offer any advantage over the absence of pleasure that comes with non-existence; such an absence is not “bad,” but merely “not bad” (p. 38 – 41). Thus, the pain of existence grounds an advantage of non-existence over existence, but the pleasure of existence does *not* ground an advantage of existence over non-existence. Non-existence thus dominates; existence is the disadvantage. Thus, a person who comes into existence is always harmed by the action that causes her to exist.

Notice that I have classified Benatar's view together with the views that accept the counterfactually worse-off condition on harming. That is because, rather than accept the counterfactually worse-off condition and deny that a determinate child is harmed (as Heyd,

Boonin, and Hare do), Benatar implicitly accepts something like the counterfactually worse-off condition and then attempts to show *that it has been met*. That is, his claim that existence is a *disadvantage* is very similar to the claim that existence makes a person *worse off*. The problem, however, is that this move involves questionable hocus-pocus. As Elizabeth Harman (2009) points out, Benatar equivocates between *value for a person* and impersonal value. Despite what Benatar claims, the absence of pain that is attributable to the non-existence of a person is not better *for* the individual who would have experienced the pain; it is only *impersonally better*. Thus, non-existence does not confer any advantages on a person. But if that is so, Benatar's dominance argument does not succeed; Benatar fails to establish that coming into existence is always a harm.

A second problem with Benatar's view is that it does not *solve* the non-identity problem so much as it *ignores* the problem. Certainly, Benatar's view has an advantage over the views of Heyd, Boonin, and Hare insofar as Benatar attempts to account for the apparent wrongness of the 14-year-old girl's action. However, Benatar's account goes *too* far. We wanted to explain why a moral reason seemed to count against the 14-year-old girl's action, but we wanted the explanation to be *contrastive*: we wanted to see why the reason appears to hold in non-identity cases *but not* in ordinary cases. Benatar's view does not distinguish between non-identity cases and ordinary cases. Therefore, it resembles both the biting-the-bullet approach and Hare's approach in its failure to either account for or successfully explain away our intuitions.

3.3. Denying the Counterfactually Worse-Off Condition on Harming

In the previous section I considered three specific proposals that attempt to solve the non-identity problem by appealing to or at least implicitly endorsing the counterfactually worse-off condition on harming. I argued that there were particular problems with each view. Nevertheless, I believe that there is a more *general* problem with all of those approaches. The general problem is that the counterfactually worse-off condition on harming is false. In this section, I will explain why we ought to reject the counterfactually worse-off condition on harming.

I will use the term, “harm-based response,” to refer to any solution to the non-identity problem that rejects the counterfactually worse-off condition on harming. Those who opt for the harm-based response hold that an action can harm an individual even if it fails to make her worse-off in any respect than she would have been, had the action not been performed. In place of the counterfactually worse-off condition, many proponents of this response endorse a *non-comparative account* of harm, according to which harm is a state of affairs that is bad for an individual, such that the badness of the state does not derive from its being worse for the individual than some other state that the individual was or would have been in.⁴⁰ Many of them also hold that harming an individual is equivalent to causing that individual to suffer harm. The

40 I am following Matthew Hanser (2008) in defining “non-comparative harm” this way. Note that the term “non-comparative harm” is misleading because its definition does not rule out all comparisons. For example, the following qualifies as a non-comparative account of harm:

The Hot Tub Account of Harm: A state of affairs A is a harm for S if and only if A is *worse* for S than a state of affairs in which S is in a hot tub drinking root beer.

The only comparisons that the definition of non-comparative harm rules out are comparisons between the state that S is actually in and either (1) the state S was in at an earlier time or (2) the state S would have been in if the harmful event or action had not occurred.

combination of these two claims yields the following:

(H) Harming an individual, S, is causing a state of affairs, A, to obtain, such that A is non-comparatively bad for S.⁴¹

H seems to justify the claim that in *The 14-Year-Old Girl*, the girl harms her child. After all, the child's having a bad start in life is a state of affairs that is bad for him, and the girl's action causes that state of affairs to obtain. More generally, H appears to justify the claim that in any non-identity case, the child is harmed and thereby in some way wronged.⁴²

There are three considerations that support acceptance of H over acceptance of the counterfactually worse-off condition on harming. The first is that the counterfactually worse-off condition is *less plausible* than the claim that individuals in non-identity cases are in some way wronged. If solving the non-identity problem requires us to either reject the worse-off condition on harming or reject the claim that such individuals are in some way wronged, we ought to reject the worse-off condition on harming. And if we do reject the worse-off condition on harming, then we have reason to accept something like H in its place.

The second consideration that supports H over the worse-off condition is that people can

41 As I have formulated it, H combines multiple ideas, a subset of which many philosophers endorse, but it is not an exact statement of anyone's view, in particular. Proponents of the harm-based response (i.e., the response that rejects the worse-off condition on harming) include Seana Shiffrin (1999), Elizabeth Harman (2004, 2009), Matthew Hanser (2008, 2009), and Fiona Woollard (2012). Philosophers who accept the non-comparative account of harm include Shiffrin, Harman, and possibly Woollard, but not Hanser, who holds that harm is an event rather than a state. Philosophers who accept the claim that harming is causing harm include Shiffrin, Judith Jarvis Thomson (2011), Woollard, and Hanser, but Thomson combines this claim with a comparative account of harm, and Harman endorses causing harm as a sufficient but not a necessary condition for harming.

42 Notice that H invokes the concept of causation. Causation is compatible with both overdetermination and preemption, so if H is true, an action can harm an individual even if, had the action not been taken, the same bad state of affairs would have obtained anyway. In cases where the bad state of affairs would have obtained anyway, the individual is no worse off in any respect than she would otherwise have been; that is why H is incompatible with the counterfactually worse-off condition on harming.

be harmed in causal preemption and causal overdetermination cases, even though the worse-off condition implies that this is impossible. Fiona Woollard (2012) offers the following example:

Shooting Match: Through no fault of his own, Victor has made two terrible enemies, Adam and Barney, who have both sworn deadly vengeance upon him. Barney is just about to shoot and kill Victor. Barney is protected by a bullet-proof, sound-proof shield so that Adam can neither stop him forcibly nor dissuade him. Adam knows this, but Victor's death by another's hand will not satisfy his thirst for vengeance. Adam shoots Victor and Victor dies from the bullet wound (p. 8).

Shooting Match is a preemption case. As Woollard points out, Adam's action does not make Victor worse off than he would otherwise have been, for if Adam had *not* shot and killed Victor, then Barney would have. Nevertheless, Woollard's judgment is that Adam harms Victor. She takes the case to be evidence against the counterfactually worse-off condition on harming and in favor of something like H.

There is also a third consideration that supports rejecting the counterfactually worse-off condition on harming. The consideration has to do with cases in which there is no fact of the matter about how well-off an individual would have been if a harmful action had not been performed. Surprisingly, since he does not endorse anything close to H, Alastair Norcross describes such a case. (He uses the case to illustrate a point that is very different from the point I am illustrating, but the case works perfectly well for my purposes):

No Fact of the Matter: Two prisoners of conscience, Smith and Jones, are slated for execution in a small totalitarian republic. The president, Shrub, troubled by the effect of an Amnesty International campaign on his public image, decides to spare one of the prisoners. He cannot make up his mind whose life to spare, so he devises the following apparatus: Smith and Jones are placed in separate cells, each with air vents leading to a canister of poison gas, which is set to release its contents at noon. A computer is programmed to select a three-digit number at random at one second before noon. If the number is even or zero, the computer will close the air vent in Smith's cell; if the number

is odd, the computer will close the air vent in Jones's cell. The random-number selection process is truly indeterministic. The vice-president of the republic, Fowl, does not approve of Shrub bowing to liberal pressure. Fowl unplugs the computer at one minute before noon. The gas is released at noon, and both Smith and Jones die (Norcross 1990, p. 270).

Intuitively, Fowl's action harms both Smith and Jones. Nevertheless, there is no fact of the matter about how well-off Smith would have been, had Fowl not unplugged the computer. Nor is there a fact of the matter about how well-off Fowl would have been, had Fowl not done what he did. Thus, it is not the case that Fowl's action makes either Smith or Jones worse off than he otherwise would have been.

I conclude that we have more than enough reason to reject the counterfactually worse-off condition on harming. Nevertheless, a full solution to the non-identity problem and a full defense of DV2 against the non-identity objection cannot stop here. First, there are some objections to rejecting the counterfactually worse-off condition on harming, and in the remainder of this section I will consider and provide a preliminary response to those objections (a more thorough response will come in the next chapter). Second, there are also some problems that arise when we attempt to endorse H as an alternative to the counterfactually worse-off condition on harming. I will discuss those objections in the next chapter.

The primary objection to rejecting the counterfactually worse-off condition on harming is that when we do so, we undermine the original reason we had for thinking that harming is morally objectionable. If harming is a matter of making someone worse off, then it is easy to see why the fact that an action would harm someone is a strong reason against performing that action. However, once we *deny* the claim that harming necessarily makes a person worse off, we

seem to lose some of our justification for thinking that there is always—or at least usually⁴³—a strong reason against harming someone.

There is also a secondary, related objection. The objection is that we need a way to measure the severity of a harm. If being harmed were simply a matter of being made worse off than one otherwise would have been, then we could simply say that the degree of any harm is equal to the difference between how well-off the harmed individual *is* and how well-off she would have been, had the harm not been caused. This, I think, would have been the wrong thing to say, but no matter: we certainly can't say it if we deny the counterfactually worse-off condition on harming. And if we can't say it, it appears as though we are at a loss to differentiate between, say, the harm of having one's legs chopped off versus the harm of losing a pizza coupon.

Fiona Woollard advances a more nuanced version of the first objection. She distinguishes between “harming in the overall comparative sense,” and “non-comparative harming.” The former makes a person worse off than she otherwise would have been, and the latter does not. For ease of reference, I will modify Woollard's terminology slightly and use the term “comparative harming” instead of “harming in the overall comparative sense.” Woollard's Shooting Match case suggests that there is at least a weak reason against non-comparative harming. Intuitively, it is wrong for Adam to shoot Victor, even though Victor would have died anyway. Woollard contends that, even if Adam's shooting Victor would save a third party, Sarah,

43 Some people might hold that there is no reason against harming someone who deserves to be harmed. This view is compatible with the claim that there is a reason against harming anyone to whom we owe a duty of nonmaleficence; it may be that we simply do not owe such a duty to those who deserve to be harmed. Nevertheless, in this dissertation I leave it open whether there are cases of justified punishment and, if so, whether in such cases there is a reason against harming that is outweighed by the reason to punish, or whether in such cases there is no reason against harming at all.

from a scraped knee, it would *still* be wrong for Adam to shoot Victor.

Nevertheless, Woollard contends that the reason against comparative harming is much stronger than the reason against non-comparative harming. To support this claim, she argues that if Adam could save Sarah's *life* by shooting Victor, then Adam would be justified in doing so. On the other hand, it is *not* permissible for Adam to save Sarah's life by killing some other individual—call him Wendell—who was *not* going to die anyway. Adam's shooting Victor qualifies as non-comparative harming. Adam's shooting Wendell qualifies as comparative harming. It takes much more to justify shooting Wendell than it does to justify shooting Victor. Woollard concludes that, other things being equal, the reason against non-comparative harming is weaker than the reason against comparative harming. She then suggests that such a conclusion threatens harm-based solutions to the non-identity problem, for it shows that the harm-based reason against the action in a non-identity case may be “weaker than we originally thought” (p. xx).

However, there are two problems with Woollard's argument. The first is that, even if it is true that, other things being equal, reasons against non-comparative harming are weaker than reasons against comparative harming, this does not mean that the reasons against non-comparative harming are weaker than we originally thought; nor does it mean that they are too weak to explain our judgments about non-identity cases. The reason against causing 5 hours of pain is weaker than the reason against causing 10 hours of equally intense pain, but this does not mean that the reason against causing 5 hours of pain is weaker than we thought, or that we cannot invoke the badness of pain in our explanation of why there is a strong reason against causing someone 5 hours of pain.

