

# University of Alberta

Aretaic Assessment and the Non-Identity Problem

by

Keith Underkoffler

A thesis submitted to the Faculty of Graduate Studies and Research  
in partial fulfillment of the requirements for the degree of

Master of Arts

Philosophy

©Keith Underkoffler

Fall 2013

Edmonton, Alberta

Permission is hereby granted to the University of Alberta Libraries to reproduce single copies of this thesis and to lend or sell such copies for private, scholarly or scientific research purposes only. Where the thesis is converted to, or otherwise made available in digital form, the University of Alberta will advise potential users of the thesis of these terms.

The author reserves all other publication and other rights in association with the copyright in the thesis and, except as herein before provided, neither the thesis nor any substantial portion thereof may be printed or otherwise reproduced in any material form whatsoever without the author's prior written permission.

## **Dedication**

I've always been interested in ethical questions about how to live and be a good person. I'm not sure I'm any closer to arriving at any answers, but I'm grateful for the opportunity to give it some thought.

To my parents, for their endless support (both emotional and financial), and for not laughing when I said I wanted to pursue a Master's degree in Philosophy. To my friends, especially Megan, Octavian, Hande, Nika, Grace, Justin, James, Jay, Andrew and Yasemin, for making the past two years so great.

## **Abstract**

The Non-Identity Problem arises in cases in which an agent causes someone to exist in an undesirable state, and it intuitively seems they have acted wrongly by doing so. It can be difficult to locate the source of this wrongdoing, however, since in these cases there is no way for the agent to cause the future person to come into a more desirable existence. If the agent had acted otherwise, the future person simply would not have existed, so the agent's action was not worse *for them*. In light of this, it is natural to revisit the intuition that the agent acted wrongly in order to see what the precise nature of the original disapprobation was. In this thesis, I argue that it may not be impermissible to bring the future person into existence, but it may be morally bad, depending on the reasons for which the agent acted.

## **Acknowledgements**

First and foremost, I'd like to thank my committee for their time and support throughout this process. I'm especially indebted to my supervisors Howard Nye and Jennifer Welchman. Howard provided extensive, challenging, and detailed comments (both written and verbal). Jenny provided foresight, clarifying remarks, and much needed structure. I am grateful to them both, as this thesis would not exist without either of them. Thanks also to Glenn Griener and David Boonin for their interest and participation in this project.

For academic support and engagement with this work I'd like to thank my colleagues in the Philosophy Department's Publication Support Group and an audience at the 2013 meeting of the Canadian Philosophical Association Annual Congress in Victoria, BC. Attendance of that conference was supported by a J Gordin Kaplan Graduate Student Award, for which I am very grateful. Thanks also to Octavian Ion and Megan Dean for helpful written comments on earlier drafts of this work and to Jeff McMahan and Julia Driver for humoring me in discussion about this project.

Finally, thanks to the Department of Philosophy, University of Alberta, FGSR, Chloë Taylor, and David and Barb Underkoffler for their financial support of my studies, and to Sussanne McDonald for her invaluable administrative support and kind words.

## Table of Contents

<b>Introduction: The Non-Identity Problem and Procreative Decisions.....</b>	<b>1</b>
<b>Chapter One: Outcome-Based Considerations.....</b>	<b>9</b>
Harm-Based Considerations.....	10
Impersonal Considerations.....	22
Rights-Based Considerations.....	33
Accepting the Implausible Conclusion.....	42
<b>Chapter Two: Intention-Based Considerations.....</b>	<b>49</b>
Respect-Based Considerations.....	50
Role-Based Considerations.....	59
Intentions and Permissibility.....	65
<b>Chapter Three: Aretaic Assessment and the Non-Identity Problem.....</b>	<b>70</b>
Deontic Assessment and Wrongdoing.....	71
Beliefs and Blameworthiness.....	77
Motivation and Moral Badness.....	85
Implications.....	88
<b>Conclusion: .....</b>	<b>91</b>
<b>Bibliography.....</b>	<b>92</b>

## Introduction: The Non-Identity Problem and Procreative Decisions

First discovered in the late 1970's by Derek Parfit, Thomas Schwartz, and Robert M. Adams, the Non-Identity Problem has generated a large body of literature and gained a reputation as one of the most difficult problems in normative ethics.<sup>1</sup> The problem starts with the recognition that our actions have the ability to affect not only the quality of life that people will experience in the future, but also the identities of the people who will exist. It seems intuitive to many of us that this ability carries with it an ethical obligation to choose to bring some people into the world and not others. Consider, for example, the following case:

*The 14 Year Old Girl:* a fourteen year old girl (call her Alma) decides to have a child (Beatrice) while foreseeing that, because of Alma's youth, Beatrice will have a "bad start" in life. Though this will have bad effects throughout Beatrice's life, her life will, predictably, be "worth living." It is also true that, if Alma waits until she is older, she will have a *different* child, who will have a much better start in life (cf. Parfit 1984, 358, and Woodward, 805).

Most people share the intuition that it would be objectionable for Alma to give Beatrice a bad start in life now when she could easily give a different child a good start in life instead. But it is difficult to identify anyone who would be made worse off by her action. This, in turn, makes it difficult to vindicate the intuition that choosing to have Beatrice now would be objectionable. If Alma waits until she is older to have a child, a different combination of sperm and egg will come together and result in a different child. So, the only life available to Beatrice is the one she will have if Alma brings her into existence now. As long as this life is not

---

<sup>1</sup> See Derek Parfit (1976; 1982; 1984), Schwartz (1978), and Adams (1979).

worse than a state of non-existence, it does not seem that she would be made worse off by being born. The only other possible victims are Alma or the possible child she could have if she waits. But failing to bring possible children into existence doesn't seem to be problematic. We would not object, for example, if Alma decided to have no child at all. And the claim that Alma makes herself worse off seems to miss the core of our objection. There seems to be more to our objection than the claim that Alma is making a bad decision about *her own* future. But if we cannot identify anyone else who would be made worse off by Alma's action, it isn't easy to see what else this might be.

The problem is that most people seem to have fairly resilient intuitions that there is something objectionable going on in cases like The 14 Year Old Girl, but these intuitions conflict with a number of theoretical commitments most ethicists are reticent to give up. To resolve the problem, then, it seems we will either have to revise our theoretical commitments or give up on our specific moral judgments. This thesis is an attempt to work out that tension. It is situated within a common methodology which takes intuitions about cases as an important piece of data in ethical theorizing, and works back and forth between specific case intuitions and general ethical principles – sometimes revising the intuitions, other times modifying the principles – in an effort to achieve a 'reflective equilibrium', in which all of the intuitions one appeals to withstand philosophical scrutiny and all of the normative principles one invokes are independently plausible.

Accordingly, the main theme of this thesis will be: how can we best capture our intuitions about the Non-Identity Problem while appealing only to an

independently plausible set of normative principles? To get a better handle on this, it will be useful to lay out the theoretical commitments at issue. The Non-Identity Problem is generated by combining three independently plausible moral principles. The first is the

*Counterfactual Account of Harm*: “someone suffers harm if and only if there occurs an event *e* such that had *e* not occurred, either his life as a whole would have gone better in some respect or he would have fared better in some respect over some interval of time” (Hanser 2008, 439).

On this conception of harm, in order for Alma to have harmed Beatrice there must be some act Alma could have performed such that, had she performed that act, Beatrice’s life would have gone better in some respect. Because Beatrice will not exist if Alma does not bring her into existence now, there is no such action available to Alma. Thus, bringing Beatrice into existence now would not harm her. This conclusion is then combined with two further moral principles. The first is the

*Individual-Affecting Restriction*: an act is only morally wrong if it wrongs someone.<sup>2</sup>

The second is the

*Harm Principle*: an act only wrongs someone if it harms her.

If we accept all three of these principles, it seems that we must reject the intuition that it would be objectionable for Alma to choose to have Beatrice now instead of waiting. Since we cannot claim that Alma has harmed Beatrice (and harms to herself or the potential child she could have had by waiting seem to be

---

<sup>2</sup> Parfit calls this the “person-affecting” intuition. I prefer “individual-affecting,” as it naturally extends the claim to cover all welfare-subjects (not just persons).



beside the point) we cannot claim that choosing to have Beatrice now would be wrong. This result strikes many as implausible; herein lies the problem.

Given this understanding of the Non-Identity Problem, it seems that there are four basic strategies for addressing it: we could reject one of the principles which generate the problem, or embrace the counterintuitive conclusion that it would not be wrong for Alma to bring Beatrice into existence. The first strategy accepts the latter two principles but rejects the Counterfactual Account of Harm. This allows for the claim that Alma's act would be morally impermissible because it would harm Beatrice, even though it would not make her worse off than she otherwise would have been. The second strategy accepts the Counterfactual Account of Harm and the Harm Principle, but rejects Individual-Affecting Restriction. This allows for the claim that Alma's action would be wrong, even though it would not *wrong Beatrice* in particular. The third strategy accepts the Counterfactual Account of Harm and the Individual-Affecting Restriction, but rejects the Harm Principle. This view concedes that Alma's act would not harm anyone and that "wrong acts" must be "wrong for" someone, but denies that harming someone is the only way to wrong them. On this view Alma would wrong Beatrice even though she would not harm her, and her action would be wrong on that account. The final strategy retains all three principles, thus accepting that bringing Beatrice into existence now would be morally permissible because it would not harm anyone (cf. Roberts and Wasserman, xx-xxi).

For taxonomical purposes, it is useful to divide these strategies into two further groups. Some of the ethicists who claim that bringing Beatrice into

existence now would be wrong locate their objection in the effects it would have on the world. They might claim, for example, that having Beatrice now would be wrong because it will make her badly off, or make the world a sadder and more limited place. Claims of this sort appeal to what I will call *outcome-based considerations*. Others, however, object not to the effects of Alma's action, but to the attitude that would be embodied in the choice to have Beatrice now. They might claim, for example, that it would be wrong to bring someone into existence without considering their interests first, and Alma's act would be objectionable on that basis. Claims of this sort appeal to what I will call *intention-based considerations*.

In order to develop my own view of what we should say about the Non-Identity Problem, it will be useful to consider the strengths and weaknesses of the other accounts that have been advanced. In the first chapter I will discuss the most prevalent accounts which appeal to outcome-based considerations. In the second chapter I will discuss the most prevalent views which appeal to intention-based considerations. A fully adequate solution to the Non-Identity Problem would capture our pre-theoretical intuitions by appealing only to independently plausible principles. In my view, none of the existing solutions are able to do this. I do not mean to suggest that the problems I raise for these views are altogether conclusive; perhaps there is some way to suitably revise one of the accounts I will discuss. I leave that possibility open. But, as things stand, it seems to me that solutions which follow one of the first three strategies (rejecting one of the premises which generate the Non-Identity Problem) are incapable of providing a

theoretical justification for why we should say that Alma's action would be wrong, while solutions which follow the fourth strategy (accepting that it would be permissible for Alma to bring Beatrice into existence) are unable to satisfactorily account for the widespread and deep-seeded intuition that Alma would be doing something wrong.

In the third chapter I propose a solution which is aimed at resolving this state of affairs. I draw attention to the fact that, while virtually all of the literature on the Non-Identity Problem is focused on the question of whether or not we should make the negative deontic judgment that it would be *impermissible* for Alma to bring Beatrice into existence now, this is not the only negative moral judgment we might make of her action. It is also possible to make the negative aretaic assessment that, while Alma's action may have been permissible, it was nonetheless blameworthy, vicious, or morally bad. It is possible that the common intuition that choosing to bring Beatrice into existence now would be objectionable is better understood as an aretaic judgment of moral badness than a deontic judgment of impermissibility. I explain why I believe moral badness – rather than blameworthiness or viciousness – is the proper assessment to make of Alma's decision, and why I believe this solution provides a valuable contribution to the existing literature on the Non-Identity Problem.

The view that Alma's action would be morally bad is compatible with nearly every proposed solution to the Non-Identity Problem. However, its significance does in some way depend on which of the existing strategies we find most plausible. When combined with some views, my position will seem

redundant; when combined with others it will seem to add an important insight. Thus, the literature survey I offer in the first two chapters is necessary to illustrate what can and cannot easily be captured by the existing views. This provides a motivation for my view, and allows us to hone in on what exactly needs to be explained.

Before getting into the thesis one more clarification is necessary. The Non-Identity Problem admits of two different versions: a direct version, in which a choice directly determines the identity of the person who will exist after the choice is made, and an indirect version in which a choice initiates a complex chain of events that eventually have an equally decisive effect on the identity of the people who exist after the choice is made (Boonin, Chapter 1, p. 6). Although it is possible that my view may be extended to the indirect version, I will only discuss the direct version. An account of what we should say about the indirect version may need to engage in discussions which I will consider beyond the scope of this thesis.<sup>3</sup> As we will see, dealing with the direct version of the problem is difficult enough.

Finally, it is worth noting that an account of what we should say about the direct version of the problem will have practical consequences that make it more than worthy of our detailed consideration. There are a host of controversial procreative decisions whose permissibility seems to depend to a significant degree on the Non-Identity Problem, such as cloning and genetic screening. Whether there is a moral obligation to undergo preconception testing, whether it is morally

---

<sup>3</sup> For an argument in favor of treating these two versions as distinct problems, see Hanser (1990).

permissible to bring people with disabilities into existence, and whether it is permissible to genetically select for traits that may make the child worse off on the whole all seem to depend on what we say about the direct version of the Non-Identity Problem. My goal in this thesis is to provide an account of moral judgment which will help us explain and cultivate our reactions to these difficult cases.