The second problem is that the argument does not show that, other things being equal, reasons against non-comparative harming *are* weaker than reasons against comparative harming. That is because whether the harming is comparative or non-comparative is not the only feature that varies between the case in which Adam shoots Victor and the case in which Adam shoots Wendell. Another difference between the cases is that if Adam had not shot Victor, then Victor would still have suffered the same harm; however, it is not true that if Adam had not shot Wendell, then Wendell would have suffered the same harm. When an action causes a harm that the victim would have suffered anyway had the action *not* been performed, let us refer to the action as “redundant harming.” Adam's shooting Victor qualifies as redundant harming, and this explains why the reason Adam has against shooting Victor is relatively weak.

Of course, if all cases of non-comparative harming were also cases of redundant harming, then we might still be tempted to infer that an action's having non-comparative harming status is an indication that the action is more justifiable than a similar case of comparative harming. However, although all cases of redundant harming are cases of non-comparative harming, not all cases of non-comparative harming are cases of redundant harming. For example, the 14-year-old girl's action of having Billy qualifies as non-comparative harming, but it is not true that if she had not had Billy, then he would have had a bad start in life anyway.

Finally, we should not infer that the reason we have against non-comparative harming will always be as weak as the reason Adam has against shooting Victor. In many cases, the reason against non-comparative harming is *stronger* than the reason against redundant harming. Compare the following two cases:

Inducing Paralysis: It is Leland's 5th birthday, but unfortunately, a piece of debris, A, is flying through the air towards Leland's midsection. If Wesley does nothing, A will hit Leland and paralyze him from the waist down. Wesley's only other option is to push Leland into the path of another piece of debris, B, which would cause Leland a qualitatively identical injury; however, pushing Leland into the path of B would prevent Sarah from having a broken arm. Wesley pushes Leland into the path of B.

Selecting for Paralysis. Wendy is choosing which embryo to implant into her uterus. Embryo B will become an individual named Sheldon. B has a mutation that will cause Sheldon to become instantly and permanently paralyzed from the waist down on his 5th birthday. None of the other embryos have that mutation. However, implanting embryo B will prevent Sarah from having a broken arm. Wendy implants B.

Wesley's action qualifies as both non-comparative harming and redundant harming because it (1) makes Leland no worse off than Leland would have been, had Wesley not pushed him into the path of debris B and (2) causes Leland to suffer a harm that Leland would have suffered anyway. By contrast, Wendy's action qualifies as non-comparative harming but not redundant harming. Her action makes Sheldon no worse off than Sheldon would have been, had Wendy not acted as she did, but her action does *not* cause Sheldon to suffer a harm that Sheldon would have suffered anyway. Notice that both actions result in the existence of an individual who is paralyzed from the waist down on his 5th birthday. Nevertheless, my judgment is that, although Wendy's action *may* be permissible, it is still a morally worse action than Wesley's. This suggests that other things equal, an action that qualifies as non-comparative harming but not redundant harming is *morally worse* than an action that qualifies as redundant harming. Since most non-identity cases are cases of non-comparative harming but not redundant harming, we have reason to think that, contrary to Woollard's assertion, the reason against non-comparative harming is relatively strong.

Of course, we might still want an account of *why* the reason against non-comparative

harming is relatively strong. And in response to the second objection to rejecting the counterfactually worse-off condition on harming, we would also like to know how to measure the *degree* of various harms. In the next chapter, I will provide a positive account of harming that satisfies both desiderata.

3.4. Conclusion

In this chapter I introduced the non-identity problem and showed how one approach to solving the problem—namely, by accepting something like the counterfactually worse-off condition on harming—motivates an objection to DV2 that I called the non-identity objection. I argued that such an approach to the non-identity problem is unsuccessful. I further argued that we ought to reject the counterfactually worse-off condition on harming. I finally considered and provided a preliminary response to some objections to rejecting the counterfactually worse-off condition on harming.

Chapter 4: The Existence Account of Harming

4.0. Introduction

Recall that my primary goal in Chapters 2, 3, and 4 is to defend DV2, the claim that in the present circumstances, our duty to future humans of nonmaleficence grounds a reason for us to refrain from damaging certain features of the natural environment. DV2 is true only if, by damaging the environment in certain ways, we can harm future people. The non-identity objection to DV2 threatens to undermine that claim. According to the objection, our actions make future humans no worse off than they would have been, had we acted differently, for if we had acted differently, those future humans would not have existed; thus, since there is a counterfactually worse-off condition on harming, our environmentally destructive actions cannot harm future people.

In response to the non-identity objection, I argued that we should reject the counterfactually worse-off condition on harming, and I gave some counterexamples to the condition. However, merely rejecting the counterfactually worse-off condition on harming is not a satisfactory stopping point. It would be better if I could offer an account of harming that fully vindicates the harm-based solution to the non-identity problem, explains why there is a reason against harming, and illuminates the reasons why some harms are more severe than others. That is what I will do in this chapter.

Here is how the chapter will be organized. In section 4.1 I will argue that the claim I referred to in the last chapter as “H” will not work as a satisfactory account of harming, despite the fact that it does not entail the counterfactually worse-off condition on harming. In sections

4.2 and 4.3 I will advance my own account of harming, and I will explain how the account can help us determine why there is a reason against harming and why some harms are worse than others.

4.1. Why H is not a Satisfactory Account of Harming

Recall from Chapter 3 that H is the following claim:

(H) Harming an individual, S, is causing a state of affairs, A, to obtain, such that A is non-comparatively bad for S

As I said at the time, H has a number of virtues. First, because it does not endorse the counterfactually worse-off condition on harming, H is compatible with the view that harming can occur, not only in cases of overdetermination and cases of preemption, but also in cases in which there is no fact of the matter about how well-off an individual would have been had the event that harmed her not occurred. Second, H seems to justify the claim that individuals are wronged in non-identity cases. In such cases, the actions that bring individuals into existence also impose constraints on their life prospects. The state of affairs in which an individual experiences constraints on her life prospects is bad, and an action who cause the existence of such an individual is causing that state of affairs. This result is particularly auspicious for my purposes because it seems to vindicate the claim that we can harm future humans with the very actions that are the condition of their own existence.

Despite the promise that it holds as both a solution to the non-identity problem and a vindication of DV2, H faces a number of problems. First, H gets the wrong result in cases in

which an action or event causes an individual to be in a state that is bad for her without harming her. Second, H gets the wrong result in cases in which an action or event harms an individual without causing her to be in a state that is bad for her. Third, it is difficult to determine just what it is that makes a state of affairs non-comparatively bad *for some individual*, especially in cases where a similar state of affairs would not be bad for another individual.

The first problem is motivated by cases like the following:

Dim Vision. Jones has been blind for many years as a result of retinal damage. Recently, Dr. Smith has developed a surgical operation that can repair some but not all of the damage. Dr. Smith operates on Jones and improves his vision from a state of blindness to a state in which Jones can see, but not very well: Jones now has what we will call *dim vision*.⁴⁴

It seems that Jones's having dim vision is a state of affairs that is bad for Jones. Dr. Smith causes that state of affairs to obtain. Thus, it would seem that if H is true, then Dr. Smith harms Jones by performing the operation. However, it is intuitively clear that Dr. Smith does not harm Jones.⁴⁵

The second problem for H is illustrated by the following:

Death. Sam, a healthy and happy 25-year-old, is attempting to cross the train tracks when his foot gets stuck. A train hits him and he dies instantly.⁴⁶

44 This example is an adaptation of cases that appear in Thomson (2011) and Hanser (2009).

45 Harman (2009) responds to this objection by distinguishing between a “particular bad state” and a “general bad state.” She claims that causing an individual to be in a general bad state is a sufficient condition for harming him or her. Presumably, Harman would say that dim vision is a *particular* bad state, so it does not follow from the claim that Dr. Smith caused Jones to be in such a state that Dr. Smith harmed Jones. This response is unsatisfying. The notions of generality and particularity that Harman relies upon face a problem analogous to the generality problem for reliabilism: we need either a criterion or firm intuitions to help us distinguish general bad states from particular bad states, and we lack both

46 This case is inspired by Hanser's (2008) argument that in general, state-based accounts of harm cannot successfully explain why death is a harm. A state-based account of harm is one according to which a harm is a state of affairs; the non-comparative account of harm is one such account. As an alternative to state-based accounts, Hanser favors an event-based account of harm. I do not critique his account here, but see Thomson (2011) for some objections to Hanser's view.

The event that consists of the train's hitting Sam does not appear to cause a state of affairs to obtain that is bad for Sam. That is because the state of affairs that the accident brings about is one in which Sam no longer exists, and it is difficult to see how a state of affairs can be bad for an individual who no longer exists. However, it is intuitively clear that the train accident harms Sam.

One might think that if we had a more precise account of what it is for a state to be non-comparatively bad for someone, we could explain away the apparent counterexamples and thereby vindicate H. However, a dilemma arises when we attempt to make the account of non-comparative badness more precise. Non-comparative badness for an individual must either be absolute or relative. But the most plausible account of badness as an *absolute* notion still renders H susceptible to apparent counterexamples, and the most plausible accounts of badness as a *relative* notion either undercut the original motivation for H, or else, like the absolute account, they render H susceptible to apparent counterexamples.

The most plausible account of non-comparative badness as an *absolute* notion holds that a state of affairs is non-comparatively bad for an individual if and only if it would be bad for *anyone*. What it is to be *bad for anyone* is defined extensionally: states that involve pain, suffering, disease, disability, and so on are bad for anyone. The problem with this account is that there appear to be cases in which an event harms someone by causing him to be in a state that would *not* be bad for just anyone. Consider the following:

Loss of Fortune. Jeeves was once a world-renowned physicist with extraordinary

intellectual abilities. He then had a stroke and suffered brain damage. The brain damage left him with average intellectual abilities.⁴⁷

It seems that having average intellectual abilities would not be a bad state for just anyone.

Nevertheless, the stroke *harmed Jeeves*.

An alternative way to explain what it is for a state to be non-comparatively bad for an individual is to relativize the badness to something about the individual. However, it is difficult to determine what it is about the individual that the badness should be relativized to. Although Elizabeth Harman suggests that badness should be relativized to the norm for the individual's species, Loss of Fortune illustrates the problem with that approach. Jeeves is a member of the human species, and it does not *seem* to be bad for a human to have average intellectual abilities, even though coming to have average intellectual abilities is a harm for Jeeves.⁴⁸

Could non-comparative badness be relativized to the individual's will or desires?⁴⁹

Neither choice is promising. A dog with cancer has neither the will nor the desire to be cancer-free, for she does not have the concept of cancer that is a necessary condition for having such a will or a desire. Nevertheless, when she gets cancer, she is harmed.

47 This case is modeled upon the Bertrand Russell case in Jeff McMahan (1996) and the case of the Nobel Prize winner in Hanser (2008).

48 This objection to H does not apply to Harman's (2009) view. That is because Harman argues that causing someone "to be in pain, to be in mental discomfort, to be in physical discomfort, to have a disease, to be deformed, to be disabled, or to die" is a *sufficient condition* for harming, not a necessary condition (p. 149). However, as I point out in note 10, Harman's view would be more persuasive if it were grounded by a complete analysis of harm, such that the analysis explained *why* causing someone to be in pain, discomfort, and so on was sufficient for causing harm.

49 Seana Shiffrin (1999) and George Pitcher (1984) appear to relativize harm to a person's will, desires, or interests. Shiffrin writes that harms are "conditions that generate a significant chasm or conflict between one's will and one's experience, one's life more broadly understood, or one's circumstances" (p. 123). Pitcher writes, "An event or state of affairs is a harm for someone (or harms someone) when it is contrary to one or more of his more important desires or interests" (p. 184).

How about interests?⁵⁰ If we hold that a state is non-comparatively bad for someone if and only if it is contrary to one of her interests, we are only pushing the problem back a step. For we must now decide whether to take an absolute or relativistic approach to deciding *which* interests an individual has. If we attribute the same interests to all sentient individuals, then we face the kind of problem that arose for making the concept of non-comparative badness absolute: the set of universal interests will inevitably be too broad or too narrow. Either we will say that blindness is contrary to the interests of a bat, and thus, that a bat's mother harms him when she conceives him (and thereby causes him to exist in a state of blindness), or we will say that having average intellectual abilities is *not* contrary to Jeeves's interests, so the stroke that causes Jeeves to have average intellectual abilities does *not* harm him.

On the other hand, if we relativize interests, we must decide what to relativize them to. If we relativize interests to species membership, desires, or will, we will again have trouble explaining why Jeeves or the dog with cancer is harmed. Suppose we relativize interests to the level of well-being that individuals *previously* had, or to the level of well-being that individuals *would have had* if some event had not occurred: we say that Jeeves has an interest in having extraordinary intellectual abilities because having such abilities would restore his well-being to pre-stroke levels, or because having such abilities would boost his well-being to the level at which it *would* have been if he had not had the stroke. If we take this option, we undermine the original motivation behind H, which was to vindicate the notion that individuals can be harmed by actions that restrict their lifetime prospects of well-being even if such individuals *never had*

⁵⁰ Joel Feinberg (1984) defends a highly influential account of harm that relativizes harm to an individual's interests. However, he endorses a comparative account of harm.

or *would not have had* unrestricted prospects.

Jeff McMahan suggests that whether an individual is “unfortunate” can be relativized to her “native potential” for having a certain level of well-being, where native potential has something to do with “the physical constitution of the individual.”⁵¹ Let us attempt to formulate this idea as an account of a non-comparatively bad state:

The Native Potential Account of Non-Comparative Badness: A state of affairs is non-comparatively bad for someone when it is worse for her than some other state of affairs that could have obtained, given her native potential.

This suggestion is again at odds with the motivation for H. We wanted to endorse H as a solution to the non-identity problem. However, there are many non-identity cases in which the constraints imposed on an individual's life prospects are genetic. Consider, for example, the following:

Selecting for Disability. Barbara wishes for her child to have disability *X*. She uses *in vitro* fertilization and selects an embryo for the gene that causes *X*. The embryo becomes a child named Leonardo, who develops *X*. Having *X* causes Leonardo to experience hardship, pain, and suffering. However, Leonardo has a life worth living.

In *Selecting for Disability*, Leonardo's disability is a consequence of his physical constitution; he is genetically predisposed to develop the disability. Thus, since he does not have the potential *not* to have the disability, the native potential account implies that Leonardo's having the disability is not bad for him.

H is in trouble. H says that harming an individual is causing a state of affairs that is non-comparatively bad for that individual. But it seems that, however we define what it is for a state to be *non-comparatively bad for* someone, either we undercut the original motivation for H, or else we invite apparent counterexamples. And if H is in trouble, then so is the harm-based

51 McMahan (1996), p. 22.

response to the non-identity problem, for without something like H, the harm-based response lacks a solid theoretical grounding. And if the harm-based response to the non-identity problem is in trouble, then DV2 is on shaky ground.

4.2. The Existence Account of Harm

In this section, I will present what I call the *existence account of harm* and the *gradational account of inevitability*. These two accounts can be incorporated into a non-comparative account of harming that avoids the pitfalls of H. The accounts are as follows:

The Existence Account of Harm (EH): A state of affairs, A, is a harm for an individual, S, just in case if it were true that both S existed *and* A did not obtain, then S would have a higher level of lifetime well-being.

The Gradational Account of Inevitability (GI): For any state of affairs, A, A is *more inevitable (less avoidable)*, the more the world would be different if A did not obtain, and A is *less inevitable (more avoidable)*, the less the world would be different if A did not obtain.

As we shall see, the existence account of harm is designed to fill the role played by the accounts of either harm or misfortune that are already on offer. However, the gradational account of inevitability does not, as far as I know, have competitors. I use “inevitability” and “avoidability” as terms of art; according to ordinary usage, the terms “inevitable” and “avoidable” do not admit of degree, but as I use them, they do.

The existence account of harm differs in important ways from other accounts of harm and misfortune. Consider, for example, the account of harm that is presupposed by the counterfactually worse-off condition on harming. Such an account is something like this:

The Counterfactually Worse-Off Account of Harm: A state of affairs, A, in which an individual, S, exists is a harm for S if and only if S is worse off in some respect for being in A than she would have been if A had not obtained.

Both the existence account of harm and the counterfactually worse-off account of harm invoke counterfactuals, but the antecedents of those counterfactuals are different. The antecedent of the counterfactual in the counterfactually worse-off account of harm picks out the possibility that A does not obtain. On the other hand, the antecedent of the counterfactual in the existence account of harm picks out the possibility that S exists *and* A does not obtain.

Because of this difference in how the antecedents are formulated, the existence account of harm allows for the possibility that in *The 14-Year-Old Girl*, Billy's having a bad start in life is a harm; however, the counterfactually worse-off account of harm does not allow for that possibility. To see why, suppose that the following claims can both be truthfully asserted by Billy:

(B1) If I did not have this bad start in life, then someone else would have come into existence in my place.

(B2) If I existed but did not have this bad start in life, then my mother would have gone to the fertility clinic and had the zygote from which I came frozen for 10 years and then implanted when she was 24.

B1 refers to a state of affairs in which Billy does not exist, and in which, *a fortiori*, Billy is not worse off in any respect. The conjunction of B1 and the counterfactually worse-off account of harm implies that Billy's having a disability is not a harm. However, B2 refers to a state of affairs in which Billy *does* exist *and* has a higher level of lifetime well-being. As a consequence, B2 and the existence account of harm imply that Billy's having the disability is a harm.

The gradational account of inevitability provides us with a new dimension along which we can measure harms: one harm may be more or less inevitable than another. As we shall see, the degree to which a harm is inevitable affects the strength of the reason against causing it.

4.3. The Existence Account of Harming

In the previous section I introduced the existence account of harm and the gradational account of inevitability. In this section and the following section, I will explain how the two accounts can be used to defend DV2 against the non-identity objection. I begin in this section by presenting some supplementary principles about harming. I will refer to the conjunction of these principles with the existence account of harm and the gradational account of inevitability as the *existence account of harming*. I will show how the existence account of harming explains why there is a reason against what the 14-year-old girl does. I will then show how the existence account of harming avoids the problems that Dim Vision, Death, and Loss of Fortune present for H.

Here are the supplementary principles about harming:

Harming as Causing (HC): An action or event harms an individual, S, if and only if that action or event causes a harm for S to obtain.

The Inevitable Harming Principle (IHP): Other things being equal, the reason against causing a harm for an individual to obtain is weaker, the more inevitable the harm.

The existence account of harming consists of the conjunction of HC, IHP, the existence account of harm, and the gradational account of inevitability. As a reminder, here again are the latter two principles:

The Existence Account of Harm (EH): A state of affairs, A, is a harm for an individual, S,

just in case if it were true that both S existed *and* A did not obtain, then S would have a higher level of lifetime well-being.

The Gradational Account of Inevitability (GI): For any state of affairs, A, A is *more inevitable (less avoidable)*, the more the world would be different if A did not obtain, and A is *less inevitable (more avoidable)*, the less the world would be different if A did not obtain.

Here is how the existence account of harming can help explain why there is a reason against what the 14-year-old girl does. First, it is reasonable to think that Billy would have a higher level of lifetime well-being if he were to exist without having had a bad start in life. For example, if Billy's mother had frozen the zygote from which Billy came and then implanted it when she was older, Billy's lifetime level of well-being would be higher. Since his level of lifetime well-being would be higher if he existed without having had a bad start in life, the existence account of harm implies that the bad start in life is a harm for Billy. The 14-year-old girl's act of having a baby causes the state of affairs in which Billy has a bad start in life; thus, HC implies that she harms him. We can then complete reasoning that solves the non-identity problem: the fact that the girl's action harms Billy explains why she wrongs him. This reasoning generalizes to any non-identity case in which an individual would have a higher level of lifetime well-being if one of the other consequences of the action that brought her into existence did not obtain.

An advantage of the existence account of harming is that it is immune to the objections illustrated by Dim Vision, Death, and Loss of Fortune. Recall that in Dim Vision, Dr. Smith causes Jones to be in a state of dim vision. Such a state appears to be bad for Jones, so if H is true, then Dr. Smith harms Jones. However, the existence account of harming does not imply that

Dr. Smith harms Jones. Consider the following counterfactual:

(J) If it were true that both Jones existed *and* he did not have dim vision, then Jones would have a higher level of lifetime well-being.

J states a necessary condition for dim vision's being a harm for Jones. However, J is false. If Jones were to exist without dim vision, then Jones would be completely blind; that would result in his having a *lower* level of lifetime well-being. The state that Dr. Smith causes Jones to be in fails to satisfy a necessary condition for being a harm, and since it is not a harm, Dr. Smith does not harm Jones in causing it.

This result may seem worrisome, for as I noted above, the state of dim vision *appears* to be bad for Jones. However, it is easy to explain away the appearance of badness. What is *actually* bad for Jones is not dim vision, in particular, but rather the more general state of not being able to see very well. And sure enough, the existence account of harm implies that the state of not being able to see very well *is* a harm for Jones. Consider the following counterfactual:

(W) If it were true that both Jones existed *and* he were *not* not able to see very well, then Jones would have a higher level of lifetime well-being.

W is true. To say that Jones is *not* not able to see very well is to say that Jones *is* able to see very well. And if Jones were to see very well, then he would indeed have a higher level of lifetime well-being.

The existence account of harming is also immune to the objection illustrated by Death. Recall that in Death, Sam dies instantly. The event of his death does not result in a state that is bad for Sam, for the state that results is not a state *in which Sam exists*. That is why H cannot explain why the train accident harms for Sam.

However, according to the existence account of harm, the state of affairs that is the standard for assessment is not the actual state of affairs in which Sam no longer exists, but rather the possible state of affairs in which Sam exists without having been killed by the train. The question is whether, in *that* state of affairs, Sam has a higher level of lifetime well-being. And since in that state of affairs, Sam would live happily and healthfully past the age of 25, he *would* have a higher level of lifetime well-being: living longer contributes more well-being to his lifetime total. Thus, the existence account of harm implies that it is a harm for Sam that he was killed by the train. Moreover, HC implies that the accident harmed him.

Finally, the existence account of harming is immune to the objection illustrated by Loss of Fortune. Recall that in Loss of Fortune, Jeeves suffers brain damage that results in his having average intellectual abilities. The proponents of H were at a loss to explain why having average intellectual abilities was a bad state for Jeeves, so they could not use H to explain why the stroke harmed Jeeves. However, the existence account of harm implies that having average intellectual abilities is a harm for Jeeves, for the following counterfactual is true:

(I) If it were true that both Jeeves existed *and* he did not have average intellectual abilities, then Jeeves would have a higher level of lifetime well-being.

After all, if Jeeves did not have average intellectual abilities, then he would not have had the stroke, and so he would have retained his extraordinary intellectual abilities; thus, he would have a higher level of lifetime well-being. The existence account of harming implies that the stroke harmed Jeeves.

So far, I have not invoked the IHP. The existence account of harm and the HC, alone, get

the right results in *The 14-Year-Old Girl*, *Dim Vision*, *Death*, and *Loss of Fortune*. Nevertheless, the IHP is an important part of the existence account of harming. To see why, consider the following objection to the view.

The objection is that, as the example of *not being able to see very well* illustrates, the existence account of harm is compatible with the claim that states of affairs that are either positively described (being blind) or negatively described (*not* being able to see very well) can be harms. But if we allow states of affairs that are negatively described to be harms, then presumably, not having almost *anything* good would qualify as a harm. For example, it is a harm for me that I do not have telekinetic abilities, since if both (1) I existed and (2) I did *not not* have telekinetic abilities, I would have a higher level of lifetime well-being. Yet it is not a harm for me that I lack telekinetic abilities.