## Chapter 1: Outcome-Based Considerations

In this chapter I will survey some of the most prevalent outcome-based solutions to the Non-Identity Problem. These views all share the common element of appealing to the state of the world that will result if Beatrice is brought into existence now, and so can meaningfully be grouped together. Beyond this, however, there is not much similarity between them. Each of the solutions considered in this chapter respond differently to the premises which generate the Non-Identity Problem. The first group of solutions suggest that the Counterfactual Account of Harm is mistaken. These views hold that if Alma brought Beatrice into existence now she would harm her, and her action would be wrong on that account. The second type of solution rejects the Individual-Affecting Restriction. These views hold that some acts may be wrong even though they do not wrong anyone, and bringing Beatrice into existence presents just such a case. The third type of solution rejects the Harm Principle. These views hold that, although bringing Beatrice into existence now would not harm her, it would wrong her in some other sense. The final views I will discuss accept all of these premises. These views hold that the Non-Identity Problem is not really a problem; rather than revising our theoretical commitments we should give up the intuition that bringing Beatrice into existence now would be wrong. I will critically assess each of these views in turn.

## Harm-Based Considerations

I have said so far that the Non-Identity Problem rests on the Counterfactual Account of Harm. While this is the most common view of harm in the philosophical literature, it is not the only one. Some philosophers instead prefer the

*Temporal Account of Harm*: “a person is harmed at state  $t_2$ , relative to an earlier time  $t_1$ , if and only if he is in some respect worse off at  $t_2$  than he was at  $t_1$ ” (Hanser 2008, 425).<sup>4</sup>

The nature of harm is a contentious metaphysical issue, discussion of which might easily distract us from the task at hand. However, for our purposes it is not necessary to assess the pros and cons of the Temporal Account of Harm. We need not subject a view to scrutiny unless it affects the conclusion which is drawn from the Non-Identity Problem. Whatever the merits of the Temporal Account of Harm, it seems clear that adopting it would not change the conclusion we are committed to in Non-Identity Cases. Like the Counterfactual Account of Harm, the Temporal Account of Harm is essentially *comparative*.<sup>5</sup> The only difference is that we are now comparing temporal states rather than counterfactual ones. This does not make a difference in Non-Identity Cases. On both views, Alma would not harm Beatrice by bringing her into existence.

On the Temporal Account of Harm, actions like Alma’s harm no one because there is no earlier time,  $t_1$  at which the person (Beatrice) existed. If we assume that non-existence has neither positive nor negative value, we are either barred from making comparisons with it, or must concede that a “life worth

---

<sup>4</sup> For a defense of the Temporal Account of Harm see Velleman, pp. 23.

<sup>5</sup> I will often use the label “Comparative Accounts of Harm” to refer to either of these views, or any others which essentially involve making comparisons between two different states.

living” is “not worse” than nonexistence. This is the same result we get on the Counterfactual Account of Harm. On either account, we cannot say that Beatrice would be harmed by being brought into existence because it is not true that her life would have been better but for Alma’s action. If Alma does not bring her into existence now, she simply will not exist, and non-existence does not seem to be better than a life worth living.

The concept of “a life worth living” is important because it seems possible to have a life which is filled with such excruciating pain that it is worse than non-existence.<sup>6</sup> Consider the

*Wretched Child:* Alma knows that, if she has a child, he will be so multiply diseased that his life will be worse than nothing. He will live only a few years in agonizing pain that can’t be controlled and makes compensating goods impossible (including even ordinary experiences).

Most ethicists agree that it would be wrong for Alma to bring the wretched child into existence. The most natural explanation of why is to appeal to a Comparative Account of Harm. Just as it seems difficult to explain why it would be wrong to bring Beatrice into existence since her life will be worth living, it seems easy to explain why it would be wrong to bring the wretched child into existence since her life will not be. A life is not worth living when the bad things it contains outweigh the good things it contains, or ever will contain. We can truly say of the person who has such a life, that it would have been better *for them* if they had never existed.<sup>7</sup> In order to generate the Non-Identity Problem, then, it is

---

<sup>6</sup> A paradigmatic example of such a case can be found in the description of Dystrophic epidermolysis Bullosa in Jonathan Glover (1992).

<sup>7</sup> David Benetar has argued that it is always worse for someone to be brought into existence because, whereas the absence of goods is not bad for a person, the absence of bads is good for



crucial to suppose that, whatever Beatrice's life will be like, it will not be so bad that it would not be worth continuing. Her life will contain more goods than bads, but fewer goods than the life of the child Alma could have had if she had waited. If bringing Beatrice into existence would harm her, it is not because her existence would be worse than non-existence.

Those who believe that Alma would harm Beatrice, even though her life will not be worse than non-existence, are thus forced to abandon any Comparative Account of Harm. Instead, they argue that we should adopt a

*Non-Comparative Account of Harm:* "a person suffers harm at a time if and only if he is in a certain kind of non-comparatively bad state at that time" (Hanser 2008, 425).

On this view, we do not determine whether or not a person has been harmed by comparing how well-off they are with how well off they would be had an act not occurred. Instead, we simply look at the absolute state of their welfare. Some states, on this view, are non-comparatively bad. To harm someone is simply to cause them to be in one of these states. If it could be shown that bringing Beatrice into existence now would cause her to be in a non-comparatively bad state, we could then say that it would be wrong for Alma to bring her into existence now because she would harm her. This state need not be so bad that it is worse than non-existence. Perhaps some lives, though bad in some respect, are still worth continuing. The non-comparative harm that one's child would suffer in

---

them. Given that all lives will contain some bads, it is always a harm to be brought into existence, and the presence of more goods in their lives cannot compensate for this (Benetar, 30-1). For Benetar, although Beatrice's life will be worth continuing when she exists, this does not mean that being brought into existence does not harm her (*Ibid*, 22-3).

life might provide a strong reason not to bring them into existence even though once they exist their life will be worth living.

Unlike the Temporal Account of Harm, the Non-Comparative Account would alter the conclusion we draw in Non-Identity Cases, and so must be carefully scrutinized. The Non-Comparative Account is often motivated by pointing to cases which the Counterfactual and Temporal Accounts seem to get wrong. Seana Shiffrin, for example, offers the following case:

*Gold Bullion*: a wealthy person (“Wealthy”) flies over an island full of less-wealthy but comfortably off inhabitants and drops a hundred cubes of gold bullion, each worth \$5 million. One person (“Unlucky”) is hit by a falling cube. The impact breaks his arm. Had the cube missed him, it would have landed at someone else’s feet (Shiffrin, 127).

It is clearly wrong for Wealthy to recklessly drop gold cubes and the reason seems to be that she exposes the island’s inhabitants to harm. It seems intuitive to say that Wealthy has harmed Unlucky by causing him to have a broken arm. However, neither the Counterfactual nor the Temporal Account can get this result. It is not true that Unlucky is now worse off than he would have been but for Wealthy’s action. The costs associated with the broken arm will be more than compensated for by the possession of \$5 million, so proponents of the comparative views will have to say that Unlucky has not been harmed. Advocates of the Non-Comparative Account of Harm object that this oversimplifies the situation and obscures what is really going on. What we should say, they claim, is that Unlucky was harmed by having his arm broken, and benefitted by receiving the \$5 million. However, this analysis can have controversial implications. Consider, for example:

*Surgery*: “A doctor cuts a hole in my abdomen in order to remove my swollen appendix. Cutting open my abdomen causes me pain (as I recover); but if the operation had not been performed, I would have suffered worse pain and died very soon” (Harman 2004, 91).

It seems to most people that the doctor does not harm me by cutting a hole in my abdomen in order to remove my appendix. So, this case most naturally seems to favor a Comparative Account of Harm over a Non-Comparative Account. On a Comparative Account, the doctor does not harm me because removing my appendix benefits me on the whole; the doctor does not make me worse off than I otherwise would have been. Elizabeth Harman, however, has disputed this analysis. She says:

While it might appear to be intuitively obvious that the doctor does not harm me, what is really clear is that what the doctor does is permissible... Suppose that instead we say, as I think we should, that the doctor does harm me. He harms me because he causes significant damage to my body. Then there is a reason against performing the surgery in virtue of the harm to me. But consider what reasons there are against not performing the surgery: there is a reason against not performing the surgery in virtue of the fact that I would suffer more severely and die if the surgery is not performed... The reasons against not performing the surgery are stronger than the reasons against performing the surgery; and so performing the surgery is permissible (*Ibid*).

For Harman, Gold Bullion and Surgery are alike in that both the doctor and Wealthy harm someone by causing significant damage to their body. In the Gold Bullion case, the inflicted harm does not seem permissible: the \$5 million benefit that Unlucky receives does not justify the harm he suffers by having his arm broken.<sup>8</sup> In the Surgery case, by contrast, the inflicted harm is permissible: it is permissible to prevent me from suffering greater pain down the road from my

---

<sup>8</sup> That is, unless Unlucky consented to have his arm broken.

swollen appendix by cutting a hole in my abdomen and causing me to experience a smaller amount of pain now. On this analysis there are two features of an action which need to be addressed separately: whether or not it harms someone and whether or not it is permissible. Thus, proponents of the Non-Comparative Account of Harm must engage in two separate tasks. First, they need some way of specifying when a state is harmful and when it is not. Second, they need to be able to explain when it is permissible to cause someone to be in such a state and when it is not. I will address these in reverse order.

Shiffrin argues that the key difference between the Surgery and Gold Bullion cases, which explains why Surgery is permissible by Gold Bullion is not, is that in the Surgery case the harm is delivered to *avoid a greater harm*, while in the Gold Bullion case the harm is delivered in order to bestow a *pure benefit* – a good which is not also a removal from or prevention of harm (Shiffrin, 124). She says, “In the [Surgery] case, the injury is necessarily inflicted to prevent greater harm. Although we sometimes speak as though removing someone from harm *benefits* that person, it does not follow that the *beneficial* aspect of the [surgery] does the moral justificatory work for inflicting the lesser harm. Rather, I believe the fact that a greater *harm* is averted performs the justificatory service” (*Ibid*, 126).

In the Gold Bullion case, by contrast, the island’s inhabitants are not experiencing any pain, and would have been fine if Wealthy had not dropped his bullions. For Shiffrin, whether or not it is okay for Wealthy to drop her bullions depends on whether or not the islanders consent to her dropping them. It seems

that many of us would consent to having our arms broken in exchange for \$5 million, but Shiffrin argues that this is not enough. She says: “[h]arm is objectively bad in such a way that it is morally problematic to inflict (unsolicited) a significant level of it on another for the sake of conferring a benefit, although a person may reasonably decide to undergo the same level of harm to retain the same level of benefit” (*Ibid*, 130). Although it would be reasonable to consent to have one’s arm broken for \$5 million, it is also within one’s rights to refuse. We cannot just assume that benefitting someone on the whole makes inflicting harm on them okay.

In order to see whether or not bringing Beatrice into existence is permissible, then, we need to see whether it is more like the Gold Bullion case or more like the Surgery case. Shiffrin argues that it is more like Gold Bullion. If we grant for the sake of argument that bringing Beatrice into existence now would harm her in a non-comparative sense, we need to see whether this harm would be inflicted in order to bestow a pure benefit or to prevent Beatrice from suffering greater harm. Because non-existence is not a harmful state, it is not plausible to say that bringing Beatrice into existence prevents her from suffering greater harm. Rather, bringing her into existence bestows her with a pure benefit: a life worth living. If pure benefits cannot justify inflicting harms without consent, then bringing her into existence now is not permissible.

Consent is a difficult issue in procreative cases, since there is no way to ask Beatrice whether or not she consents to being brought into existence before she exists. What Shiffrin seems to imply is that since bestowing benefits cannot

justify inflicting harms without consent, and it is impossible for anyone to consent to being brought into existence, it is never permissible to bring someone into existence when doing so will harm them. The difficulty with this proposal is that all cases of childbirth will cause the child to be non-comparatively harmed at some point in their life. It seems arbitrary to claim that the child is only harmed in cases like *The 14 Year Old Girl*. Shiffrin's argument thus seems to leave us with the result that it is never permissible to bring anyone into existence. David Boonin claims that three facts lead to this result:

The first is that all people who exist are non-comparatively harmed in many ways over the course of their existence: they end up enduring conditions that are bad for them to be in. The second is that if a particular person who might exist in fact never exists, then that person will never suffer any non-comparative harm at all. The third is that a person who does exist will never suffer any non-comparative harm before they exist. It follows from these three facts that any couple who considers conceiving a child will cause the child to suffer more non-comparative harm by conceiving it than by not conceiving it: more than it suffered before it existed and more than it would have suffered if it had never existed (Boonin, Chapter 3, 26-7).

Although the conclusion that it is never permissible to bring someone into existence has its defenders,<sup>9</sup> it strikes most people as highly implausible, and is not the result that Shiffrin seems to be after. If she wants to claim that normal cases of childbirth are permissible, it seems that Shiffrin needs to admit that sometimes it is permissible to cause someone to suffer a non-comparative harm, even though that person does not consent and the harm is not inflicted to prevent one from suffering a greater harm.<sup>10</sup> But once she does this, she will no longer be

---

<sup>9</sup> Cf. Benetar (2006).

<sup>10</sup> Perhaps one defense available to Shiffrin would be to claim that, although everyone will non-comparatively harmed at some point in their lives, and their being brought into existence is in

able to explain why it would be wrong for Alma to bring Beatrice into existence. Unless Shiffrin is able to draw a principled distinction between normal cases of childbirth and the 14 Year Old Girl case, she will either have to claim that neither action is impermissible, or that they both are. Appealing to a Non-Comparative Account of Harm cannot explain this difference, because virtually all children will suffer non-comparative harm at some point in their life. So, there must be some other explanation of the relevant difference, and it's not clear what this could be.