My response to this objection is that the intuition that not having telekinetic abilities is *not* a harm is not very robust, and consequently, we should disregard it, since we otherwise have strong reason to accept the existence account of harm. The same is true for our intuitions about the lack of other goods. Not having wings, not being able to live to the age of 200, not having the intellectual abilities of a telepathic super genius: these are harms for most, if not all, of us.⁵² I noted in Section 4.1 that it does not *seem* to be bad for humans who are not Jeeves to lack extraordinary intellectual abilities. I also noted that, intuitively, a mother bat does not harm her

52 Cf. Thomas Nagel (1970), pp. 73 – 80. Nagel writes, “The question is whether we can regard as a harm any limitation, like mortality, that is normal to the species. ... A man's sense of his own experience ... does not embody this idea of a natural limit. His existence defines for him an essentially open-ended possible future, containing the usual mixture of goods and evils that he has found so tolerable in the past. ... Viewed in this way, death, no matter how inevitable, is an abrupt cancellation of indefinitely extensive possible goods. Normality seems to have nothing to do with it, for the fact that we will all inevitably die in a few score years cannot by itself imply that it would not be good to live longer” (p. 80).

baby when she causes him to exist in a state such that he is unable to see. I now claim that such intuitions are misleading: we have a principled reason to hold both that it *is* a harm for many people to lack the intellectual abilities of Jeeves and that a mother bat *does* harm her baby when she brings him into a typical bat existence. Nevertheless, we can still claim that any *reason* against what the mother bat does, or against what any human parents would do, were they to bring into existence an individual of ordinary intellectual ability, would be vanishingly weak. The justification for these claims comes from the IHP.

Recall that the IHP says that, other things being equal, the reason against harming an individual by causing her a harm is weaker, the more inevitable the harm. For a bat, being unable to see is a harm that is highly inevitable. That is because, in a world where a bat could see, there would be extraordinary developments in genetic engineering, and such a world is very different from our own. Given how inevitable blindness is for a bat, IHP implies that, other things being equal, the reason against causing a bat to be born blind is much weaker than the reason against causing a human to be born blind.

Moreover, not having extraordinary intellectual abilities is more inevitable for the average person than it is for Jones. Therefore, according to IHP, the reason against causing the average person not to have extraordinary intellectual abilities is much weaker than the reason against causing Jones not to have extraordinary intellectual abilities. Compare the following two cases:

Inducing Average. Alice is already pregnant with the fetus who will become Jeeves. She decides to take a drug that will cause the fetus brain damage such that Jeeves will have average intellectual abilities instead of extraordinary ones. After he is born, Jeeves

develops average intellectual abilities.

Selecting for Average. Barbara is using *in vitro* fertilization. She would like to have a child who has average intellectual abilities, so she selects an embryo on this basis. After he is born, her child, whom she names Reeves, develops average intellectual abilities.

Having average intellectual abilities is more inevitable for Reeves than it is for Jeeves. Given that other factors are equal between the cases, IHP implies that the reason against what Alice does is stronger than the reason against what Barbara does.⁵³ Since we can distinguish between the strength of reasons in this way, we need not bother with further distinctions implying that blindness is a harm for a human but not for a bat, or that not having extraordinary intellectual abilities is a harm for Jeeves but not for Reeves.⁵⁴

Here are some further considerations that support the existence account of harming. Recall that non-comparative harming is harming in which the individual harmed is no worse off than she would have otherwise been. The existence account of harming implies that non-comparative harming *is* a form of harming. However, as I noted in Chapter 3, we would like an account of why the reason against non-comparative harming is relatively strong. And in response to the second objection to rejecting the counterfactually worse-off condition on harming, we would also like to know how to measure the *degree* of various harms.

53 Indeed, for any action that causes an independently existing person to suffer a harm, it is likely that, according to IHP, the reason against the action is strong. Consider a typical case in which an action causes a person who would exist regardless of the action to suffer a harm. In such a case, the harm is highly avoidable: if the harm did not obtain, the only thing that would be different would be that the action was not performed. Since the harm is highly avoidable, the reason against the action (other things being equal) is strong.

54 We might also think that blindness is less of a harm for a bat than it is for a human because being able to see would not improve a bat's well-being as much as it would improve a human's well-being. The existence account of harm is compatible with this view.

Here is why the reason against non-comparative harming is relatively strong. The reason is that non-comparative harming, like all other forms of harming, is causing a state of affairs that detracts from an individual's well-being. Moreover, we can also use the *amount* by which a state of affairs detracts from an individual's well-being as a measure of the severity of a harm.

Suppose, for example, that my legs have been violently chopped off. This is a harm for me because if I existed and the state of affairs did not obtain, then my total level of lifelong well-being would be higher. Indeed, it would be *much* higher—let us say, by 90 units. On the other hand, suppose that I lost a coupon for a free medium pizza. This is also a harm for me because if I existed and the state of affairs in which I lost the coupon did not obtain, then my total level of lifelong well-being would be higher. However, it would not be much higher; having the coupon would add about 2 units of well-being, at most, to my lifelong total. We can therefore use the difference in well-being to measure the severity of various harms.

Keep in mind, however, that the reason against harming is not merely a function of how severe a harm is. The considerations I raised in Chapter 3 suggest that other things being equal, non-redundant harming is morally worse than redundant harming. It is probably also true that, other things being equal, non-comparative harming is morally better than comparative harming but worse than harming that is both non-comparative and redundant. Finally, the IHP also influences the strength of the reason against harming. For example, an act of non-comparative harming that causes a bat to be blind is morally better than an act of non-comparative harming that causes a human to be blind.

4.4. Conclusion

I conclude that we ought to accept the existence account of harming. Such an account solves the non-identity problem, accounts for the strength of the moral reason against harming, and escapes the objections that plagued H. But if we accept the existence account of harming, then we no longer need to worry about the non-identity objection to DV2; the existence account of harming implies that the first premise of the non-identity objection is false.

Chapter 5: Causes and Mere Conditions

5.0. Introduction

The existence account of harming implies, roughly, that harming is causing a harm. Thus, the matter of whether a state of affairs has been *caused* by an event or action partly determines whether an instance of harming has occurred. Still, one might think that my account need not invoke any substantive claims about causation; perhaps any account of causation can be plugged into the existence account of harming to yield morally acceptable results.

It turns out, however, that this is not so. Just as we saw in Chapter 1 that the metaphysics of time made a material difference to whether future humans have moral standing, so we will see in this chapter that the metaphysics of causation make a material difference to the plausibility of various accounts of harming. To show why the metaphysics of causation matter to the metaphysics of harming, I will use Elizabeth Harman's account of harming as a test case.

Harman's account is very similar to my own; Harman holds that harming is causing harm and that harm is a non-comparatively bad state of affairs. Nevertheless, Harman makes a substantive assumption about the metaphysics of causation, and this assumption creates problems for her view; I call one of these problems the *problem of too much harm*. I will argue that, to handle these problems, Harman supplements her otherwise plausible account of harming with *ad hoc* moral principles that stand in tension with the original motivation for her account. The account would be more plausible if it were premised upon an alternative assumption about causation.

In this chapter, I will first motivate the claim that the metaphysics of causation matter to the metaphysics of harming by explaining the problem of too much harm and using this problem

to develop a critique of Harman's view. In Section 5.2, I will explain why my own account of harming would be subject to the problem of too much harm if I were to endorse Harman's assumption that something I call *forward dependence* is sufficient for causation. However, I will argue that I need not assume that forward dependence is a sufficient condition for causation; there are alternative accounts of causation, compatible with my account, which do not imply that forward dependence is sufficient for causation. Moreover, I will argue that at least one of these accounts supports a metaphysical distinction between causes and mere conditions, and that this distinction can insulate my view from the problem of too much harm. Finally, in Section 5.3, I will argue that even those who reject my solution to the non-identity problem and endorse *comparative* accounts of harming should nonetheless accept only those accounts of causation that support a metaphysical distinction between causes and mere conditions.

5.1. The Problem of Too Much Harm

Harman takes what I have called the “harm-based” approach to the non-identity problem. That is, as I do, she rejects the counterfactually worse-off condition on harming. In place of the counterfactually worse-off condition, Harman defends a sufficient condition for harming that she formulates as follows:

(H1) One harms someone if one causes him to be in pain, to be in mental discomfort, to be in physical discomfort, to have a disease, to be deformed, to be disabled, or to die (2009, p. 149).

Although she does not offer a full account of the necessary *and* sufficient conditions for harming, she claims that H1 is the only principle we need to solve the non-identity problem. If we make

certain assumptions about causation, this at first seems plausible. Consider, for example, The 14-Year-Old Girl. Assume that by causing Billy to have a bad start in life, the 14-year-old girl also causes Billy to be in mental discomfort. If so, then H1 allows us to say that the 14-year-old girl harms Billy. The fact that she harms him then appears to explain the intuition that she in some way wrongs him.

But why, exactly, should we think that the 14-year-old girl causes Billy to be in mental discomfort? Harman writes, “I assume that counterfactual dependence is sufficient for causation: I assume that if an action is such that if it had not been performed, a particular event would not have occurred, then the action causes the event” (2009, p. 140). If the 14-year-old girl had never had Billy, he would not have experienced mental discomfort. So if counterfactual dependence is sufficient for causation, then the very action that causes Billy to exist also harms him.

However, Harman's substantive assumption about the nature of causation adds a new complication. Not only does her account imply that Billy is harmed by the action that brings him into existence, but it also implies that *everyone* is harmed by the action that brings him or her into existence. After all, if counterfactual dependence is sufficient for causation, then *every* act of procreation causes the resultant individual to experience whatever pain, physical discomfort, mental discomfort, diseases, deformities, or disabilities he or she suffers from throughout her life. Harman's account therefore implies that not only in non-identity cases, but also in ordinary cases, procreative actions cause the resultant individuals a great deal of harm. On the face of it, this implication of the view is counterintuitive. I call the implication that *every person* is harmed by the action that brings her into existence the *problem of too much harm*.

Harman's response to the problem of too much harm, and her full solution to the non-identity problem, begins with the observation that if counterfactual dependence is sufficient for causation, then not only can procreative actions harm the resultant individuals, but so can such actions can benefit them. She then argues that while individuals are harmed by being conceived in both non-identity cases and in ordinary cases, the benefits that the procreative actions cause those individuals outweigh the harms in most ordinary cases and in some non-identity cases. In the ordinary cases in which the benefits outweigh the harms, the fact that the benefits outweigh the harms can justify the act of conceiving the individual. However, in many of the non-identity cases in which the benefits outweigh the harms, the fact that the benefits outweigh the harms *cannot* justify the act of conceiving the individual. This is because, in such non-identity cases, there's some other individual that the agent could conceive instead, and that other individual would not suffer the same harms. Harman formulates the principle behind such reasoning as follows:

(H2) If an action harms someone who does not independently exist, then the fact that the action also benefits the person harmed, and benefits him more than it harms him, is *ineligible* to justify the harm if *failing to perform the action* would similarly benefit someone (2009, p. 139, italics are Harman's).

H2 implies that if a potential parent faces a choice between having Baby A or having Baby B, and if the choice is such that having A would cause A to incur certain harms and having B would not cause B to incur those particular harms, then the benefits in the life of A do not justify having A instead of B.

This appeal to H2 is problematic for two reasons. First, it purports to explain the

appearance of wrongness *only* in those non-identity cases in which the agents could have brought another individual into existence. Recall from Chapter 3 that a non-identity case is any case in which an individual appears to be in some way wronged by the action that brings her into existence; it is not essential to a non-identity case that the agents could have brought some other individual into existence. At most, then, H2 can explain our intuitions about *some* non-identity cases. Since it cannot explain our intuitions about *all* non-identity cases, it fails as a general solution to the non-identity problem.

Harman anticipates this objection and responds that non-identity cases fall into at least two different classes. In what she calls “same-number cases,” agents face a choice between alternatives, all of which will result in the same number of people coming into existence. In what she calls the “hard cases,” the agents know that performing the procreative action will cause more people to exist than failing to perform the action. Harman acknowledges that the hard cases fall outside the scope of H2; for all that H2 says, the benefits that procreative actions confer upon the resultant individuals in hard cases *may* or *may not* justify the harms that those actions cause the individuals. She argues, however, that the fact that H2 leaves the permissibility of the hard cases unsettled is a virtue of her view; it explains why different people have different intuitions about such cases.

This response to the objection is unsuccessful. It may well be that different people have different intuitions about whether procreation in the hard cases is all-things-considered permissible. But I take it that the intuition to be explained in non-identity cases is not an intuition about all-things-considered permissibility. Rather, it is an intuition about a particular *reason* that

counts against the action. Intuitively, such a reason exists in *all* non-identity cases, not just the same-number cases.