Even if we could devise a more adequate account of when it is permissible to cause someone to be non-comparatively harmed, however, the task of specifying when a state constitutes a non-comparative harm and when it does not proves very difficult. Shiffrin suggests that we can identify harms by consulting “a list of evils like broken limbs, disabilities, episodes of pain, significant losses, death” (Shiffrin, 123). These are harms, she claims, because they hinder one’s ability to live an authentic life in which there is not a significant chasm between one’s will and lived experience (*Ibid*, 123-4). Harman develops a similar list by invoking the notion of “a healthy bodily state,” defined as “the normal healthy state of an organism of the species in question” (Harman 2004, 97). She claims that “an action harms someone if it causes the person to be in a state, or to endure an event, that is worse than life with a healthy bodily state.... pain, early death, bodily damage, and deformity... are clear cases of states and events that are worse than life with a healthy bodily state” (*Ibid*, 96-7). In order to account for the 14

---

some sense causally responsible for this, there is another sense in which many of these harms are not attributable to Alma’s act. Perhaps Alma is only culpable for the non-comparative harms that are attributable to her act of having Beatrice (cf. Hanser 1990, 60-1).

Year Old Girl case, which doesn't seem to involve any of these conditions, she adds that "[s]uffering the trauma and emotional problems resulting from inadequate parental support, even if it involves no physical injury, constitutes suffering a harm" (*Ibid*, 94). Thus, we have two different proposals for specifying when a life is not worth beginning: when one's ability to control one's own life would be hindered and when one's life would be worse than that of a normal member of her species.

These standards are worryingly vague, as it's not clear what constitutes a "normal" state for our species or sufficient control over one's own life. Even if we could neatly specify these thresholds, however, we will not be able to solve the Non-Identity Problem without an account of why they are ethically relevant. It is difficult to see what this account might be. If Beatrice is brought into existence, she will be better off than most people in earlier epochs, when the average lifespan was much shorter and the average quality of life was much lower. It's not clear what grounds there could be for objecting to bringing Beatrice into existence but not objecting to bringing people into existence in the distant past, since these people had less control over their own lives than Beatrice will, and lived lives that were far below what we now think of as a normal healthy life. If proponents of Non-Comparative Accounts of Harm want to avoid the conclusion that most historical cases of childbirth were impermissible, they will need to be able to explain why someone in the distant past wouldn't be harmed, and someone in the present would be, even though they are caused to exist with the exact same level of well-being.



This presents advocates of Non-Comparative Accounts of Harm with a dilemma: they can either claim that the conditions that count as non-comparative harms are relative to the society in which one lives, or that there is an absolute standard which transcends historical eras. The second option seems to entail the problematic result that most historical cases of childbirth were wrong. It also seems to arbitrarily privilege our current society, since it is unable to explain why it wouldn't be wrong to cause someone to exist with what is now an average amount of well-being in the future, when people will likely be much better-off than they presently are (cf. Harman 2009, 778). We can imagine a case much like the 14 Year Old Girl arising in such a society, where Alma\* could cause someone to have a higher level of well-being, but instead chooses to cause someone to exist with what is now a typical amount of well-being. According to the second option, which claims that there is an absolute standard which typical human beings in contemporary society meet, there is no problem in this case. But it's not clear why this case is importantly different from the original 14 Year Old Girl case. The first option, which relativizes the judgment of when a person is non-comparatively harmed to the epoch in which the person lives, requires an explanation of why the permissibility of bringing any given person into existence depends on how well off other people are in her society. It seems implausible to claim that given the well-being which is typical in our current society, it is impermissible to bring Beatrice into existence; but if, unbeknownst to Beatrice, everyone in society was much worse off, it would be permissible to bring her into existence.

It thus does not seem that there is any non-arbitrary to claim that Alma would harm Beatrice by bringing her into existence without claiming that most parents harm their children by bringing them into existence or that most historical cases of childbirth were wrong. If these are conclusions that we want to resist, we will need to seek some other solution to the Non-Identity Problem. Rather than rejecting the Counterfactual Account of Harm, we will either need to dispute the Individual-Affecting Restriction, the Harm Principle, or accept the conclusion that it would not be wrong for Alma to bring Beatrice into existence. In the remaining sections of this chapter, I will discuss these options in order.

## Impersonal Considerations

If the arguments made in the last section are correct, Alma would not harm Beatrice by bringing her into existence. Theorists who retain the judgment that it would be wrong for Alma to bring Beatrice into existence thus need some other way of explaining why it would be wrong. The only possible explanation seems to be that one of the remaining principles – the Individual-Affecting Restriction or the Harm Principle – is mistaken. If we can identify a class of actions that are wrong even though they do not make anyone worse off, and explain why these actions are wrong by appealing to independently plausible principles, then we will be able to solve the Non-Identity Problem. In the next section I will consider views which reject the Harm Principle, which holds that an act is morally wrong only if it harms someone. In this section I will discuss views which reject the Individual-Affecting Restriction, which holds that wrong acts must wrong someone. These views accept some form of the

*Impersonal Explanation:* “The correct explanation of the impermissibility of [Alma’s] action is not that it harms: it does not harm. Rather, the action is impermissible because the world is *better* if the action is not performed; it is *impersonally* better, though it is not better for any person” (Harman 2004, 90).

There are two different versions of the Impersonal Explanation which have been proposed as solutions to the Non-Identity Problem: restricted versions and unrestricted versions. Unrestricted Impersonal Explanations “focus on the maximization of *aggregate* wellbeing. Whether we do that (a) by creating additional wellbeing for a particular person, or (b) by bringing a nonidentical but better off person into existence instead, is immaterial. If we have no way of

accomplishing the former, then we *must* do the latter” (Roberts 2009).<sup>11</sup> According to these views, if we can make the world a better place at trivial cost to ourselves, we ought to do so. A world that contains more well-being is a better world, so if we can promote the aggregate amount of well-being in the world at trivial cost to ourselves, we ought to do so.

An Unrestricted Impersonal Explanation of the Non-Identity Problem has been offered by Peter Singer, who suggests that bringing Beatrice into existence would be wrong because it would be a case of “bringing into existence a child with a less satisfactory quality of life than another child whom one could have brought into existence. In other words, [Alma would] have failed to bring about the best possible outcome” (Singer, 124-5). For Singer, ‘the best possible outcome’, needn’t be better *for* any particular individual. If Alma chose to bring Beatrice into existence now when she could easily have a better-off child instead, she would fail her duty to maximize the amount of well-being in the world. Thus, we can explain why choosing to bring Beatrice into existence would be wrong even though it would not harm her.

There are obvious benefits to this sort of view. It does not seem possible for Alma to make Beatrice better off, but it does seem possible for Alma create a better state of affairs, since if she waits to have a child she will bring a better-off child into existence, thus causing more well-being to exist in the world. If there is a moral obligation to create better states of affairs, then perhaps we can explain

---

<sup>11</sup> Derek Parfit calls this view the *Impersonal Total Principle* (Parfit 1984, 387). In the literature on the Non-Identity Problem it is often called *Totalism* (cf. Wasserman and Roberts, pp. xxi). I will refer to it as the Unrestricted Impersonal Explanation in this chapter chiefly to highlight the contrast between this view and what I am calling Restricted Impersonal Explanations.

why bringing Beatrice into existence would be wrong. At first it seems that there is such an obligation: the principle of beneficence can be plausibly interpreted as holding that if we can create a better state of affairs at trivial cost to ourselves, we ought to do so. Note, though, that the Unrestricted Personal Explanation requires us to interpret this obligation in impersonal terms. That is, it requires an obligation to create a better state of affairs *even if* this will not benefit anyone in particular. This proposal turns out to have some extremely counterintuitive implications.

The most obvious challenge to the Unrestricted Impersonal Explanation is that it seems to imply that, if one can increase the aggregate well-being by bringing someone into existence, one is morally required to do so. Since we don't typically regard deciding not to have children as a moral failure, most theorists have found this implication to be implausible.<sup>12</sup> The common slogan is that, while we have an obligation to make individuals happy, we do not have an obligation to make happy individuals. Jan Narveson makes this point when he says:

[T]hat we are to aim at the greatest happiness of the greatest number, does not imply that we are to aim at the greatest happiness and the greatest number. In order to make this perfectly clear, note that the classical utilitarians' view may be put this way: everyone should be as happy as possible. Cast into modern logical form, this reads, "For all persons x, x should be as happy as possible", and this is equivalent to, "if a person exists, he should be as happy as possible." This last shows clearly that the classical formulation does not imply that as many happy people as possible should be brought into existence (Narveson, 62).

---

<sup>12</sup> R.M. Hare seems to be one of the few exceptions. See Hare (1998).

Narveson illustrates the implausibility of understanding act-consequentialist principles in impersonal terms by invoking the following example:

Suppose that we live in a certain country, say, Fervia, and we are told by our king that something is about to happen which will greatly increase the general happiness of the Fervians: namely that a certain city on Mars, populated by extremely happy Martians will shortly become a part of Fervia. Since these new Fervians are very happy, the average happiness, hence the “general happiness” of the Fervians will be greatly increased. Balderdash. If you were a Fervian, would you be impressed by this reasoning? Obviously not... The argument that an increase in the general happiness will result from our having a happy child involves precisely the same fallacy. If you ask, “whose happiness has been increased as a result of his being born?”, the answer is that nobody’s has (*Ibid*, 66).

Narveson’s argument is extremely compelling, but it is in some ways weaker than it could be. This becomes clear when we recognize that the Martians need not be very happy at all in order to increase the aggregate amount of well-being in Fervia. The mere addition of more people will result in there being more aggregate well-being in Fervia. What’s more, if enough people are added to the population – even if their lives are barely worth living – their addition to the aggregate well-being could in principle outweigh whatever costs may be imposed on Fervia’s present population. Derek Parfit has called this the

*Repugnant Conclusion*: for any finite population of individuals at a given quality of life, things would be much better if there was a much larger population of individuals, all of whom have lives just barely worth living (Parfit 1984, 388).

An Unrestricted Impersonal Explanation which appeals only to the amount of well-being in the world seems to entail that a world in which there is a large number of people, each of whom has a very small amount of well-being, would be better than a world in which there is a smaller number of people, each of whom

has a very high level of well-being. Given the implausibility of this result, and the implausibility of a moral requirement to bring happy people into existence, it seems that the principle of beneficence is best understood in person-affecting terms.<sup>13</sup> Thus, it cannot help us solve the Non-Identity Problem.

However, there is an important difference between the sort of case which gives rise to the Repugnant Conclusion and the 14 Year Old Girl case. The Repugnant Conclusion only arises in what Parfit calls Different Number Choices – choices in which an act will affect not only the identities of people who will exist, but the number of people who will exist too (*Ibid*, 361). The 14 Year Old Girl Case, by contrast, is a Same Number Choice – the same number of people will exist whether Alma chooses to bring Beatrice into existence now, or to wait to have a child. Her decision to bring a child into existence now will affect the identity of the person who will come into existence, but not the number of people who will exist.

In order to address Same Number Choices like the 14 Year Old Girl, many theorists have thus attempted to provide a Restricted Impersonal Explanation. These views invoke “principles that require agents to create additional wellbeing ‘for persons’ (even if not for particular persons) but restrict that obligation to the case where they can substitute in a single, better off, nonidentical child in place of the less well off child” (Roberts 2009). The difference between the Restricted and Unrestricted versions of the Impersonal Explanation is that, whereas the Unrestricted view claims that there is always an obligation to maximize aggregate

---

<sup>13</sup> For further arguments to this effect, see Jonathan Bennett (1978), M.A. Roberts (2002), and Shelly Kagan and Peter Vallentyne (1997).

well-being, the Restricted view claims that this obligation only holds in cases where the same number of people will ever exist. This view is most famously captured by Derek Parfit's

*Same Number Quality Claim, or Q*: "If in either of two outcomes the same number of people would ever live, it would be bad if those who live are worse off, or have a lower quality of life, than those who would have lived" (Parfit 1984, 360).<sup>14</sup>

Although this principle is restricted to cases in which the same number of people would ever live, it is not *individual-affecting*: it is not about what would be good or bad for the particular people who will be affected by our action (*Ibid*, 386). It is concerned with good and bad outcomes, and is indifferent to the particular identities of the people who exist in each outcome. It claims that, of two possible worlds, the one which contains more well-being is better, even if it is not better for any particular individual.

Restricted Impersonal Explanations seem well suited to deal with the Repugnant Conclusion. Because they only claim that we have an obligation to maximize aggregate well-being when the same number of people would ever exist, they are not forced to make the implausible assertion that a world which contained more people with lower quality lives would be better than a world which contained fewer people with much higher quality lives. Restricted Impersonal Explanations are also able to avoid the problematic implication that

---

<sup>14</sup> A similar view is Principle N: "Individuals are morally required not to let any child or other dependent person for whose welfare they are responsible experience serious suffering or limited opportunity or serious loss of happiness or good, if they can act so that, without affecting the number of persons who will exist and without imposing substantial burdens on others, no child or other dependent person for whose welfare they are responsible will experience serious suffering or limited opportunity or serious loss of happiness or good" (Buchanan, et al. 2000).



there is a moral obligation to bring happy people into existence. Their claim is rather that, if you are going to bring someone into existence, it is better if they are happier. Of course, this cannot be better *for* the particular child, since if Alma waits to bring a better-off child into existence Beatrice will not exist. Instead, Restricted Impersonal Explanations hold that it would be better in some impersonal sense for Alma to wait to have a better-off child.

Although Restricted Impersonal Explanations seem to fit well with many people's intuitions about the Non-Identity Problem and the Repugnant Conclusion, it is difficult to ground the obligations on which they rely. To begin with, it is not easy to see why the number of people who will exist after we perform an action is morally relevant, so unless there is some explanation of why Same Number Choices are importantly different from Different Number Choices, Restricted Impersonal Explanations will seem arbitrary. Perhaps more pressingly, there needs to be some theoretical support for invoking impersonal obligations at all. What Narveson's argument makes clear is that the moral obligation to make the children you have created better off is importantly different from the supposed moral obligation to make better-off children. Because *Q* suggests that, once you decide to have a child, there is an obligation to make a better-off child, even though that particular child will not be better off, it requires some further theoretical justification. Parfit acknowledges this point when he says: "If [*Q*] is intuitively plausible, and provides objections in our four examples, does it solve the [Non-Identity Problem]? Only superficially. [*Q*] merely restates our intuitions" (Parfit 1982, 123). Many theorists have attempted to find a theoretical

rationale which would support a Restricted Impersonal Explanation of the Non-Identity Problem. It seems to me that all of the existing proposals fail, for a variety of reasons. Although it may be possible to devise a Restricted Impersonal Explanation which overcomes these worries, the following arguments illustrate the considerable difficulties any proponent of this solution will have to overcome.

One possible rationale, proposed by Gregory Kavka, relies on the notion of a 'restricted life'. Kavka defines a restricted life as "a life that is significantly deficient in one or more of the major respects that generally make human lives valuable and worth living" (Kavka, 105). He then appeals to a principle which holds that, other things being equal, "conditions of society or the world are intrinsically undesirable from a moral point of view to the extent that they involve people living restricted lives" (*Ibid*). The idea here is that although it would not be bad *for Beatrice* to have a restricted life, there is a sense in which it would be bad *from the moral point of view*. Restricted lives are impersonally bad, even though they are not bad for the people who live them.

There are a number of problems with this proposal. David Wasserman notes that it is "uncomfortably eugenic, making prospective parents the gatekeepers for a society with fewer unhappy and limited people, a reduction to be achieved by the exclusion of seriously impaired gametes and embryos" (Wasserman, 142). The idea that a society is better, the fewer people with restricted lives it contains, has struck many commentators as an implausible and discriminatory proposal. This becomes even more striking when we consider that Kavka's principle seems to

imply that there is a sense in which it would be better if people with restricted lives never existed at all.<sup>15</sup> Parfit makes this criticism rather forcefully:

Consider someone whose life is well worth living, but who dies at 35. Setting aside effects on other people, is it bad that this person lived? It is of course bad, even tragic, that his life is cut short. But this is not the claim made by [Kavka]. According to [Kavka], though this person's life is well worth living it would have been better if he had never lived. This is not plausible. Consider next parenthood, one of the 'major respects that generally make human lives... worth living.' Is [Kavka's principle] plausible when applied to those who cannot have children, but have lives well worth living? Is it bad that such people ever live? No. Consider next a severe and lifelong handicap. Think of someone born blind whose life is well worth living. Setting aside effects on other people, is it bad that such a person ever lives? Once again, this is not plausible. (Parfit 1982, 120-1).

It does not seem plausible to say that bringing Beatrice into existence would be bad from the moral point of view and, because she would not otherwise exist and her life would be worth living, it doesn't seem that bringing her into existence would be bad from her own perspective either. The only remaining option seems to be that bringing a less well-off child into existence would somehow be bad from the point of view of the person bringing her into existence. Julian Savulescu makes this suggestion in his argument for the

*Principle of Procreative Beneficence*: "couples (or single reproducers) should select the child, of the possible children they could have, who is expected to have the best life, or at least as good a life as the others, based on the relevant, available information" (Savulescu, 415).

Savulescu attempts to support this principle by appealing to rational choice theory. He claims that, if a procreator has reason to believe that one

---

<sup>15</sup> In this respect, Kavka's proposal is unlike Q, N, or the Principle of Procreative Beneficence. Whereas those principles only apply in Same Number Choices, Kavak's principle seems to imply that, in some cases, it would be better if there were fewer numbers.

possible child will be at risk of having a worse life than another possible child one could easily bring into existence instead, it would be *irrational* for the procreator to choose to have the former child. The intuitive idea is that, whereas the child's particular identity should not seem important to the person bringing her into existence,<sup>16</sup> the level of well-being she will experience should seem important. Thus, when deciding which of two embryos one should implant, one should choose the one which is more likely to produce a better-off child. He says: "Embryo B has nothing to be said in its favour over A and something against it. Embryo A should (on pain of irrationality) be implanted" (*Ibid*, 416).

There are two difficulties with this proposal. First, it is difficult to identify the sense in which Embryo B (which we can assume will become Beatrice) has 'something against it'. If it will not be bad *for* Beatrice to be brought into existence, and it will not be bad from the moral point of view, it is difficult to see why Alma should see the fact that Beatrice will likely have a lower quality of life than the child she could have if she waits as an important reason which counts against bringing Beatrice into existence. This connects to the more general problem which seems endemic to all Restricted Impersonal Explanations: if there is a moral reason to cause well-off people to exist, it will be importantly different from the moral reason to cause existing people to be well-off. Savulescu simply assumes that the fact that someone will be worse off counts against bringing them into existence. He doesn't offer any further support for why there is an obligation to cause better-off people to exist instead of worse-off people.

---

<sup>16</sup> For support of this idea, see Hanser's claim that we cannot choose to bring *particular* children into existence (Hanser 1990).

It thus seems to me that both the Restricted and Unrestricted versions of the Impersonal Explanation stand in want of a strong theoretical backing. If we want to claim that the state of affairs which would result from Alma having the better-off child is impersonally better than the state of affairs which would result from Alma bringing Beatrice into existence, then we will need to identify an ethically relevant sense in which this would be better. As we have seen, this is incredibly difficult to do. Those who nonetheless retain the judgment that it would be wrong for Alma to bring Beatrice into existence will need to reject one of the other principles which lead to the Non-Identity Problem. In the next section I will discuss views which reject the Harm Principle. In the final section of this chapter I will discuss views which accept all of the principles which lead to the Non-Identity Problem and the conclusion that they entail.

## Rights-Based Considerations

We have seen that the outcome-based solutions which attempt to solve the Non-Identity Problem by rejecting the Counterfactual Account of Harm or the Individual-Affecting Restriction face serious problems. Accordingly, many have sought to challenge the Harm Principle, which holds that an act is morally wrong only if it harms someone, instead. There are several ways to dispute this principle, each of which gives rise to a distinct solution to the Non-Identity Problem. In the next chapter I will consider solutions which claim that we can wrong someone without harming them by adopting certain intentions or attitudes towards them. These solutions claim that bringing Beatrice into existence would be wrong because of the way in which Alma would *relate* to her. In this section I will consider views which object to the outcome of Alma's decision, rather than the intention with which it was made. On these views, bringing Beatrice into existence would be wrong because of the effects Alma's action would have on Beatrice, even though she would not harm her. The most common solutions of this sort appeal to a class of non-welfarist rights which allow us to make judgments of wrongness independently of welfare assessments of "better" or "worse" off. James Woodward, for example, argues that:

[P]eople have relatively specific interests (e.g. in having promises kept, in avoiding bodily injury, in getting their fair share) that are not simply reducible to some general interest in maintaining a high overall level of well-being and... many moral requirements function so as to protect against the violations of such specific interests. That an action will cause an increase in someone's overall level of well-being is not always an adequate response to the claim that such a specific interest has been violated (Woodward, 809).

Woodward illustrates these rights by appealing to cases of racial discrimination, promise-breaking, and pollution, in which the discriminatory or reckless action has not made the victim worse off – in fact, the victim may have been made better off on the whole – but there still seems to be something objectionable about the action. For him, what is wrong in these cases is that the offending party has violated a right that the victim had in virtue of a specific interest.

Many philosophers have disputed the intelligibility of these sorts of rights. If it can be seen in advance that an act will not harm another in any way, and may even benefit them on the whole, it can seem odd to claim that they have a right against that act being performed. This charge is brought out by Parfit's discussion of a letter in *The Times* in which the child of a fourteen year old girl – who had in fact grown up to have a life worth living – disputed the notion that it would have been better if his mother had waited. Parfit writes: “this man’s letter shows that he was glad to be alive. He denies that his mother acted wrongly because of what she did to him. If we had claimed that her act was wrong, because of a right that cannot be fulfilled, he could have said, ‘I waive this right’. This would have undermined our objection to the mother’s act” (Parfit 1984, 364). Alma cannot ask Beatrice whether she would object to being brought into existence before she is born. If Beatrice does not object to Alma’s decision after coming into existence, it can seem odd to claim that Alma has nonetheless violated her rights.

The most common response to this sort of objection is to note the frequency with which we grant non-welfarist rights in other aspects of life. Doran Smolkin says:

[O]ften we are prepared to recognize that an act is grounds for complaint even though it may not be worse, on balance, for that person. For instance, it is generally recognized as prima facie wrong, and grounds for complaint, for a physician to lie to her patient, and this is the case even if the lie is in the patient's best interest. And, it may be wrong, and grounds for complaint, to take away one's right to vote, even if the person does not know that her right has been taken away, and even though she may never want to exercise her right to vote (Smolkin, 202).

Smolkin's view is that, even if someone did not mind having their right to vote taken away, and would agree to retroactively waive their right, it would still be wrong to take away their right to vote. It is possible that there is something similar going on in the 14 Year Old Girl case. Perhaps Beatrice has a birth-right that it would be wrong for Alma to violate, even though Beatrice would not afterwards object to having had her right violated. For now I will leave aside the discussion of whether or not these sorts of rights are intelligible. I will grant for the sake of argument that there are rights of this sort and turn to the question of whether we can identify any right of Beatrice's which Alma would violate by bringing her into existence.<sup>17</sup>

Woodward argues that the specific interests in virtue of which individuals have rights generate two sorts of obligations for others. First, it is wrong to thwart an interest that another person has, as I do when I actively discriminate against another person. Second, it is wrong, on Woodward's view, to cause another

---

<sup>17</sup> For a helpful discussion of the rights-waiver argument, see Boonin, Chapter 5, pp. 18-40.



person to have a specific interest that I know I will not be able to satisfy, as I do when I make a promise that I know I won't be able to keep (Woodward, 810).<sup>18</sup> It is wrong of me to make such a promise even if making it will result in the other party being better off on the whole (perhaps because I will be able to carry out part, but not all, of what the promise entails). Some theorists attempt to solve the Non-Identity Problem by appealing to the first sort of obligation; others appeal to the second. I will discuss these solutions in order.

Smolkin argues for the first sort of rights-based solution by invoking what he calls the

*Complaint-Warranting Condition thesis*: “a future person has been wronged by some act A, if A was a necessary condition for her to come into being, and A resulted in her having a life that will in some serious sense be defective” (Smolkin, 202).

For Smolkin, bringing someone into existence with a complaint-warranting condition would violate that individual's rights. It would be akin to taking their right to vote away in that, even if they had a life worth living and were glad to have been brought into existence, bringing them into existence would still wrong them. He defines a complaint-warranting condition by appealing to an objective list theory of welfare, according to which things like “physical well-being, understanding, autonomy, liberty, pleasure, and love” (*Ibid*) are what make life worth living. He says that “a future person's life is seriously defective and she can legitimately complain that she has been wronged by an

---

<sup>18</sup> Of course, several authors have argued that it is possible, and perhaps preferable, to give an account of promissory obligation that does not reduce them to general concerns with promoting or respecting the interests of others. Cf. Ross (1930) and McNaughton and Rawling (1998).

earlier person if she was caused by this person to exist with a life that lacked one or more of these goods to the extent that she was unable to lead a flourishing life” (*Ibid*, 203).

This solution bears many similarities to the harm-based solutions we saw in the first section of this chapter, only instead of insisting that Beatrice would be harmed, Smolkin claims that she would be wronged. Smolkin’s solution thus inherits the problems the harm-based solutions faced. It is easy to see why it would be wrong for Alma to bring Beatrice into existence if Beatrice's life would contain *none* of the things that made life worth living. This is because if Beatrice’s life contained none of the things that made life worth living, then her life would be worse than non-existence, and bringing her into existence would harm her. But, by hypothesis, her life will be worth living. To get the result that bringing Beatrice into existence would be wrong even though her life would be worth living Smolkin is forced to appeal to the notion of “a flourishing life,” which contains more goods than the bare minimum which are required to make life worth living, but needn’t contain all of the goods a life could possibly contain. Smolkin claims that this is the standard we must appeal to when making procreative decisions.

It seems likely that all human lives will lack some of the objective goods at some point. In order for Smolkin to avoid the implication that it is always wrong to bring people into existence, he has to say that it is only wrong to bring someone into existence if their life will lack *so many* goods as to prevent them from flourishing. But how many goods must one be deprived of in order to fall

below this threshold? Recall that Beatrice's life will be better off than most people in earlier epochs, whose lives contained less autonomy, liberty, understanding, etc. than Beatrice's will. It seems that, depending on how he specifies the threshold of a flourishing life, Smolkin's solution will either have the implausible implication that most historical cases of childbirth were impermissible, or that the level of well-being that counts as complaint-warranting depends on the average level of well-being in one's society. These worries apply with equal force to Smolkin's view and Non-Comparative Accounts of Harm.

Perhaps, then, we would be better off appealing to the second rights-based strategy, which rejects the claim that bringing Beatrice into existence now would violate a right she presently has, and instead claims that bringing Beatrice into existence now would be wrong because it is foreseeable that Alma would violate her rights later, after she is born. This is the strategy that Woodward advocates.

He says:

Alma knows that if she has a child she will incur certain duties and obligations [that is, parental duties to provide basic forms of care to her child] which she would not otherwise have and which she is very unlikely to meet adequately... the likely failure to fulfill these duties and obligations constitutes an important reason, distinct from... comparative considerations... for Alma *not* to have a child (Woodward, 815).