Harman might respond that I have simply misconstrued the problem; she might claim that to the contrary, the non-identity problem is one of explaining why procreation appears to be all-things-considered wrong in some but not all non-identity cases, rather than a problem of explaining why there appears to be a reason against procreation in all non-identity cases. But even if I grant this way of construing the non-identity problem, I contend that there is still a *related* problem to be solved, namely, the problem of explaining why there appears to be a reason against procreation in cases like the following:

Hard Case: Alberta has a condition that makes her unable to bear more than one child in her lifetime. If she bears a child, that child will suffer from a genetic condition that makes him suffer from a moderate amount of pain and hardship throughout his life, but he will still have a life worth living. She bears this child and names him William.

Alberta's decision in this case might be morally justified, all things considered. Nevertheless, her action appears to in some way wrong William,⁵⁵ and Harman's proposal for solving the non-identity problem cannot explain the appearance of wrongdoing in this kind of case. This is a relative failing of Harman's view because, other things being equal, we would prefer to have a unified explanation for both the appearance of all-things-considered wrongness in some non-identity cases and the appearance of a reason against the procreative action in all of them.

The second problem with Harman's proposal for solving the non-identity problem is that even in the sorts of cases with which Harman is most concerned, her proposal gets the right

⁵⁵ Recall from Chapter 3 that I use the term, "in some way wrong" to indicate that there is at least an fh-reason against the action, even though the action might still be all-things-considered permissible.

answers for the wrong reasons. We had wanted to appeal to harm as the wrong-making feature in non-identity cases. Harman's view allows us to say that non-identity individuals are harmed but then appeals to additional features to explain the intuition that these individuals are wronged: on Harman's account, another important wrong-making feature in a non-identity case is the fact that someone else could have been better off. The second problem, then, is that Harman's view cannot accommodate both of the following claims: (1) in the non-identity cases in which procreation is wrong, what *primarily* makes it wrong is *harm to the child*; and (2) the primary feature that makes conception wrong in those cases is *not equally operative* in cases of ordinary procreation.

To address these problems, Harman should renounce her claim that both non-identity procreation and ordinary procreation harm the child who is conceived. Instead, she should say that non-identity procreation harms the child conceived, but ordinary procreation does not. Saying this is more in keeping with the intuition that actions in non-identity cases are in some way more objectionable than actions in ordinary procreation cases. Nevertheless, if Harman continues to hold that harming is causing harm, she cannot also claim that non-identity procreation harms the individual while ordinary procreation does not unless she insulates her account from the problem of too much harm. And to avoid the problem of too much harm, she will have to renounce the claim that counterfactual dependence is sufficient for causation.

5.2. Causes and Mere Conditions

In the previous section, I used Harman's account of harming to show that the metaphysics of causation matter to the metaphysics of harming. Specifically, I pointed to the implausible moral

implications of Harman's substantive assumption about the nature of causation. In this section, I will explain why my own account of harming would also face the problem of too much harm if I had to endorse the assumption that something I call *forward dependence* is sufficient for causation. However, I need not assume that forward dependence is a sufficient condition for causation; there are alternative accounts of causation that are compatible with my account of harming. Moreover, at least one of these accounts—namely, Alex Broadbent's (2007a) account—supports a distinction between causes and mere conditions, and this distinction can insulate my account from the problem of too much harm.

I define *forward dependence* as the relation that holds between an event or state of affairs, *a*, and a later event or state of affairs, *b*, just in case if *a* had not obtained or occurred, then *b* would not have obtained or occurred. For example, there is a relation of forward dependence between my birth many years ago and the parking ticket I received two weeks ago: If I had not been born, then I would not have received the ticket.

Forward dependence is a species of counterfactual dependence, which is the relation that holds when one event or state of affairs depends counterfactually upon another. Another species of counterfactual dependence is *reverse dependence*, which is the relation that holds between an event or state of affairs, *a*, and a later event or state of affairs, *b*, just in case if *b* had not obtained or occurred, then *a* would not have obtained or occurred. Suppose that Alice passes her driving test in the morning and gets her driver's license in the afternoon. The claim that if she hadn't gotten her driver's license, she wouldn't have passed her driving test expresses a relation of reverse dependence.

Recall that on Harman's view, counterfactual dependence is sufficient for causation. Since forward dependence is a species of counterfactual dependence, Harman's view implies that forward dependence is sufficient for causation. While I have no objection to the claim that reverse dependence is sufficient for causation, Harman's acceptance of forward dependence as a sufficient condition for causation is what leads to the problem of too much harm.

Notice that if I took forward dependence to be sufficient for causation, then my own account of harming would also face the problem of too much harm. According to the existence account of harming, an action or event harms an individual, S, if and only if that action or event causes a harm for S to obtain, and a harm for S is a state of affairs A such that if S existed and A did not obtain, then S would be better off on the whole. Now, consider the example in which I received a parking ticket. It is true that if I existed and had not received the parking ticket, I would have experienced less mental discomfort in my life, and would therefore have been better off on the whole.⁵⁶ Thus, receiving a parking ticket was a harm for me. If forward dependence is sufficient for causation, then my parents caused me that harm by causing me to exist. After all, I would not have received the parking ticket if I had not been born.

Fortunately, there is no need for me to accept forward dependence as a sufficient condition for causation. There are a number of accounts of causation that do not imply that forward dependence is sufficient for causation. Of course, some of these accounts have other commitments that would still, for other reasons, saddle my view with the problem of too much

⁵⁶ Someone might object that not getting the ticket might have made me worse off. It's possible, for example, that I went to city hall to pay the ticket and met the love of my life. I grant that this possible, but the truth of the claim that the ticket *might* have made me better off on the whole is compatible with the truth of the more plausible claim that the ticket *would not* have made me better off on the whole.

harm. So in addition to rejecting accounts that endorse forward dependence as a sufficient condition for causation, I must endorse only those accounts of causation that would not saddle my view with the implication that everyone is harmed by the action that brings her into existence.

The accounts that would not have this implication, when paired with my claim that harming an individual is causing a harm for her, would need to posit a metaphysical distinction between causes and mere conditions. An event or state of affairs A is a *condition* of another event or state of affairs B if and only if B bears a relation of forward dependence to A or B bears to A a relation grounded in relations of forward dependence. Honore and Hart offer an intuitive example of the distinction between causes and mere conditions: what causes the match to light is the striking of the match, and a mere condition of the lighting is the oxygen in the room. However, Honore and Hart take the distinction to be a mere pragmatic one, as do many other philosophers. If the distinction were merely pragmatic, then the distinction between harm in non-identity cases and harm in ordinary cases would also be pragmatic, and a mere pragmatic distinction is not sufficient for avoiding the problem of too much harm.

There are, however, accounts of causation that take the distinction between causes and mere conditions to be something more than a pragmatic distinction. Notable attempts to justify a metaphysical distinction between causes and mere conditions include attempts by Jonathan Schaffer (2005), Peter Menzies (2004), and Alex Broadbent (2007a).

In principle, any one of these accounts of a metaphysical distinction between causes and mere conditions could be an appropriate supplement to my view of harming. However, the

account I find most plausible is Broadbent's (2007a) account, which posits reverse dependence as one of the conditions on causation.⁵⁷ Broadbent's account also posits two other necessary conditions for causation, but those conditions make a difference only in unusual kinds of preemption cases. I will therefore discuss the following simplified version of the view:⁵⁸

Reverse. For any events C and E, C causes E iff if E hadn't happened, then C wouldn't have happened.

To see how the account works, suppose that Billy throws a rock at a bottle, and two seconds later, the bottle shatters. The throwing happens before the shattering. Nevertheless, if the bottle had not shattered, then Billy would not have thrown the rock. According to Reverse, the fact that there is a relation of reverse dependence between the throwing of the rock and the shattering of the bottle implies that Billy's throwing caused the shattering.

Here is how Reverse helps insulate my solution to the non-identity problem from the problem of too much harm. As I argued in Section 5.1, the solution to the non-identity problem should invoke harm as a wrong-making feature that obtains in non-identity cases but not in ordinary cases. Since harming is causing harm, we will want to say that the action that brings an individual into existence in a non-identity case causes harm to the individual, but the action that brings an individual into existence in an ordinary case *does not* cause harm to the individual.

Now compare the following two cases:

Medical Condition. Xena has a temporary medical condition that will cause any individual she conceives in the next two weeks to have a painful genetic condition. She

57 This account is further defended and elaborated in Broadbent (2007b), Broadbent (2008), and Broadbent (2009).

58 Broadbent's own statement of his view is the following (where the arrow indicates a relation of counterfactual dependence): "C causes E iff $(\sim E \rightarrow \sim C)$ and there is no chain of intermediary events X_1, \dots, X_n , actual and distinct from both C and E, such that $(\sim E \rightarrow \sim X_n), \dots$ and $\sim(\sim X_1 \rightarrow \sim C)$ " (2007a, p. 179).

conceives a child, Bobby, who has the condition. Nevertheless, Bobby has a life worth living.

Ordinary Procreation. Yvette, who does not suffer from any unusual medical conditions, conceives a child whom she names Henry. Henry has a happy childhood. However, when Henry turns 26, he is injured in an industrial accident. The injury causes him to experience pain for the rest of his life. Still, Henry has a life worth living.

Medical Condition is a non-identity case, ordinary procreation is not. Now compare the harm of Bobby's genetic condition to the harm of Henry's accident. Reverse implies that while Xena's procreative action causes Bobby's genetic condition, Yvette's procreative action does *not* cause Henry's industrial accident. That is because the first of the following two counterfactuals is true and the second, false:

(R1) If Bobby had not had the genetic condition, then Bobby would not have existed.

(R2) If Henry had not had the accident, then Henry would not have existed.

R1 is true because any world in which Bobby existed without having the condition would be more distant than various worlds in which Bobby did not exist. The egg and sperm from which Bobby originated contained the genes that caused the condition, so if he didn't have the condition, something about his origins or his living conditions would be strange. This is not to say that it is metaphysically impossible for Bobby to exist without the condition. Indeed, my account of harm implies that the condition is a harm for Bobby *because* it is metaphysically possible for him to exist without it. Nevertheless, the possibility in which he exists and does not have the condition is a remote one. Such remoteness does not matter to the issue of whether the condition is a harm, but it matters to the issue of whether the condition was caused by the action that brought Bobby into existence.

Here is why R2 is false. If Henry had not had the accident, then there are almost any number of possible events before the accident that might have been different. For example, the machinery that injured Henry could have been shut down more quickly, or Henry might have pursued a differently line of work altogether. Certainly, Henry might not have been born; a world in which Henry's parents procreate at a slightly different time and thereby bring another individual into the world instead of Henry seems just as close as a world in which Henry feels ill and decides to stay home from work on the day of the accident. But if the world in which Henry doesn't exist and doesn't have the accident is as close as the worlds in which Henry feels ill, gets help shutting down the machinery, or pursues a different line of work and doesn't have the accident, then it is false that if Henry had not had the accident, he *would* not have been born. What is true is that if he had not had the accident, then he *might* not have been born, and that is a very different inference.

This reasoning about Medical Condition and Ordinary Case generalizes. In any non-identity case, Reverse implies that the action that brings the individual into existence also causes that individual harm. However, in any ordinary case in which an individual's life is worth living, Reverse implies that the action that brings an individual into existence *does not* cause the individual harm.⁵⁹

5.3. Problems for the Comparative Account

I have so far shown that my view need not be subject to the objection I raise against Harman's

⁵⁹ Or rather, it doesn't cause the individual any harm *except death*, but I take the claim that death is the *only* harm that procreation causes in ordinary cases to be an acceptable implication of the view.

account. The reason that my account is safe from the problem of too much harm is that I do not endorse the substantive account of causation that Harman endorses. Instead, I claim that my account will get the right results as long as it is paired with any plausible account of causation that supports a metaphysical distinction between causes and mere conditions. In this section, I will argue that my reliance on some kind of distinction between causes and mere conditions does not put my account of harm at a disadvantage, compared to other accounts of harm. For one, as I have already argued, my account is not at a disadvantage compared to Harman's non-comparative account of harming, for Harman's account faces the problem of too much harm. But neither is my account of harming at a disadvantage when compared to *comparative* accounts of harming. For although the comparative accounts of harming do not face the problem of too much harm, they face a problem that is very similar.