Unlike Smolkin, Woodward does not claim that there is some standard of "a flourishing life" which is normative for procreative decisions. Instead, his view piggy-backs on a commonsense notion that parents have obligations to do such things as "feed, clothe, and shelter their children adequately... provide their children with love, affection, and certain kinds of education and training, and...

not abuse their children physically or emotionally” (*Ibid*). Whether these obligations are best cashed out in terms of rights the child has is debatable, but the regularity with which we refer to parental obligations such as these seems to provide a strong foundation for Woodward’s view, which Smolkin’s view lacks.

The primary difference between Woodward’s account and the accounts advanced by Smolkin and proponents of Non-Comparative Accounts of Harm is that, instead of claiming that Alma’s action would be wrong because it would cause Beatrice to exist in an undesirable state, Woodward claims that it would be wrong because it would cause Alma to incur obligations that she knows she will not be able to adequately meet. There are two problems with this view. First, it is not clear that Alma would be any less likely to meet the parental obligations that Woodward lists than most parents are. While it is likely that Alma’s immaturity will lead her to wrong her child in minor ways, it is not obvious that every 14 year old girl who has a child will fail the more general obligations to “feed, clothe, and shelter their children adequately... provide their children with love, affection, and certain kinds of education and training, and... not abuse their children physically or emotionally” (*Ibid*).

In order to make good on the claim that Alma will knowingly incur obligations that she will not be able to meet, it thus seems that we need to specify in more detail exactly what Alma’s obligations to Beatrice will be. It seems plausible to claim that, as Beatrice’s parent, Alma will have an obligation to provide the best life that is available to Beatrice, or ensure that her life is worth living. Alma would be able to meet either of these obligations. So, in order for

Woodward's strategy to solve the Non-Identity Problem, we will need to claim that parental obligations require something more than this. It's hard to see how we can do this without running into the problems that faced Smolkin and proponents of Non-Comparative Accounts of Harm. For example, Woodward might claim that parents have an obligation to make sure that their children's lives meet a certain standard, say, the standard of a flourishing life. If this is the proposal, we could simply point out that, because of the society she lives in, Alma will be able to provide more care for her child than many others, such as slaves or people in the distant past. Woodward's account would thus imply that each of these people have violated their child's rights by bringing them into existence even though they knew that, because of their circumstances, they would not be able to provide their child with certain forms of basic care. If we try to avoid this worry by claiming that parents only have an obligation to ensure that their children are as well-off as a typical human being in their historical era, we will need to explain why parental obligations depend on how well-off others are in one's society. As we have seen, this is a formidable task.

It seems plausible to claim that parental obligations do not require parents to ensure that their children meet any specific threshold of well-being, but rather require parent's to adopt certain attitudes and intentions towards their children. I will discuss this proposal in the next chapter, as it would provide an intention-based reason, rather than an outcome-based reason, to object to Alma's decision. For now the question is whether or not it is possible to identify an outcome-based

right that Alma would violate. It does not seem that we can do so without running into the same problems that confronted Non-Comparative Accounts of Harm.

These problems seem endemic to all outcome-based solutions which retain the Individual-Affecting Restriction. There does not seem to be any non-arbitrary way to claim that bringing Beatrice into existence would wrong her without entailing that many other procreative decisions, which are intuitively unobjectionable, wronged the individuals they created as well. It thus seems that we must reject the Individual-Affecting Restriction (which we have seen has problems), accept the conclusion that it would not be wrong for Alma to bring Beatrice into existence, or search for some non-outcome-based way to claim that Alma has wronged Beatrice. In the next section I will discuss views which accept the conclusion that it would not be wrong for Alma to bring Beatrice into existence. In the following chapter I will consider the most prominent intention-based solutions to the Non-Identity Problem.

## Accepting the Implausible Conclusion

The Non-Identity Problem results from the fact that when we combine three plausible principles which most people are inclined to accept, we are left with an implausible conclusion which most people are inclined to reject. In the previous sections of this chapter I argued that it is difficult to identify any plausible revisions to these three principles which would explain why Alma's decision in the 14 Year Old Girl case would be wrong. If these arguments are correct, it is natural to revisit the conclusion that it wouldn't be wrong for Alma to bring Beatrice into existence, to see just how implausible this claim really is. According to the views I will discuss in this section the answer is: it really isn't that implausible after all. Rather than revising our theoretical commitments, we should accept the conclusion that they entail.

Several philosophers have accepted that, since there doesn't seem to be any theoretically supportable way to avoid the implausible conclusion, we must accept it.<sup>19</sup> Insofar as this solution simply instructs us to give up the intuition that Alma would be doing something wrong, though, it can seem unsatisfactory, as it does little to render the conclusion more plausible than it originally seemed. For a solution of this sort to be successful, one needs to do more than assert that we ought to give up our original intuition, or report that one doesn't share the intuition that Alma would be doing something wrong. One needs to provide

---

<sup>19</sup> Cf. Schwartz (1978; 1979), Heyd (1992), Roberts (1998), and Boonin (forthcoming).

grounds for thinking that the implausible conclusion isn't so implausible, and explain why those who do share this intuition ought to give it up.

David Boonin has tried to meet these tasks by engaging in a two-step process. First, he argues, in much the same way that I have so far, that there is no satisfactory account of how to revise the principles that lead to the Non-Identity Problem in such a way that they would avoid the implausible conclusion. Boonin claims that engaging in this process has two effects: first, it provides an argument for the truth of the implausible conclusion; second, it renders the implausible conclusion more plausible. He says: “[t]he process of coming to see the problems that arise when we try to avoid the Implausible Conclusion should lead us to see that the conclusion itself is not as implausible as it at first seemed” (Boonin, Chapter 7, 4). Although nearly everyone who first encounters the Non-Identity Problem has a clear and visceral negative reaction to Alma's decision, Boonin claims that once we have engaged in the process of philosophically scrutinizing the principles which lead to this conclusion, we will be less sure that this reaction is philosophically supportable, and more willing to give it up.

While it may be true that realizing that our pre-theoretical intuitions are difficult to support should make us more willing to revise these intuitions, it remains the case that nearly every philosopher who has discussed the Non-Identity Problem in the past thirty years has regarded the intuition that it would be wrong for Alma to bring Beatrice into existence as an intuition that any successful moral theory must be able to account for. So, in order to render the conclusion that Alma would not be doing anything wrong more plausible, we need to do



more than philosophically scrutinize the principles that lead to the Non-Identity Problem. The second part of Boonin's strategy is to compare the 14 Year Old Girl case with other cases that have certain structural similarities, in order to show that Alma's action is much like other actions about which most people don't have such negative reactions. Highlighting these similarities can undermine the certainty with which people claim that Alma would be doing something wrong, and put pressure on them to explain why we ought to feel worse about Alma's action than others which are similar in important ways. Three of Boonin's cases strike me as particularly well-suited to this task. First, consider

*Betty*: Betty has a slight preference to have a child. She goes to her doctor for a pre-conception checkup and the doctor tells her that if she has a child, no matter when the child is conceived, the child will be poorly off. Knowing this, Betty decides to conceive and, as a result, she has a child who is as poorly off as [Beatrice] (*Ibid*, 14).

Most people would not react negatively to Betty's decision. But her decision, like Alma's, results in a poorly off child existing. The only difference is that, whereas it is possible for Alma to cause a better-off child to exist, it is not possible for Betty to cause a better-off child to exist. In order to explain why we should have such different reactions to these two cases, we need an account of why the ability to cause a better-off child to exist is morally important. Most theorists who address the Non-Identity Problem think that this difference is morally important, but at the very least Betty's case should cast doubt on the general claim that it is wrong to bring poorly off people into existence, and this should make us more receptive to the possibility that it might not be wrong for Alma to bring Beatrice into existence. Next, consider

*Fred*: Fred has two children, Billy and Timmy. Billy is less well-off than Timmy. Fred finds himself in a situation in which one of them will drown, and he cannot save them both. He decides to save Billy rather than Timmy because it is a bit more convenient for him to do so (*Ibid*, 9).

Again, most people would not react negatively to Fred's decision. But Fred's decision is like Alma's in many ways. Both Fred and Alma choose between two options which will determine who will exist in the future, and they both choose the option which results in a less well-off person existing instead of a better-off person existing, because it is a bit more convenient for them to do so. Unlike Betty, Fred *could* cause a better-off person to exist, but this doesn't make his action seem objectionable.

There are three ways we could respond to this situation: we could claim that both Fred and Alma would be acting objectionably and that we should revise the intuition that Fred would not be doing anything wrong (this is the option we will be drawn to if we understand the principle of beneficence in impersonal terms); we could claim that both of their actions are permissible and that we should revise the intuition that Alma would not be doing anything wrong (this is the option that Boonin advocates); or we could try to find a principled difference between the cases which explains why Alma's action would be impermissible but Fred's action would not be. Invoking the case of Fred does not provide a direct argument in favor of the view that we should accept the implausible conclusion, since it cannot tell us which of these options we should choose. But this is not the point of the case. The point of Fred is simply to illustrate that the second option, which requires us to revise the intuition that Alma would be acting wrongly, is not obviously any more implausible than the other options. If we can make the

implausible conclusion seem less implausible by comparing it with other cases that don't seem problematic, and combine this with the fact that it is difficult to see how we could revise theoretical commitments in a way that would avoid the implausible conclusion, then we will have a compelling argument to accept the conclusion. Finally, consider

*Mary*: Mary lives on a farm with her two children and four horses. After deciding that she wanted to conceive one more child or breed one more horse but not both, she consulted with a doctor and a veterinarian. The doctor assured her that if she conceived a child, the child would be perfectly healthy and the veterinarian assured her that if she bred another horse, the horse would be perfectly healthy. After spending a long time considering whether she would rather have one more child or breed one more horse, she decided that she had a slight preference to breed one more horse. As a result, she bred one more horse and did not have any more children (*Ibid*, 25).

Most people would not object to Mary's decision. But it is like Alma and Fred's decisions in that Mary chooses between two options which will determine who will exist in the future, and chooses the option which result in a less well-off creature existing instead of a better-off creature.<sup>20</sup> Unlike Fred's case, Mary chooses which creature to bring into existence rather than which creature to save; her case is like Alma's in this respect. The only difference between her case and Alma's is that, whereas the less well-off creature that Alma chooses to bring into existence is a human being, the less well-off creature that Mary chooses to bring into existence is a horse. In order to explain why we should react differently to these cases, we need an account of why species membership is morally important.

---

<sup>20</sup> I take it as fairly uncontroversial that the child will be better-off than the horse, if only because she will have a much longer life expectancy.

Again, simply invoking the case of Mary does not provide a direct argument in favor of accepting the implausible conclusion, but it does seem to put pressure on those who want to retain the intuition that Alma would be acting objectionably, and make the claim that Alma would not be acting objectionably seem less implausible. When combined with the fact that it is difficult to devise a theoretically supportable way to revise our theoretical commitments in a way which would avoid the implausible conclusion, this provides a strong argument for the claim that we should accept the implausible conclusion.

Boonin's strategy makes an important contribution to the literature on the Non-Identity Problem because it has too often simply been assumed that Alma's action would be wrong and that any satisfactory moral theory must be able to account for this. It is important to be open to the possibility that our pre-theoretical intuitions are mistaken, and to revisit the intuition that Alma would be doing something objectionable after scrutinizing the relevant principles, in order to see whether or not this is actually the appropriate reaction. I agree with Boonin that it is much less obvious that Alma would be doing something wrong than is commonly supposed, and it seems to me that his argument at least succeeds at taking some of the sting out of accepting the implausible conclusion.

However, it is worth noting that, when presented with these cases, most philosophers will either argue that there is a principled difference between Mary, Fred, Betty and Alma, or will claim that we should revise our intuitions about one of these other cases and retain the judgment that Alma would be acting

objectionably.<sup>21</sup> This seems to indicate that most people are more confident that Alma's action would be objectionable than they are that Mary, Fred, and Betty's actions would not be. Before we give up this intuition entirely, it is worth investigating Alma's decision in more detail to see if we can identify any morally relevant features which would account for the widespread and recalcitrant negative reaction that most people have. In the final chapter of this thesis, I will argue that much this intuition may have to do with Alma's intentions and beliefs, and that the appropriate reaction to Alma's decision may be somewhat different than is commonly supposed. This solution is largely compatible with Boonin's view, and the two support each other in important ways.

Those who think that we should retain the intuition that Alma would be acting objectionably and that we should reject the implausible conclusion need some way of theoretically supporting this claim. I have argued so far that it is difficult to do this by appealing to outcome-based considerations, as all of the major existing attempts to do so have proven unsatisfactory. Unless one of these solutions can be adequately revised, it seems to me that the only remaining option is to attempt to explain why Alma's action would be wrong by appealing to intention-based considerations. I will discuss solutions which attempt to do this in the next chapter. After pointing out a number of problems with the existing intention-based solutions, I will provide my own account in the final chapter.

---

<sup>21</sup> There is some evidence that Woodward and some proponents of Non-Comparative accounts of Harm would object to Betty's decision (see especially Woodward, 816) and that advocates of the Impersonal Explanation would object to Fred's decision. Others have insisted that the fact that Betty could not act otherwise makes her decision permissible, and that there is an important difference between causing people to exist and saving people's lives (see McMahan 2012).

## Chapter 2: Intention-Based Considerations

In the last chapter I surveyed some of the most prevalent outcome-based solutions to the Non-Identity Problem and raised a number of problems those solutions face. If the arguments considered in that chapter are correct, it seems difficult to claim that there would be any outcome-based reasons to object to Alma's action.<sup>22</sup> Those who nonetheless retain the intuition that it would be objectionable for Alma to bring Beatrice into existence now thus need to find some other way to support this intuition. In this chapter I will survey the most prevalent views which attempt to do this by appealing to intention-based considerations. These views reject the Harm Principle, according to which an act wrongs someone only if it harms her, and claim that Alma's action would wrong Beatrice because of the intentions and attitudes she would display toward her. In the first section I will discuss views which hold that Alma's action would be wrong because it would violate a principle of mutual respect. In the second section I will discuss views which hold that Alma's action would fail a parental obligation to show concern for her child. In the third section I will consider the nature of intention-based considerations, and whether the intentions and attitudes with which Alma would have acted could render her action impermissible.

---

<sup>22</sup> That is, assuming we leave out any effects Alma's action may have on third parties.

## Respect-Based Considerations

All of the intention-based solutions in the literature on the Non-Identity Problem deny the Harm Principle, according to which an act is only morally wrong if it harms someone. These solutions claim that there is an important distinction between harming and wronging and that, while Alma's action would not harm Beatrice, it would wrong her.<sup>23</sup> Rahul Kumar gives voice to this idea when he says:

The kinds of considerations that are relevant for determining whether or not a person has been harmed have primarily to do with the state of the person who claims to have been harmed. Whether or not another has wronged one, on the other hand, has to do with facts concerning the character of the wrongdoer's regulation of her conduct with respect to how she related to the wronged (Kumar, 100).

The character of an agent's regulation of her conduct does not simply have to do with the effects her action has on another. It is a broader category which includes the reasons for which she acted and the way in which she saw others as factoring into her decision-making process.<sup>24</sup> It is possible for an agent to have wronged another, in this sense, without having harmed them, and it is possible that this is what explains the negative reaction to Alma's decision. A successful explanation of this sort must provide a plausible account of what would be wrong with 'the character of Alma's regulation of her conduct with respect to how she

---

<sup>23</sup> It is of course possible to have a Non-Comparative view of harm while claiming that Alma's action would be wrong for intention-based reasons. For simplicity, though, and because it seems to be the dominant position in the literature, I will speak as if those who propose intention-based solutions to the Non-Identity Problem accept a Comparative Account of Harm.

<sup>24</sup> Kumar claims that one person harms another when they fail to comply with that person's legitimate expectations, where failing to comply with someone's legitimate expectations "is not just a matter of having failed to conduct oneself in a certain way. It can also be understood as a failure to have been responsive to certain considerations that it was legitimate to expect one to have been responsive, or to have taken into account considerations that it was reasonable to have expected one to have disregarded as irrelevant for one's deliberations at that time" (Kumar, 107).

related to Beatrice', given that Alma will not harm Beatrice by bringing her into existence. According to the view I will discuss in this section, it would be wrong for Alma to bring Beatrice into existence because something about the character of her regulation of her conduct would display a lack of respect toward Beatrice.

Kumar argues for this sort of view by appealing to a contractualist framework on which moral principles "fix what the general terms are for relating to one another on a basis of mutual respect for one another's value as persons, by establishing certain legitimate expectations concerning consideration and conduct between persons" (*Ibid*, 106). These principles hold between the parties to certain types of relationships, and are meant to specify the requirements of respecting one another as rational self-governors. The capacity for rational self-governance opens one up to being wronged without being harmed because it exposes one to "a distinct kind of vulnerability, a vulnerability to what another's reasons, or reasoning, concerning how it is appropriate to relate to oneself, says about oneself" (*Ibid*, 107). Because rational self-governors are vulnerable to the way in which we reason about them and factor them into our decisions, we ought to relate to them in certain ways and not others.

This framework allows Kumar to propose an alternative to the Harm Principle. He claims that acts can wrong others either by harming them *or* by violating contractualist principles. This happens when "the wrongdoer has, without adequate excuse or justification, violated certain legitimate expectations with which the wronged party was entitled, in virtue of her value as a person, to have expected her to comply" (*Ibid*). What individuals can legitimately expect of



one another depends, in part, on the kind of relationship which holds between them. These relationships are to be understood as holding between certain types of people – employer and employee, teacher and student, mother and daughter, etc. For example, if a teacher is holding office hours at the end of the day, but is tired and wants to go home a little bit early, she does not need to know whether any particular student will be adversely affected in order to decide whether or not it would be wrong to go home early. What she needs to do, instead, is reflect on the legitimate expectations which govern the relationship of student and teacher. If someone fitting under the type description ‘student’ could legitimately expect her to stay in her office, and she leaves anyway, she will have wronged all tokens of that type, even if none of them come to her office hours or are adversely affected in any way.

Appealing to contractualist principles which hold between people in interpersonal relationships gives Kumar’s strategy an advantage over outcome-based solutions because, whereas harm- and rights-based solutions need to show that Beatrice would be adversely affected in some way, Kumar simply needs to show that Alma and Beatrice fit under type descriptions which generate legitimate expectations that Alma would violate by bringing Beatrice into existence now. Whether or not Beatrice is actually made worse off is beside the point.

It can be difficult at first to see how this could provide a solution to the Non-Identity Problem, because at the time of Alma’s action Beatrice does not exist, and so is not capable of having any expectations of Alma’s conduct. But Kumar claims that a contractualist solution can overcome this worry because

“knowledge of the particular identity of the other is not essential. What is essential is that one has reason to take the other to be of a certain type” (*Ibid*, 112). People can be wronged, as tokens of types, by acts that pre-date their existence. For instance, if I plant a bomb now that explodes in 200 years, I will have wronged those harmed by the explosion. The victims of the explosion are tokens of the type ‘person’, and there is a principle governing the person/person relationship which tells us not to harm others. Before I plant the bomb I can reference this principle in order to know that doing so would be wrong. If the bomb fails to explode and does not harm anyone, I will still have done something wrong because of the way in which I related to the potential victims of the explosion. Just as the teacher does not need to adversely affect any students in order to wrong them, I do not need to harm any future people to have wronged them, and they do not need to exist at the time of my action in order for me to have violated their legitimate expectations.

Appealing to the requirements of the person/person relationship is not enough to solve the Non-Identity Problem, however, because unlike the bomber Alma does not expose anyone to risk of being harmed. Alma knows that, whoever her child turns out to be, her life will not be so bad that it is not worth living. If the arguments made in Chapter 1 are correct, this means that whatever child she chooses to have, that child will not be harmed. In order to provide a contractualist solution to the Non-Identity Problem, then, Kumar needs to show that there is another relationship which holds between Alma and Beatrice, in virtue of which

Beatrice could legitimately expect Alma to refrain from bringing her into existence.

Kumar says that “[w]hat an individual may legitimately expect of others, and what may be legitimately demanded of her, at a given time turns on both (a) what expectations can in fact be defended on the basis of the relevant principle, and (b) the relevant type descriptions that happen to fit her and her circumstances at that time” (*Ibid*, 111). He attempts to solve the Non-Identity Problem by claiming that the relationship between Alma and Beatrice fits under the type description ‘caretaker and dependent’, which entails the following principle:

M: Those individuals responsible for a child’s, or other dependent person’s, welfare are morally required not to let her suffer a serious harm or disability or a serious loss of happiness or good, that they could have prevented without imposing substantial burdens or costs or loss of benefits on themselves or others (*Ibid*, 112).

For Kumar, Alma would violate the legitimate expectations which govern the caretaker/dependent relationship by bringing Beatrice into existence now, and her action would be wrong on that account. Just as a teacher does not need to know whether or not any particular student will be made worse off in order to determine what can legitimately be expected of her, Alma does not need to know whether or not she would harm Beatrice in order to determine whether or not having a child now would be wrong. All that is essential is that Alma has reason believe that she and her child will fit under the type descriptions of caretaker and dependent, and that this relationship entitles Beatrice to expect Alma not to relate to her in the way that she would by bringing her into existence. Alma knows that if she chooses to bring someone into existence now, that child will fit the type description “Alma’s child.” “Alma’s child” is entitled to expect her caretaker to

reason and act in ways which protect her from suffering a serious harm or disability. By choosing to have a less well-off child now, when she easily could have waited to have a better-off child, Kumar believes that Alma has violated “her child’s” legitimate expectations.

It seems to me that there are three main problems with this view. The first two problems have to do with Kumar’s framework which grounds claims to have been wronged in the legitimate expectations which hold between rational self-governors; the third problem has to do with the claim that we owe obligations to people that meet certain type descriptions. Appealing to a framework which grounds claims to have been wronged in the legitimate expectations which hold between rational self-governors is problematic because it seems to imply that only rational self-governors can be wronged, but Beatrice is not yet a rational self-governor, and will not be for some time. Kumar’s ability to make sense of duties to potential rational self-governors is limited. He says:

[T]here will be important constraints on how one relates to one’s child even at the early stage of the developmental process where the child has only the potential to develop a capacity for rational self-government. One may, therefore, be entitled to *terminate* the caretaker-dependent relationship with one’s child, but as long as one allows the relationship to continue, it may be that one is *not* entitled to *interfere* in certain ways with the child’s development (*Ibid*, 115).

For Kumar, Alma’s obligations are owed to a future version of Beatrice as a rational self-governor and, until she has reached this status, Alma can terminate her relationship with Beatrice at any time. This implies that if Alma does choose to become pregnant with Beatrice, she can absolve herself of any obligations she will owe Beatrice by simply having an abortion. This can’t count as wronging

Beatrice until she has become a rational self-governor. What's more, Beatrice will not develop into a rational self-governor for some time after she is born. Kumar's account thus seems to license infanticide until she reaches this point. Abortion and infanticide do not wrong anybody because rather than violating the principles which govern relationships between existing rational self-governors, they terminate relationships with merely potential rational self-governors. But it seems odd to claim that choosing to have a child with a life worth living would be disrespectful, while aborting them or killing them after they are born would not be.

The claim that Alma can terminate her relationship with Beatrice without wronging her also seems to imply that whether or not Alma has wronged Beatrice depends on whether or not she expects a caretaker/dependent relationship to hold between them. Because the contractualist principles that Kumar invokes hold between types of people in interpersonal relationships, it seems that Alma only has an obligation to refrain from bringing Beatrice into existence so long as she expects a relationship to hold between them which fits under the type description of caretaker and dependent. However, there may be situations in which the procreator knows that such a relationship will never take place. For instance, imagine that Alma\* knows that, if she conceives now, she will die during childbirth. She also knows that, if she decides to have a child anyway, her child will grow up with the same level of parental care as Beatrice, and will be equally as poorly off. Kumar's account seems to imply that, whereas choosing to bring Beatrice into existence in the original case wrongs her, choosing to bring

Beatrice\* into existence does not wrong her, because Alma\* knew that a relationship which falls under the caretaker/dependent type description would never hold between them, and so was not subject to those principles. It seems odd to claim that Alma has acted disrespectfully toward her child but Alma\* has not, since Alma and Alma\* both knowingly bring children into existence with the exact same level of welfare.

Perhaps there is some way to revise this framework so that it does not have these implications. I leave that possibility open. However, even if we do this, it seems that there will still be problems with Kumar's claim that there are ethical obligations which hold between 'types' of people. Parfit addresses a view like Kumar's in *Reasons and Persons*, which holds that although Alma does not harm anyone in particular, there is still a sense in which she harms the type "her child," since "her child" is worse off than she could have been. He says that claiming that "Alma's child" was made worse off "does *not* explain the objection to this girl's decision. This becomes clear after she has had her child. The phrase 'her child' now naturally refers to this particular child. And this girl's decision was *not* worse for *this* child" (Parfit 1984, 359). If Beatrice in particular will not be harmed it is difficult to see why the fact that "Alma's child" will be made worse off is of moral significance.

Boonin brings out this challenge by noting that, if Alma chose to adopt a relatively unhealthy child instead of a very healthy child, there is a sense in which she would make "her child" worse off than she could have been (Boonin, Chapter 2, 12). But choosing to adopt the less healthy child is not problematic. It seems

like this is best explained by noting that, while it is true that adopting the less healthy child will make “Alma’s child” worse off than she could have been, it does not make any particular child worse off. It is the effects that our actions have on particular individuals that are morally significant. Although it may be true that choosing to bring Beatrice into existence now would make “Alma’s child” worse off than she could have been, it would not make any particular token worse off, and this seems to be what is morally important.

It thus seems to me that Kumar’s solution cannot solve the Non-Identity Problem. If bringing Beatrice into existence would be wrong, it is because it would wrong Beatrice now, as she is brought into existence, and not as a future rational self-governor or a token of the type “Alma’s child”. Still, the claim that Alma owes special obligations to Beatrice in virtue of the role she inhabits, and that these obligations require more than just providing for certain outcomes, has struck many highly plausible. In the next section, I will discuss a solution which attempts to justify the claim that bringing Beatrice into existence now would be unfitting of a caretaker by appealing to the role morality of parents rather than a principle of mutual respect which holds between Alma and “her child.” After laying out this solution and discussing its pros and cons, I will raise a problem which I believe is devastating for all intention-based solutions to the Non-Identity Problem.

## **Role-Based Considerations**

To this point we have seen two solutions which attempt to appeal to the special obligations of parents or caretakers in order to solve the Non-Identity Problem. In Chapter 1 I discussed Woodward's claim that, if Alma chose to bring Beatrice into existence now, she would knowingly incur obligations that she would not be able to fulfill, and her action would be wrong on that account. This view is unpromising insofar as it requires parents to provide a standard of care which is measured in terms of outcomes, as such a solution will either entail that most historical cases of childbirth are wrong or that the level of care parents are required to provide is relative to the society in which they live. In the last section we saw a view that appealed to the special obligations of caretakers but attempted to specify these obligations by appealing to principles of mutual respect between rational self-governors. This framework also seemed unpromising. Perhaps, though, there is some way to capture the intuition that Alma would be failing in her role as a caretaker or parent without running into either of these problems.