Recall from Chapter 3 that proponents of a comparative account of harming (“comparativists”) accept the following:

(C1) An action or event, A, harms an individual, S, if and only if A causes a state of affairs, W, to obtain, such that W is a harm for S with respect to A.

They also endorse at least one of the following:

(C2) A state of affairs, W, is a harm for an individual, S, with respect to an action or event A if and only if W is worse for S in some respect than some other state, X, such that X obtained immediately before A.

(C3) A state of affairs, W, is a harm for an individual, S, with respect to an event or action A if and only if W is worse for S in some respect than some other state, X, such that X would have obtained if A hadn't happened or been performed.

To get an intuitive understanding of the comparative account, consider a paradigmatic case of

harming: Jones punches Smith in the face, and Smith gets a broken nose. A proponent of C2 would say that Smith's broken nose is a harm for Smith because it is worse for him in some respect than the unbroken nose he had before Jones punched him. A proponent of C3 would say that Smith's broken nose is a harm for Smith because it is worse for Smith in some respect than the unbroken nose Smith would have had if Jones had not punched him. A proponent of C1 and C2 or C1 and C3 would then add that since Jones caused Smith's broken nose, Jones harmed Smith.

Recall that the problem of too much harm is the implication that *every person* is harmed by the action that brings her into existence. This is not an implication of the comparative account because comparativists deny that *anyone* can be harmed by the action that brings her into existence. There is no procreative action that can make the resultant individual worse off than she was before, since if it is a procreative action, the individual did not exist before the action. Likewise, there is no procreative action that can make the resultant individual worse off than she would otherwise have been because, but for the procreative action, she would not have existed at all.

Even though comparativists do not face the problem of too much harm, those who fail to countenance a metaphysical distinction between causes and mere conditions face a similar problem. Consider the following case:

Money for College: Aunt Gracie gives her nephew Leonard \$15,000 for college. Leonard goes to college and does exceedingly well. Nevertheless, at a party celebrating his college graduation, he eats some spoiled potato salad and ends up in the hospital for two days, severely dehydrated from vomiting. If Leonard had not received the money from his aunt, he would not have gone to college, and he would not have eaten the potato salad.

If all conditions are causes, then Aunt Gracie causes the particular episode of food poisoning that Leonard endures. After all, the college money she gives him is a condition of his getting the food poisoning. If she hadn't given him the money, he wouldn't have eaten that particular serving of potato salad.

Given this causal assumption, both versions of the comparative account of harming imply that the food poisoning is a harm for Leonard. To see this, let us use “t1” to refer to the time immediately before Aunt Gracie gives Leonard the money and “t2” to refer to Leonard's second day in the hospital. At t2, Leonard is worse off in some respect than he was at t1. The respect in which he is worse off is his health. Therefore, if all conditions are causes and if C2 is true, the food poisoning that puts in in the hospital is a harm for him.

Similar reasoning holds for C3. Having food poisoning makes Leonard worse off in some respect than he would have been, had Aunt Gracie not given Leonard the money. For if Aunt Gracie hadn't given him the money, he wouldn't have been poisoned by that particular serving of potato salad. Therefore, if all conditions are causes and if C3 is true, the food poisoning is a harm for Leonard.

If the food poisoning is a harm for Leonard, and if Aunt Gracie causes that harm by giving him the college money, then she harms him by giving him the money. But the claim that she harms him by giving him college money is counterintuitive. I will refer to the implication that an action that appears intuitively non-harmful is actually harmful as the problem of *sneaky harms*.

To avoid the problem of sneaky harms, the proponent of the comparative account might try to revise C2 or C3 by replacing the qualifier, “in some respect” with the phrase, “on the whole.” The modified versions of C2 and C3 would be the following:

(C2*) A state of affairs, W, is a harm for an individual, S, with respect to an action or event A if and only if W is worse for S *on the whole* than some other state, X, such that X obtained immediately before A.

(C3*) A state of affairs, W, is a harm for an individual, S, with respect to an event or action A if and only if W is worse for S *on the whole* than some other state, X, such that X would have obtained if A hadn't happened or been performed.

C2* and C3* get the right results in Money for College. Food poisoning is unpleasant, but the benefits of a college education outweigh the unpleasantness of two days in the hospital. Even on day two of his hospital stay, Leonard is no worse off on the whole than he was before Aunt Gracie gave him the college money. He is also no worse off on the whole than he would have been, had she not given him the money. Therefore, in giving him the money, she does not harm him.

However, C2* and C3* run into trouble with another case:

Torture. Felix the torturer kidnaps Steve, tortures him for a month, and then releases him. Steve suffers from post traumatic stress disorder for the rest of his life. Nevertheless, paradoxically, the experience of being tortured made Steve better off on the whole. For as he was being tortured, Steve gradually developed inner strength and deeper compassion for the suffering of others, and this strength and compassion enabled Steve to pursue a successful career as a psychologist, despite his PTSD. He would not have become a psychologist if he had not been kidnapped and tortured by Felix.

If there is no distinction between causes and mere conditions, then Felix's actions cause Steve's successes; if Felix had not done what he did, Steve would not have become a compassionate and successful psychologist. The state of affairs in which Steve has been tortured is better for Steve

on the whole than the state Steve was in before Felix acted, and it is better for him on the whole than the state Steve would have been in, had Felix not kidnapped and tortured him. If so, then when Felix causes this state, he does not harm Steve. The implication that Felix's actions of kidnapping and torture are were harmful is counterintuitive. I call the implication that actions that are intuitively harmful are *not harmful* the *problem of sneaky benefits*.

If the comparativist holds that all conditions are causes, then she will face either the problem of sneaky harms or the problem of sneaky benefits. Thus, the fact that my own account of harming relies upon a metaphysical distinction between causes and mere conditions does not put my account at a comparative disadvantage.

5.4. Conclusion

In this chapter I have identified a problem for Elizabeth Harman's comparative account of harming. I called this problem, “the problem of too much harm,” and I argued that it would also threaten my own account of harming if I accepted Harman's assumption that forward dependence is sufficient for causation. However, I argued that I do not have to accept such an assumption. There is a plausible alternative account of causation that holds that forward dependence is not sufficient for causation, although reverse dependence is. I showed that this alternative account—a modified version of Broadbent's (2007a) account—supports the claim that individuals in non-identity cases are harmed whereas individuals in ordinary procreation cases are not. More generally, I suggested that any account of causation that supports a metaphysical distinction between causes and mere conditions can insulate my account of harming from the problem of too

much harm. Finally, I argued that the reliance of my account on a distinction between causes and mere conditions is not a relative weakness of my account, for even *comparative* accounts of harming face problems when they do not distinguish between causes and mere conditions.

Chapter 6: Benefiting by Causing Existence

6.0. Introduction

I have so far argued for DV1, the claim that we have duties of beneficence and nonmaleficence to future humans, and for DV2, the claim that in the present circumstances, the duty of nonmaleficence grounds a reason for us to refrain from damaging the environment in certain ways. In this chapter, I will defend DV3, the claim that in the present circumstances, the duty of beneficence we owe to future humans grounds a reason for at least some of us to bring future humans into existence. Whereas I have been referring to a reason against harming someone that is grounded in a duty of nonmaleficence to that person an “fh-reason,” I will call a reason to benefit someone that is grounded in a duty of beneficence to that person an “fb-reason.” In section 6.1 I will present my main argument for the claim that there is an fb-reason for some of us to bring future humans into existence. I will consider and respond to two objections. The second objection is inspired by what I call the “non-identity benefit problem,” and I will provide an account of benefiting that solves this problem. In section 6.2 I will show how my account of benefiting can vindicate the claim that some of us can benefit future humans by causing them to exist. In section 6.3, I will consider and respond to objections.

6.1. The Argument

My main argument for DV3 is the following:

- (B1) We have a pro tanto duty to some future humans to benefit them.
- (B2) In the present circumstances, some of us can benefit some of those humans by bringing them into existence.

(DV3) Therefore, in the present circumstances, there is an fb-reason for us some of us to bring those humans into existence.

I have already argued for something like B1 in Chapter 1. Recall that my main argument for the claim that entails B1 is the following:

The Argument from Moral Standing

- (M1) We owe all individuals who have moral standing pro tanto duties of nonmaleficence and beneficence.
- (M2) At least some future humans have moral standing.
- (DV1) Therefore, we owe at least some future humans pro tanto duties of nonmaleficence and beneficence.

Thus, in this chapter, I will focus upon premise B2, which claims that some of us can benefit some of the future humans to whom we owe duties of beneficence by bringing them into existence.

I will argue for B2 in stages. First, I will show that it is possible, in principle, to benefit future humans. I will then argue that if it is possible, in principle, to benefit future humans, then some of us can benefit future humans by performing an action that causes both the benefit to those humans and the existence of those humans. Finally, I will argue that if some of us can benefit future humans by performing an action that causes both the benefit and their existence, then some of us can also benefit future humans *by causing their existence*.

One preliminary challenge to the claim that it is possible, in principle, to benefit future humans comes from an account of benefiting that is analogous to the account of harming that accepts C1 and C2. Recall that C1 and C2 are the following:

- (C1) An action or event, A, harms an individual, S, if and only if A causes a state of affairs, W, to obtain, such that W is a harm for S with respect to A.

(C2) A state of affairs, *W*, is a harm for an individual, *S*, with respect to an action or event *A* if and only if *W* is worse for *S* in some respect than some other state, *X*, such that *X* obtained immediately before *A*.

The combination of these claims implies that an action harms an individual only if it makes her worse off in some respect than she was before the action. There is a similar account of benefiting that consists of the following:

(CB1) An action or event, *A*, benefits an individual, *S*, if and only if *A* causes a state of affairs, *W*, to obtain, such that *W* is a benefit for *S* with respect to *A*.

(CB2) A state of affairs, *W*, is a benefit for an individual, *S*, with respect to an action or event *A* if and only if *W* is better for *S* in some respect than some other state, *X*, such that *X* obtained immediately before *A*.

CB1 and CB2 imply that an action benefits an individual only if it makes her better off in some respect than she was before the action. Now, we are supposing that the term “future humans” refers to individuals who do not exist at the time that our actions are performed.⁶⁰ If so, then no future humans exist *before* we act, either. Thus, for any action we take that affects future humans, it is necessarily false that such humans were worse off in some respect immediately before the action and better off afterward. If CB1 and CB2 are correct, then no action we take can benefit future humans.

However, CB1 and CB2 are false. To see why, consider the following case:

Trump's Offer. Donald Trump sets aside \$5 million in a special account, with the stipulation that the money be divided among all the persons born in Madison,

⁶⁰ I'm assuming that our actions are events. Certainly, those events can be described in various ways, but it is a category mistake to ask for the time of a description of an action. For example, by pulling the trigger, Jones kills Smith. “Jones's killing Smith” is a description of Jones's action, as is “Jones's causing Smith to die.” But while Jones's action—the pulling of the trigger—takes place at a time, his killing of Smith does not. See Weintraub (2003), who holds that killings are causings. She then writes, “When does X cause Y?” is nonsensical, as are ‘When does X precede Y?’, and ‘When does X last longer than Y?’ (p. 181).

Wisconsin in the first week of February, 2015.

By setting aside this money today, Trump benefits the individuals who will be born in Madison in 2015. The example shows that the conjunction of CB1 and CB2 is false, for an action can benefit someone who does not exist at the time the action is performed.

If it is possible, in principle, to benefit future humans, then there are some cases in which some of us can benefit them by performing an action that also causes their existence. Consider, for example, the following:

Successful Timing. The Bigglesworths live in Madison. Knowing about Donald Trump's fund, they try to conceive in May of 2014 so that their child will be born in the first week of February in 2015. They are successful, and their child, Bieber, receives \$24,000.

Successful Timing is an example of a procreative action with double effects. The Bigglesworths' action—namely, procreating in May of 2013—brings Bieber into existence. The same action also benefits Bieber by causing him to receive some of Donald Trump's money. Although I have used a dramatic example, the kind of action that the Bigglesworths perform is in principle very similar to actions that people perform every day. There are lots of benefits that come to people as a result of when they were born, such that the actions that caused those people to be born when they were born caused those people to receive those benefits. Thus, there are some cases in which some of us can benefit future humans by performing an action that also causes their existence.

At this stage of my argument, someone might wish to object that there is a problem with the claim that some actions can both benefit individuals and cause their existence. I will refer to this problem as *the non-identity benefit problem*. To see the problem, notice that Bieber is no better off in any respect, as a result of the Bigglesworths' procreative act, than he would have

been, had they not performed that action. After all, had the Bigglesworths not conceived Bieber when they did, Bieber would not have *been* at all: either the Bigglesworths would not have conceived a child, or else they would have conceived a child at a different time. Since a child conceived at a different time would have come from a different egg and sperm, it would not have been Bieber. But if Bieber is no better off than he would have been, had his parents not acted as they did, then it is difficult to see how their performing that action benefited Bieber. The non-identity benefit problem is the problem of accounting for the intuition that, despite such reasoning, the Bigglesworths *have* benefited Bieber.

Just as one of the proposals for solving the non-identity problem motivates an objection to DV2, one proposal for solving the non-identity benefit problem motivates an objection to my reasoning about Successful Timing. The proposal is that we should reject the claim that Bieber has been benefited and then somehow explain away the intuition to the contrary.

I take it that this proposal is motivated by the following account of benefiting:

(CB1) An action or event, A, benefits an individual, S, if and only if A causes a state of affairs, W, to obtain, such that W is a benefit for S with respect to A.

(CB3) A state of affairs, W, is a benefit for an individual, S, with respect to an action or event A if and only if W is better for S in some respect than some other state, X, such that X would have obtained if A had not occurred or been performed.

Together, CB1 and CB3 imply the counterfactually better-off condition on benefiting, which is the claim that an action, A, benefits an individual, S, only if A makes S better off in some respect than S would have been, had A not been performed.

However, just as I argued in Chapter 3 that the solution to the non-identity problem is to

reject the counterfactually worse-off condition on harming, so I now argue that the solution to the non-identity benefit problem is to reject the counterfactually better-off condition on benefiting. There are at least three reasons to opt for this approach to the non-identity benefit problem over the approach that rejects the claim that Bieber was benefited. The first is a symmetry consideration: since harming is the negative counterpart of benefiting, we should, other things being equal, seek an account of benefiting that mirrors our account of harming. The second reason is that, just as there are examples of overdetermined or preempted harming, so there are examples of overdetermined or preempted benefiting, such as the following preemption case:

Glasses. Alice is leaning over the balcony, and her glasses slip from her face. Both Barbara and Candice are on the sidewalk below. Barbara catches the glasses, but if she hadn't caught them, Candice would have.

In *Glasses*, Alice is no better off in any respect than she would have been, had Barbara not acted as she did. Nevertheless, Barbara has benefited Alice.

The third reason that supports rejecting the counterfactually better-off condition on benefiting has to do with cases in which there is no fact of the matter about how well-off an individual would have been if an intuitively beneficial action had not been performed. Consider the following:

No Fact of the Matter 2: Two contestants on a game show, Bloggs and Pinkerton, are placed in separate booths. Above each booth is an air vent, and each air vent leads to a giant chamber in which a fan is blowing around a whirlwind of \$100 bills. Over time, the bills are gradually being blown into both booths. A computer is programmed to select a three-digit number at random at one second before noon. If the number is even or zero, the computer will close the air vent to Bloggs's cell; if the number is odd, the computer will close the air vent to Pinkerton's cell. The random-number selection process is truly

indeterministic. However, the host of the game show, Weatherby, is feeling particularly generous, and so just before the number is about to be selected, he shuts down the computer. The bills continue to blow into both booths until the fan is turned off and taping ends for the day.⁶¹

Intuitively, Weatherby's action benefits both Bloggs and Pinkerton. Nevertheless, there is no fact of the matter about how well-off Bloggs would have been, had Weatherby not unplugged the computer. Nor is there a fact of the matter about how well-off Pinkerton would have been, had Weatherby not done what he did. Thus, it is not the case that Weatherby's action makes either Bloggs or Pinkerton better off than he otherwise would have been.

In place of the counterfactually better-off condition on benefiting, I endorse the following:

The Existence Account of Benefiting:

Benefiting as Causing (BC): An action or event benefits an individual, S, if and only if that action or event causes a benefit for S to obtain.

The Inevitable Benefiting Principle (IBP): Other things being equal, the reason to cause a benefit for an individual to obtain is stronger, the more inevitable the benefit.⁶²

The Existence Account of Benefit (EB): A state of affairs, A, is a benefit for an individual, S, just in case if it were true that both S existed *and* A did not obtain, then S would have a lower level of lifetime well-being.

The Existence Account of Benefiting implies that in Successful Timing, the Bigglesworths' action of procreating in May of 2014 benefited Bieber. That is because, in accordance with EB, the state of affairs that consists in Bieber's having an extra \$24,000 is a benefit: if it were true

61 I have, of course, modeled this case upon Norcross's (1990) case of the indeterminate poisoning, which I presented in Chapter 3.

62 I include this principle because it explains why there is more reason to bring into existence an intelligent human than there is to bring into existence a goat who will be equally intelligent. However, nothing in my argument for DV3 turns upon this principle, so I include it here only for completeness.

that Bieber existed and did *not* have the extra \$24,000, then he would have a lower lifelong level of well-being. Second, the Bigglesworths' action of procreating in May of 2013 caused Bieber to receive the extra \$24,000; since they caused a benefit for Bieber, BC implies that they benefited him. But some of us are in a position like the Bigglesworths, since some of us can cause the existence of humans who will incur benefits associated with the time at which they are born. If so, then some of us can perform actions that both cause the existence of future humans and cause benefits for those humans.

I have so far argued that, in principle, it is possible to benefit future humans. Moreover, I have argued that some of us can perform actions that both benefit future humans and bring those future humans into existence. What I will argue in the next section is that some of us can benefit future humans *by* bringing them into existence.

6.2. Benefiting by Causing Existence

In this section I will argue that some of us can benefit future humans by bringing them into existence. Notice that this claim does not follow straightaway from the claim that some of us can perform actions that cause both the existence of and benefits to future humans. From the fact that I can perform an action that causes X and Y, it does not follow that I can cause X *by* causing Y, just as, from the fact that I can take a pill that causes pain relief and tinnitus, it does not follow that I can take a pill that causes pain relief *by* causing tinnitus. Nevertheless, the claim that I can benefit someone by causing her to exist *does* follow from the claims that (1) I can benefit her by causing her a benefit, (2) a life worth living is a benefit for her, and (3) I can cause this person to

have a life worth living by causing her to exist. If we assume that I can cause someone to exist, then the first of these claims follows from the conjunction of that assumption with Benefiting as Causing.

Next, let me generalize claim 2: For anyone who has a life worth living, having a life worth living is a benefit for her. This claim follows from the Existence Account of Benefit. To see why, suppose that an individual, P, has a life worth living. Call the state of affairs in which she has a life worth living A. Now consider this counterfactual:

(C) If P existed and A did not obtain, then P would have a lower level of lifetime well-being.

C is true, for if P existed and did not have a life worth living, then P would have a lower level of lifetime well-being; total lifetime well-being is very low when one does not have a life worth living. But if C is true, then the Existence Account of Benefit implies that for anyone who has a life worth living, having a life worth living is a benefit for her.

What about claim 3? Claim 3 says that I can cause a person to have a life worth living by causing her to exist. This is a causal claim, and as long as we assume that I can cause a person whose life is worth living to exist, then any plausible account of causation should yield the result that 3 is true. For example, since I have already endorsed a simplified version of Alex Broadbent's (2007a) account of causation, I will show how his account gets this result.

Recall from Chapter 5 that on a simplified version of Broadbent's view, the following is true:

Reverse. For any events C and E, C causes E iff if E hadn't happened, then C wouldn't have happened.

Now let us consider an individual, P, and suppose that P's having a life worth living is one event, E, and the action that causes P to exist is another event, C. Is it true that if E hadn't happened, then C wouldn't have happened? Consider all the possible worlds in which E doesn't happen. In some of those worlds, P doesn't exist, and in others, P exists. In the worlds in which P exists but does not have a life worth living, much misfortune befalls her. Either she suffers from lifelong pain, she is rejected by everyone she knows, she fails at all she sets out to do, she is in a permanent vegetative state, or she suffers some combination of these or other misfortunes. I contend that *any* combination of these misfortunes would be rare. The commonsense view is that in the actual world, very few individuals have lives that are not worth living. The kind of combination of misfortunes that would have to obtain to render a life not worth living would be *more unusual* than a state of affairs in which P's parents procreated at a different time. Thus, the worlds in which E doesn't happen and P doesn't exist are closer than the worlds in which E doesn't happen and P has a life that is not worth living.

Such reasoning suggests that the following counterfactual is true: If E hadn't happened, then C wouldn't have happened. But if that counterfactual is true, then according to Reverse, an action that causes an individual to exist also causes her to have a life worth living. Thus, if I can cause an individual whose life is worth living to exist, then by causing that individual to exist, I can cause her to have a life worth living.

To summarize, I have just argued that if I can cause a person whose life is worth living to exist, then (1) I can benefit her by causing her a benefit, (2) a life worth living is a benefit for

her, and (3) I can cause this person to have a life worth living by causing her to exist. If all this is so, then I can benefit a person whose life is worth living by causing her to exist. But there is nothing special about me; other presently living people can cause the existence of future humans whose lives will be worth living. If so, then in the present circumstances, some of us can benefit some future humans by bringing them into existence. This completes my argument for premise B2 of my main argument for DV3.

6.3. Objections and Replies

I have argued that since we owe a pro tanto duty of beneficence to some future humans, and since some of us can benefit those humans by bringing them into existence, there is an fb-reason for some of us to bring those humans into existence. In this section, I will consider and respond to two objections.

The first objection is raised by Jonathan Bennett, who argues for the following principle:

The question of whether action A is morally obligatory depends only upon the utilities of people who would exist if A were not performed (p. 62).

Bennett's principle implies that it is never morally obligatory to bring someone into existence.

For if we failed to bring someone into existence, such a person would not exist, and so her utility would not count against the action.

In support of his principle, Bennett offers the following reasoning:

[I]f a failure to bring someone into existence is ever wrong for utilitarian reasons, these must concern the utilities of people who are at some time actual, not those of the person whose coming-into-existence didn't happen. It might be wrong for me to fail to beget a child because that would deprive my parents of the pleasures of grand-parenthood, or

because any child of mine would be sure to benefit mankind; in one case my parents are deprived, in the other mankind in general. But it couldn't be wrong because by not bringing the child into existence one deprives *it* of something (p. 61-62).

Although Bennett has formulated his argument in utilitarian terms, similar reasoning could be directed against my non-utilitarian view. One might observe that I am defending an fb-reason to bring some future humans into existence. If there is such a reason, then there must also be a possible case in which someone fails to act upon that reason without sufficient justification for that failure. Presumably, failing to act upon a reason in such circumstances would wrong someone. Yet there is no one whom it would wrong. Therefore, the supposition that there is an fb-reason to bring someone into existence implies a contradiction.

In response to this objection, I reject the assumption that there would be an fb-reason for an agent to bring a future human into existence even if the agent were not going to do it. Fb-reasons to bring humans into existence are therefore different from other sorts of reasons. Whereas other sorts of reasons can apply to merely possible actions, fb-reasons to bring humans into existence apply only to actual actions. This is simply a function of the fact that fb-reasons to procreate are owed to *actual* future humans, and not to merely possible ones.

A second objection to my claim appeals to the practical consequences of bringing future humans into existence. The objection points out that bringing another human into existence is a way of dramatically increasing one's carbon footprint. Since I have already argued that there are fh-reasons not to contribute to the carbon dioxide in the atmosphere, it would seem that these fh-reasons count against bringing future humans into existence. If both DV2 and DV3 were true, then surely some of these fh-reasons against bringing humans into existence would be grounded

by the moral standing of exactly the same individuals whom we would have fb-reasons to bring into existence. There must be, for example, at least one future human who is such that there is a nonmaleficence reason against bringing him into existence and a beneficence reason to bring him into existence. The objection is that this appears to be an incoherent implication of my view.