David Wasserman has suggested that we can accomplish this task by reflecting on the role morality of prospective parents. His solution does not rely on the obligations Alma will later have as a parent or caretaker, but rather the obligations she has now, as someone deciding whether or not to bring someone into existence. He argues that people inhabiting this position have an obligation to adopt certain intentions towards their potential children. He says:

[T]he most plausible role morality is a permissive one, placing few constraints on the kinds of children parents have, as long as they are expected to have lives worth living... What it does require is that if people bring children into the world, they do so in part for



certain reasons, reasons that concern the good of those children... parents' reasons for having a child with certain characteristics will play a crucial role in determining the acceptability of their decision and the plausibility of any complaint from the child they create (Wasserman, 135-6).

Wasserman's solution is somewhat different from the others we have considered because he does not believe it is always impermissible to choose to bring a less well-off child into existence. However, this does not lead Wasserman to accept the implausible conclusion wholesale. He says: "the decision to have an impaired child when one could have an unimpaired one need not reflect insensitivity to the expected hardship of the former" (*Ibid*, 135). But, when the decision to have an impaired child does reflect insensitivity on the part of the prospective parent, Wasserman claims that the decision is impermissible, because it constitutes a failure in one's role as a prospective parent. Prospective parents have an obligation to adopt certain intentions and attitudes toward the child they bring into existence.

In order to make good on this solution, Wasserman does not need to show that Alma's decision would adversely affect Beatrice, nor does he need to show that Beatrice in particular would be wronged as a token of a type. What he must show, instead, is that Alma has acted in a way that is unfitting of a progenitor or prospective parent. If Alma has acted in a way that is unfitting of her role others will be able to criticize her decision, and Beatrice will be able to complain against her, on that basis. In order to make this judgment we need an account of how prospective parents ought to reason and why it would be wrong for them to have a child if they do not reason in this way – then we will need to show that Alma has failed these requirements. To this end, Wasserman says:

All prospective parents should expect their children to face significant hardships – death, loss, frustration, and pain – that dwarf the specific hardships associated with most impairments. They must be able to justify the decision to subject their children to those hardships, and they can do so only if part of their reason for having those children is to give them lives good and rich enough to offset or outweigh those hardships (*Ibid*).

In the first chapter we saw that children who are brought into existence with lives worth living are not harmed, as the goods in their lives outweigh the bads. However, Wasserman claims that the fact that the child has not been harmed is not enough to determine whether or not the procreator has acted permissibly by bringing them into existence. This is because parents need to be able to justify the decision to bring the child into an existence that they knew would contain some bads. They can only do this if they have children, in part, for certain reasons. He says: “if the expected good of the child’s life plays no role in [the procreator’s] decision to have the child, that good cannot be weighed against, or used to justify, the suffering or limitation he is expected to face” (*Ibid*, 147). Wasserman thus shifts the focus from the question of whether or not the child who is brought into existence has been harmed or wronged to the question of whether or not the procreator has satisfactorily carried out their role. This requires that they are able to justify the decision to bring a child into existence – before bringing them into existence – and they can only do this if the goods the child will experience in life played a role in their decision. If they cannot justify their decision, Wasserman claims, the procreator will have done something wrong, and the child will have reason to complain against her.

The claim that parents ought to think about how well off their children will be before bringing them into existence seems intuitively attractive and, given

the problems we have seen for the claim that we ought not to bring certain people into existence, a permissive morality for procreative decisions seems more plausible than it might at the outset. What's more, it seems that Wasserman can explain the common intuition that it would be wrong for Alma to bring Beatrice into existence now since as the case is initially described it seems pretty clear that Alma would not think about Beatrice's good before bringing her into existence. When her friend comes to her and says, "[t]his is not entirely your affair. You should think not only of yourself, but also of your child. It will be worse for him if you have him now. If you have him later, you will give him a better start in life" (Parfit 1984, 358-9), Alma does not respond by explaining that Beatrice will still have a life worth living and will not be harmed. In fact, she doesn't seem to provide any justification at all. Perhaps this is what is at the core of the common intuition that it would be wrong for her to bring Beatrice into existence. The intuition may have less to do with the character of Beatrice's existence than is commonly supposed, and more to do with the way in which Alma made her decision. I am sympathetic to this idea, and in the next chapter I will propose a solution which attempts to explain and justify this claim. However, I believe that there are a couple of problems with the way in which Wasserman attempts to capture this important insight.

An initial worry concerns Wasserman's attempt to cash out the obligations of prospective parents in terms of the birth-rights of prospective children. He says that we should "understand the child as having birth-rights whose satisfaction or violation depend on the attitudes with which his parents create him" (*Ibid*, 151).

In the first chapter we saw that appeals to birth-rights have difficulty explaining why prospective children have a right to a life which meets some threshold above a life worth living. Since a better life is not available to them, and their life will be worth living, it seems odd to claim that their rights were violated by being brought into existence. Whatever hardships they might experience in life seem to be more than compensated by their worthwhile life. It seems like we could run a similar argument against Wasserman's position: if the insensitive attitudes prospective parents display by choosing to have certain children are a necessary condition of their being brought into existence, what grounds can the child have to objecting to the parent's decision? They are not worse off because of their parent's decision, and being born to insensitive parents seems to be a price worth paying for a worthwhile life.

This objection points to a more general problem with Wasserman's claim that prospective parents need to be able to justify the decision to bring someone into existence before they do so. In everyday life people make lots of decisions which do not make anyone worse off than they otherwise would have been, and do not expose them to being harmed. We don't typically think that, unless they are able to justify their decisions, it is wrong for them to act in these ways. Wasserman thinks that the decisions of prospective parents are different because they will cause the children to experience hardships in life, and so require justification. But why should we think that parents need to justify their decision *before* bringing another into existence? If Alma does not think of Beatrice's well-being before bringing her into existence, but the goods in Beatrice's life

nonetheless outweigh the bads, then there is a justification for Alma's decision. From Beatrice's perspective, it does not seem to matter whether Alma was aware of this justification before bringing her into existence. What's important is that the goods in Beatrice's life actually do outweigh the bads; whether or not this played a role in Alma's decision-making process does not seem relevant to determining whether her decision was permissible.

This discussion connects to a larger debate about the relevance of one's intentions to the permissibility of one's actions. In my view, intentions are not relevant to permissibility. In the next section I will discuss this issue at length, defend the view that intentions are not relevant to permissibility, and consider the implications that this has for intention-based solutions to the Non-Identity Problem.

## Intentions and Permissibility

In the last section I suggested that the intentions with which one acted cannot render one's action impermissible.<sup>25</sup> If this is right, it seems to pose a problem for all of the intention-based solutions in the literature on the Non-Identity Problem, since they tend to claim that Alma's decision to bring Beatrice into existence would be *wrong* because of the intentions and attitudes displayed in her decision. In this section I will explain in more detail why I do not think one's intentions are relevant to the permissibility of one's action, and sketch an outline of how we might recast the intuitively appealing features of Wasserman's account in somewhat different terms. In the next chapter I will provide a solution which does this, in an attempt to avoid the difficulties we have seen for the other solutions discussed so far.

Perhaps the most famous objection to the claim that intentions are relevant to permissibility comes from Judith Thomson, who asks us to consider what it would be like to respect intention-dependent rights of the sort that Wasserman invokes. Thomson claims that respecting such rights would implausibly require agents to "look inwards" before deciding what to do. To take a modified version of her example, imagine Alma coming to you and asking whether or not it would be permissible for her to bring Beatrice into existence now. On Wasserman's view, the appropriate response would be: "well, it depends on what your reasons would

---

<sup>25</sup> Although the arguments I make in this section support the general claim that one's intentions are never relevant to the permissibility of one's action, for the purposes of this thesis it is not necessary to defend such a strong claim. It is enough to make the weaker claim that intentions are not relevant to permissibility in the particular way, or in the particular cases, that advocates of intention-based solutions claim.

be. If you would bring her into existence now in part because her life will contain more goods than bads, then it is permissible; but if you would bring her into existence now for some other reason – say, because you think it will annoy your parents – then it is impermissible.” Thomson reacts to this sort of response by saying: “[w]hat a queer performance this would be! Can anyone really think that [Alma] should decide whether [to bring Beatrice into existence] by looking inward for the intention with which [s]he would be [doing it]?” (Thomson, 293). Simply introspecting to see why you would be doing something does not seem to affect that status of what you are about to do. In order to decide whether or not you should do something, you should consider how it will affect others; this seems to be what is relevant to deciding whether or not an action is permissible, not the intentions with which it is performed.

This becomes especially clear when we view the situation from the position of those affected by the action. Frances Kamm asks whether the potential victims of an action have a right which would make it impermissible for the agent to act unless she acts on certain reasons. She argues that “potential victims have no such right and that they can point to no grounds affecting them as reasons why the act should not be done” (Kamm, 168-9). In the abstract, most people share the intuition that acting on poor intentions is bad, but from the victim’s perspective it is not clear why the reasons on which the agent acts should matter, since there is nothing a victim could point to which reflects a difference in what has been done to her. This makes it difficult to see why we should agree that if Alma doesn’t act on certain reasons she will have violated Beatrice’s rights. If Alma will only bring

Beatrice into existence by acting with a poor intention, we might imagine Beatrice saying, as before, “I waive my right to be brought into existence with a good intention.” Given that her life is worth living this seems like a perfectly reasonable thing for her to say, yet it undermines the claim that Alma’s decision would be wrong because of her unfitting intentions.

It seems clear the intentions with which a person acts are relevant to moral evaluations of some sort. Kamm and Thomson do not dispute this. Their claim is rather that intentions are not relevant to the deontic assessment of an action as right or wrong, permissible or impermissible. Intentions *are* relevant, however, to the aretaic assessment of an action as praiseworthy or blameworthy, morally good or morally bad. In order to see this difference, it is helpful to consider T. M. Scanlon’s distinction between the deliberative and critical use of moral principles. He says:

As guides to deliberation, moral principles answer a question of permissibility: ‘May one do X?’ They also explain the answer by identifying the considerations that make it permissible or impermissible to do X under the circumstances in question... In what I will call their critical employment, however, a principle is used as the basis for assessing the way in which an agent went about deciding what to do on some real or imagined occasion. Used in this way, it provides the basis for answering a question of the form, ‘In deciding to do X under those circumstances, did Jones take the proper considerations into account and give them the right weight?’ (Scanlon, 22).

On Scanlon’s view, whether or not Alma has acted permissibly depends on whether or not there are considerations, such as harming Beatrice or violating impersonal moral requirements, which would weigh against the decision to bring Beatrice into existence. In the first chapter I argued that it is difficult to identify any such considerations. However, even if it turns out that Alma would be acting



permissibly, this does not determine whether or not Alma would be doing anything blameworthy or morally bad by bringing Beatrice into existence, because the answers to these questions do not depend on the actual considerations that are present, but rather the way in which Alma went about making her decision.

In the last section I agreed with Wasserman's claim that Alma's decision to bring Beatrice into existence now would not be made with appropriate intentions, since she seems indifferent to the well-being of the child she will bring into existence. If intentions cannot affect permissibility, however, we will need to understand the moral relevance of this in a somewhat different way than he does. Wasserman attempts to motivate the view that Alma's intentions can affect the permissibility of her action by presenting the following dilemma:

[W]e must either accept that the careless creation of an impaired child does not wrong that child or understand the child as having birth-rights whose satisfaction or violation depend on the attitudes with which his parents create him. The former denies the child a special complaint, while the latter suggest that, at least in the reproductive domain, morality prescribes not only how one should act but the reasons or motives with which one should act as well (Wasserman, 151-2).

Once we note the distinction between deontic and aretaic assessment, however, it becomes clear that these are not the only two options. It is perfectly consistent to deny that Beatrice had a birth-right that required Alma to act on the basis of certain reasons while maintaining that Beatrice can legitimately complain against Alma because of the reasons she took into consideration before bringing her into existence. This is because, although one's intentions are not relevant to the deliberative question of whether or not one's action was permissible, they are

relevant to the critical question of whether or not one's action was blameworthy or morally bad. Thus, we can vindicate the intuition that Alma has acted objectionably, without necessarily claiming that Alma has wronged Beatrice.

This, it seems to me, is the most promising way to address the Non-Identity Problem. It does not run into the difficulties that outcome-based solutions faced, and captures Wasserman's insight in a more theoretically supportable way. In the next chapter I will turn my attention to providing a more detailed account of how we can claim that Alma's decision merits a negative aretaic assessment although (for the reasons we have seen) it may not merit a negative deontic assessment. I will then discuss the implications of this account, and its relationship to the other proposed solutions to the Non-Identity Problem.

### **Chapter 3: Aretaic Assessment and the Non-Identity Problem**

In the last two chapters I raised problems for a number of views which attempt to solve the Non-Identity Problem by defending a deontic judgment that it would be wrong or impermissible for Alma to bring Beatrice into existence now. Given the difficulty of identifying any outcome-based reason (excepting effects on third parties) to object to Alma's decision to bring Beatrice into existence now, and that intention-based considerations cannot bear on the permissibility of Alma's action, I suggested that the proper judgment to make may be an aretaic assessment that Alma's action would be bad or blameworthy because of the attitude with which it would be performed, even though it may not be impermissible. In this chapter, I elaborate on and defend this claim. In the first section I explore the relationship between deontic and aretaic assessment by distinguishing five different senses in which an act can be wrong: the fact-, evidence-, belief-, moral-belief, and moral-evidence-relative senses of wrongdoing. In the second section I argue that when someone has done something wrong in the belief-, moral-belief, or moral-evidence-relative senses they are blameworthy, and that this accounts for a large part of the negative reaction to Alma's decision. In the third section I broaden my assessment of Alma's action by considering a number of factors that may help explain the intuition that she has done something morally problematic which do not strictly have to do with wrongdoing in the belief-, moral-belief or moral-evidence-relative senses. In the final section I discuss the limitations and implications of this solution, as well as its place in the literature on the Non-Identity Problem.