In response, I grant that it is likely that there is at least one future human such that there is a harm-based reason against causing him to exist and a beneficence-based reason to cause him to exist. Nevertheless, this is not an incoherent implication of my view. There are many situations in which we can have reasons for and against an action that are grounded by our duties to the same individual. For example, if we are lost together in the wilderness and I have nothing but a package of processed snack foods to share with you, then I have a harm-based reason not to share the food with you and a beneficence-based reason to share it with you. After all, the food contains preservatives and other harmful ingredients that are bad for your health, but you are hungry and have nothing else to eat, so the food would provide you with needed calories. Our situation with respect to future humans is similar. If we bring them into existence, then we will exacerbate the harms of climate change to which these humans will be subject, but we will also cause them the benefit of lives worth living. In this dissertation, I have not provided a calculus for how to weigh these reasons against each other, but I believe it is useful, as a first step, to know what the reasons are. By establishing that there are fh-reasons and fb-reasons against damaging the environment in certain ways and in favor of bringing future humans into existence, DV clarifies some of the moral considerations that are relevant to our future-regarding decisions.

Bibliography

- Adams, R. (1972). "Must God Create the Best?" *The Philosophical Review*, vol. 81, no. 3, pp. 317–332.
- Benatar, D. (2006). *Better Never to Have Been: The Harm of Coming into Existence* (Oxford: Oxford University Press).
- Bennett, J. (1978). "On Maximizing Happiness," in *Obligations to Future Generations*, R. Sikora and B. Barry (eds.), Philadelphia: Temple University Press, pp. 61–73.
- Boonin, D. (2008). "How to Solve the Non-Identity Problem," *Public Affairs Quarterly*, vol. 22, no. 2, pp. 127 – 157.
- Bourne, C. (2006). *A Future for Presentism*. Oxford: Oxford University Press.
- Broad, C. D. (1923). *Scientific Thought*. (New York: Harcourt, Brace and Co.)
- Broadbent, A. (2007a). "Reversing the Counterfactual Analysis of Causation," *International Journal of Philosophical Studies*, vol. 15, pp. 169 – 189.
- . (2007b). *A Reverse Counterfactual Analysis of Causation*. (Doctoral dissertation). University of Cambridge.
- . (2008). "The Difference Between Cause and Condition," *Proceedings of the Aristotelian Society*, vol. 108, pp. 355 – 364.
- . (2009). *No Fire Without Smoke: Causation, Counterfactuals, and How to Make a Difference*. (Unpublished book manuscript).
- Burgess-Jackson, K. (1998). "Doing Right By Our Animal Companions," *The Journal of Ethics* 2: 159 – 185.
- Collins, J., N. Hall, and L.A. Paul, eds. (2004). *Causation and Counterfactuals*. Cambridge, MA: MIT Press.
- Doerflinger, R. M. (1999). "The Ethics of Funding Embryonic Stem Cell Research: A Catholic Viewpoint," *Kennedy Institute of Ethics Journal*, vol. 9, no. 2, pp. 137 – 150.
- Earl, D. (2011). "Ontology and the Paradox of Future Generations," *Public Reason* 3(1): 60 – 72.

- Feinberg, J. (1974). "The Rights of Animals and Future Generations," in William Blackstone (ed.), *Philosophy and Environmental Crisis*. Athens, GA: University of Georgia Press.
- . (1984). *Harm to Others* (Oxford: Oxford University Press).
- Forrest, P. (2004). "The Real But Dead Past: A Reply to Braddon-Mitchell," *Analysis* 64: 358 – 62.
- Gardner, A. (2009). "Being Near Nature Improves Physical, Mental Health," *USA Today* Accessed Jan. 15, 2013.
http://usatoday30.usatoday.com/news/health/2009-10-15-nature-anxiety-exercise_N.htm
- Gardner, M. (2013). "Misfortune and Inevitability" (Unpublished manuscript).
- Gardner, M. and J. Weinberg. (2013). "How Lives Measure Up," *Acta Analytica* 28(1), pp. 31 – 48.
- Gardiner, S.M. (2011). *A Perfect Moral Storm: The Ethical Tragedy of Climate Change*. (Oxford: Oxford University Press).
- Gollier, C. (2007). "The Economics of Climate Change: A Critical Analysis of the Stern Review," <http://www.oecd.org/eco/greeneco/40133781.pdf> Accessed Jan. 15, 2013.
- Greene, M. and S. Augello (2011). "Everworse: What's Wrong with Selecting for Disability?" *Public Affairs Quarterly*, vol. 25, no. 2, pp. 131 – 139.
- Hanser, M. (2008). "The Metaphysics of Harm," *Philosophy and Phenomenological Research*, vol. 77, no. 2 (2008), pp. 421 – 450.
- . (2009). "Harming and Procreating," in *Harming Future Persons: Ethics, Genetics and the Nonidentity Problem*, ed. Melinda Roberts and David Wasserman (Dordrecht: Springer), pp. 179 – 199.
- Hare, C. (2007). "Voices from Another World: Must We Respect the Interests of People Who Do Not, and Will Never, Exist?" *Ethics*, 117: 498–523.
- Harman, E. (2004). "Can We Harm and Benefit in Creating?" *Philosophical Perspectives* 18, pp. 89 – 113.
- . (2009). "Harming as Causing Harm," in *Harming Future Persons: Ethics, Genetics and*

- the Nonidentity Problem*, ed. Melinda Roberts and David Wasserman (Dordrecht: Springer, pp. 137 – 154.
- Hart, H.L.A. and T. Honore (1985). *Causation in the Law*, 2nd Ed. (Oxford: Clarendon Press).
- Heyd, D. (2009). “The Intractability of the Nonidentity Problem,” in *Harming Future Persons: Ethics, Genetics and the Nonidentity Problem*, ed. Melinda Roberts and David Wasserman (Dordrecht: Springer), pp. 3–25.
- International Energy Agency. (2012). “Global Carbon Dioxide Emissions Increase by 1.0 Gt in 2011 to Record High,” Accessed January 15, 2013.
<http://www.iea.org/newsroomandevents/news/2012/may/name,27216,en.html>
- IPCC. (2007). Fourth Assessment Report: Climate Change (AR4).
- International Union for Conservation of Nature (2013). “Why is Biodiversity in Crisis?” Accessed January 15, 2013.
http://www.iucn.org/what/biodiversity/about/biodiversity_crisis/
- Kamm, F. M. (1999). “Famine Ethics: The Problem of Distance in Morality and Singer's Ethical Theory,” in Dale Jamieson, ed. 1999. *Singer and His Critics*. Malden, MA: Blackwell, pp. 161 – 208.
- Kavka, G. (1982). “The Paradox of Future Individuals,” *Philosophy & Public Affairs*, vol. 11, pp. 93–112.
- Lomborg, B. (2001). *The Skeptical Environmentalist*. (Cambridge: Cambridge University Press.)
- Maas, J., R.A. Verheij, S. de Vries, P. Spreeuwenberg, F.G. Schellevis, and P.P. Groenewegen (2009). “Morbidity is Related to a Green Living Environment.” *J Epidemiol Community Health*, vol. 63, pp. 967 – 973.
- McGinn, C. (1999). “Our Duties to Animals and the Poor,” in Dale Jamieson, ed. 1999. *Singer and His Critics*. Malden, MA: Blackwell, pp. 150 – 161.
- McMahan, J. (1996). “Cognitive Disability, Misfortune, and Justice,” *Philosophy & Public Affairs*, vol. 25, no. 1, pp. 3 – 35.
- . (2009). “Asymmetries in the Morality of Causing People to Exist,” in *Harming Future Persons: Ethics, Genetics and the Nonidentity Problem*, ed. Melinda Roberts and David Wasserman (Dordrecht: Springer), pp. 49–68.

- Menzies, P. (2004). "Difference-Making in Context," in *Causation and Counterfactuals*, J. Collins, N. Hall, and L.A. Paul (eds.), pp. 139 – 180.
- Moore, G. E. (1903) *Principia Ethica* (Dover, reprinted in 2004).
- Nagel, T. (1970). "Death," *Noûs*, vol. 4, no. 1, pp. 73 – 80.
- Norcross, A. (1990). "Killing, Abortion, and Contraception: A Reply to Marquis," *The Journal of Philosophy* 87(5), pp. 268 – 277.
- (2005). "Harming in Context," *Philosophical Studies* 123, pp. 149–173.
- Parfit, D. (1976). "On Doing the Best for Our Children," in *Ethics and Population*, ed. Michael Bayles (Cambridge, Mass.: Schenkman), pp. 100 – 102.
- . (1984). *Reasons and Persons* (Oxford: Oxford University Press).
- Pitcher, G. (1984)., "The Misfortunes of the Dead," *American Philosophical Quarterly*, vol. 21, no. 2, pp. 183 – 188.
- Regan, T. (2004). *The Case for Animal Rights* (University of California Press).
- Roberts, M., and D. Wasserman, eds. (2009). *Harming Future Persons: Ethics, Genetics and the Nonidentity Problem*. Dordrecht: Springer.
- Ryberg, J. and T. Tännsjö, eds. (2004). *The Repugnant Conclusion. Essays on Population Ethics* (Dordrecht: Kluwer Academic Publishers).
- Schaffer, J. (2005). "Contrastive Causation," *Philosophical Review*, vol. 114, no. 3, pp. 297 – 328.
- Schwartz, T. (1978). "Obligations to Posterity," in *Obligations to Future Generations*, ed. Richard Sikora and Brian Barry (Philadelphia: Temple University Press), pp. 3 – 13.
- Shiffrin, S. (1999). "Wrongful Life, Procreative Responsibility, and the Significance of Harm," *Legal Theory*, vol. 5, pp. 117–48.
- Sider, T. (2003), *Four-Dimensionalism: An Ontology of Persistence and Time*. Oxford: Oxford University Press.

- Sikora, R. and B. Barry, eds. (1978). *Obligations to Future Generations*. Philadelphia: Temple University Press.
- Singer, P. (1972). "Famine, Affluence, and Morality," *Philosophy and Public Affairs*, vol. 1, no. 1 pp. 229-243.
- (2009). *The Life You Can Save: Acting Now to End World Poverty*. New York: Random House.
- Slote, M. (2005). "Famine, Affluence, and Empathy," in David Boonin and Graham Oddie, eds. *What's Wrong? Applied Ethicists and Their Critics*, 2nd Edition. New York: Oxford University Press, pp. 146 – 154.
- Smolkin, D. (1994). "The Non-Identity Problem and the Appeal to Future People's Rights," *The Southern Journal of Philosophy*, vol. 32, pp. 315 – 329.
- Thomson, J. (2011). "More on the Metaphysics of Harm," *Philosophy and Phenomenological Research*, vol. 82, no. 2, pp. 436 – 458.
- Tooley, M. (1997). *Time, Tense, and Causation*. (Oxford: Oxford University Press).
- Weintraub, R. (2003). "The Time of a Killing," *Analysis* vol. 63, no. 279, pp. 178 – 182.
- WHO. (2009). "Ch. 2, Results: 2.6 Environmental Risks," *Global Health Risks: Mortality and Burden of Disease Attributable to Selected Major Risks*. Geneva, Switzerland: WHO Press, p. 24.
- Wilson, E.O. (2002). *The Future of Life*. (New York: Alfred A. Knopf).
- Wolpert, S. (2009). "Last Time Carbon Dioxide Levels Were This High: 15 Million Years Ago, Scientists Report," UCLA Newsroom.
<http://newsroom.ucla.edu/portal/ucla/last-time-carbon-dioxide-levels-111074.aspx>
Accessed January 15, 2013.
- Woodward, J. (1986). "The Non-Identity Problem," *Ethics*, 96: 804–31.
- Woollard, F. (2012). "Have We Solved the Non-Identity Problem?" *Ethical Theory and Moral Practice*, vol. 15, no. 5, pp. 677 – 690.
- World Wildlife Fund (2013). "How Does Biodiversity Loss Affect Me and Everyone Else?"
http://wwf.panda.org/about_our_earth/biodiversity/biodiversity_and_you/