## Deontic Assessment and Wrongdoing

In the last chapter I argued that while the intentions with which one acted cannot affect the deontic status of whether or not one's action was permissible they can affect the aretaic status of whether one's action was good or bad, praiseworthy or blameworthy. Perhaps the most salient question for this proposal is: how can we make the judgment that Alma's action would be bad without claiming that she has done anything wrong? In order to respond to this challenge, it is necessary to distinguish five different senses in which an act can be wrong. Once we have done this it will become clear that the claim that Alma may not deserve a negative deontic assessment does not entail the claim that she would do nothing wrong by bringing Beatrice into existence. Instead, it entails the more specific claim that Alma would not be doing anything wrong in the sense that is tied to deontic assessment. Derek Parfit says acts are:

*wrong* in the *fact-relative* sense just when this act would be wrong in the ordinary sense if we knew all of the morally relevant facts,

*wrong* in the *belief-relative* sense just when this act would be wrong in the ordinary sense if our beliefs about these facts were true,

and

*wrong* in the *evidence-relative* sense just when this act would be wrong in the ordinary sense if we believed what the available evidence gives us decisive reasons to believe, and these beliefs were true (Parfit 2012, 143).

He later adds that acts are

*wrong* in the *moral-belief-relative* sense just when the agent believes this act to be wrong in the ordinary sense (*Ibid*, 148).

There are two distinct cases in which an action can be wrong in the evidence-relative sense. Sometimes an agent has evidence about

descriptive facts, for example, that if she performs action  $\Phi$  it will have outcome  $\Psi$ . Other times, an agent has normative evidence, for example, that it would be wrong to  $\Phi$ . In order to distinguish between these two cases, it is useful to introduce a fifth sense in which an act can be wrong.

Acts are

*wrong* in the *moral-evidence-relative sense* just when the available philosophical evidence gives the agent decisive reason to believe that performing the act would be wrong in the ordinary sense.

The fact- and evidence-relative senses of wrongdoing come apart in cases in which there is overwhelming evidence that one outcome will result from an action, but a different result in fact comes about. These cases have the following structure:

Consider a first-order morality  $M$  which opposes the causing of disasters... and which pays no attention to knowability. Agent  $\phi$ s, and this starts up a causal chain that leads to a disaster; but the chain includes indeterministic elements whose outcomes were highly improbable. Not only did Agent not expect the disaster: nobody at that time could have foreseen it; conditions making it inevitable or even likely did not exist in the world at that time. Still, according to  $M$ , Agent acted wrongly in  $\phi$ ing (Bennett 1998, 50).

In this example, there is a sense in which Agent did something wrong, and a sense in which she did not. Assuming that it is wrong to cause a disaster, Agent did something wrong in the fact-relative sense, because she caused a disaster. This judgment holds even though there was no way for her to know that her action would have disastrous effects, and no one could have foreseen the outcome at the time she acted. However, there is also a sense in which Agent did something entirely unobjectionable. There was no evidence available before she acted that her act would be wrong, so her action is not wrong in the evidence-relative sense.

Jonathan Bennett proposes two tests for determining whether an action is wrong in the evidence-relative sense:

[A] well-enough-informed bystander could in principle have advised against, condemned, or deplored at the time of acting; and an agent who has acted wrongly should have at least a theoretical chance of learning from the wrongness of his action, seeing how he could have done better and may do better in future (*Ibid*, 50-1).

At the time of acting, a well-enough informed bystander would not have had any reason to object to Agent's action, and it's not clear that there is anything that Agent could learn that would help her act better in the future. She only did something wrong in the fact-relative sense because a freak accident occurred; in the evidence-relative sense her action was unobjectionable.

The belief- and evidence-relative senses of wrongdoing come apart in cases in which the agent has mistaken views about what her evidence supports. Consider, for example, a case in which a very bad doctor wants to kill her patient. The doctor gives the patient a drug which she believes will kill her. In fact the drug cures the patient, as all of the evidence available to the doctor suggested it would (Parfit 2012, 144). Because the doctor's act does not harm anyone, and a well-enough informed bystander would not have reason to object to it (since all of the available evidence suggested that it would not harm anyone), the doctor's action is not wrong in the fact- or evidence-relative senses. However, many people still share the intuition that the doctor has done something bad by performing an act that she unjustifiably believed would harm another. We can capture this judgment by noting that the doctor's action was wrong in the belief-relative sense, since she gave the patient the drug with the belief that it would kill her.

We can see how the moral-belief-relative sense differs from the belief-relative sense of wrongdoing by noting that the doctor might have different beliefs about the moral status of her act. She might believe that she is morally permitted to try to kill the patient (perhaps because she is a retributivist and the patient has wronged her in some way) or she might believe that it would be wrong to kill the patient, but try to do it anyway (perhaps because she just really doesn't like the patient). The fact that the act seemed justified to her in the first case might affect how badly we feel about her action; whereas in the second case she knowingly does something wrong, in the first case she merely makes a moral mistake.

Finally, we can note that the doctor's moral beliefs might be more or less reasonable, depending on how well they respond to the moral evidence she considers. There may be some philosophical evidence that it would be permissible for her to kill one patient in order to save five others, but there is virtually no philosophical evidence that it would be permissible for her to kill a patient when doing so will have no good effects at all. If she acts on the moral belief that it would be permissible for her to kill a patient when doing so will have no good effects, she will have done something wrong in the moral-evidence-relative sense.

With these distinctions in hand we can now turn to two tasks. First, we need to figure out which senses of wrongdoing attach to which moral judgments, such as impermissibility, moral badness, and blameworthiness. Then, we need to figure out in which, if any, of these senses Alma's action would be wrong. This will provide us with the clearest view of what the appropriate assessment of Alma's action is.

Deontic judgments about whether or not an act was permissible most naturally concern wrongdoing in the fact- and evidence-relative senses. There is a long-standing debate about whether the claim that an act either was or was not permissible is a claim about what was or was not morally supported by the facts of the case, or whether it is a claim about what the agent had decisive reason to believe given the morally relevant facts that were knowable to her before acting.<sup>26</sup> I do not wish to comment on that debate here. For present purposes, it is enough to note that claiming that an action is impermissible requires claiming either that it is wrong in the fact-relative sense, or that the agent had decisive reason to believe that her act would be wrong in the fact-relative sense.

My claim that Alma's action may not merit a negative deontic assessment is best understood, then, as the claim that Alma may not be doing anything wrong in the fact- or evidence-relative senses by bringing Beatrice into existence. If the arguments considered in the first two chapters are correct, it is difficult to identify any morally relevant facts which would bear against Alma's decision to bring Beatrice into existence. If we cannot identify any such facts, then her action would not be wrong in the fact-relative sense. And if her act will not be wrong in the fact-relative sense, it's hard to imagine how her act could be wrong in the evidence-relative sense. She has evidence that Beatrice will have a life worth living, and – if the arguments discussed in the first chapter are correct – this means that Beatrice will not be harmed. A well-enough informed observer, who knew that Alma would not harm Beatrice by bringing her into existence, would

---

<sup>26</sup> Cf. Parfit (2012), Thomson (1986), Bennett (1998).



have no reason to condemn her decision before she acts. This implies that, if her act is not wrong in the fact-relative sense, it is not wrong in the evidence-relative sense either. If we cannot identify any morally relevant facts which would bear against Alma's decision to bring Beatrice into existence, then we cannot claim that a negative deontic judgment of Alma's decision would be warranted.

I suspect that many people will object to the claim that a well-enough informed observer would have no reason (excepting effects on third parties) to condemn Alma's decision before she acts. The reason that this intuition is so recalcitrant, it seems to me, is that most people believe that Alma would be doing something wrong. We might respond to this claim, as Boonin does, simply by restating the arguments considered in the first two chapters and claiming that, because the claim that Alma would be doing something wrong in the fact-relative sense doesn't seem theoretically supportable, we should give up the intuition that she would be doing something wrong. However, I think we can do better. We have already seen that the belief-, moral-belief and moral-evidence-relative senses of wrongdoing come apart from the fact- and evidence-relative senses. If it could be shown that Alma's act was wrong in one of these former senses, we would be able to make sense of the intuition that her decision to bring Beatrice into existence would be objectionable while remaining attentive to the fact that she may not be doing anything wrong in the fact- or evidence-relative senses. I will turn to this task in the next two sections. First, though, it will be useful to consider what moral judgment attaches to the claim that an act is wrong in the belief-, moral-belief, or moral-evidence-relative senses.

## Beliefs and Blameworthiness

Perhaps the most well-recognized negative aretaic judgment is that of blameworthiness, which includes reasons for remorse and indignation. Derek Parfit claims that “[w]hat is most important here is what, when acting, people believe” (*Ibid*, 143). For Parfit, “when some act is wrong in the *belief-relative* sense, because this act would be wrong if the agent’s non-moral beliefs were true, this fact makes this act blameworthy” (*Ibid*). He later claims that performing an act that is wrong in the moral-belief-relative sense, because one believed that one would be doing something wrong, can also make an act blameworthy (*Ibid*, 149). Holly Smith provides a useful illustration of this in her analysis of lying. She says:

[I]n the case of lying, it could be argued that what is genuinely objectively wrong is making an assertion that misleads the person who hears it; what is subjectively wrong is making an assertion in the belief that it is false and will mislead; and what is blameworthy is performing an act that one believes to be subjectively wrong (Smith 2010a, 89).

On this analysis, it is possible to be blameworthy for lying even if one’s lie does not actually succeed in misleading the other person. This happens, for example, when someone makes an assertion that she believes will mislead another person, but the assertion fails to mislead. Because the assertion was made with the belief that it would mislead and be wrong, the agent is blameworthy, even though she was mistaken and did not do anything objectively wrong. A similar judgment holds for the doctor who gives her patient a drug in the belief that it will kill her. Because the doctor does not succeed in killing her patient, the status of her act is different than it would have been if she had killed her patient. It is not

impermissible to give patients drugs which cure them (as the doctor does), but it is impermissible to give patients drugs which kill them (as the doctor believes that she does). The fact that the doctor does not succeed in doing what she believes to be wrong changes what she does in the fact-relative sense, so it affects the deontic status of her action. But it does not change what she has done in the belief-relative sense, so it does not remove the patient's grounds for indignation. The doctor performed an act which was wrong in the belief-relative sense, so she is blameworthy.

It is possible that part of what undergirds the intuition that Alma's action would be wrong is the feeling that she would be doing something wrong in the belief-relative sense. Given that Alma hasn't reflected on the fact that she will have a different child if she waits, and that she therefore will not harm anyone, she does not have any moral evidence that having Beatrice now is not the wrong thing to do. So when her friend comes to her and says: "[y]ou should think not only of yourself, but also of your child. It will be worse for [her] if you have [her] now" (Parfit 1984, 358-9) this should strike her as correct. If she decides to have Beatrice now anyway, without any further justification, she would be acting objectionably. This is because she would be doing something which she believed to be wrong (or should have believed to be wrong, given her philosophical evidence).<sup>27</sup> Although, depending on whether or not she will actually harm Beatrice or violate an impersonal moral requirement, she may not be acting

---

<sup>27</sup> I won't say much about blameworthiness and ignorance here, but it seems to me that Alma cannot absolve herself of blameworthiness by simply ignoring her friend's challenge. This means that she will have *some* belief about the moral status of her action, and her belief will be indexed to the philosophical evidence she has considered (or else be a case of culpable negligence).

impermissibly, she would be blameworthy for her act, since she would be doing something that is wrong in light of her moral beliefs and evidence. We can thus justify our negative feelings towards her, regardless of whether or not her act was permissible.

Relativizing the assessment of Alma's action to her moral beliefs and evidence allows us to see how her action can be bad even if she does not actually do anything impermissible. It is enough that she has reason to believe that she would be doing something wrong, but acts anyway.<sup>28</sup> However, I suspect many will be skeptical of the claim that our negative feelings simply have to do with the fact that Alma's act would be wrong in light of her moral beliefs and evidence, since it leaves open the possibility that if she believed that her act would not be wrong and had moral evidence that supported this (perhaps she was familiar with the Non-Identity Problem) then there may be no objection to her action. If we found that we would still have a negative reaction to Alma's decision, even if her act was not wrong in light of her moral beliefs and evidence, then the account I am offering would be very limited.

In order to respond to this challenge it simply needs to be shown that we would feel differently about Alma's action if she acted with different beliefs and evidence. To this end, it will be useful to consider two different ways in which Alma might hold the belief that she would not be doing anything wrong by

---

<sup>28</sup> Holly Smith (2010b) has argued that blameworthiness depends on what an agent actually believes, not on what it would be reasonable for her to believe. I use the language 'has reason to believe', simply to indicate that she will either have this belief, or be culpable for not forming it, since it is what the philosophical evidence she considers supports.

bringing Beatrice into existence: she might be aware of the Non-Identity Problem, and the fact that she will not harm Beatrice by bringing her into existence, or she might not believe that she will be making Beatrice worse-off than another child she easily could have. First, consider a case in which Alma knows that she would bring a different child into existence if she waits, and that this means she won't harm the child she chooses to bring into existence now. In this case, we might imagine Alma responding to her friend by saying: "I thought that choosing to have a child now would be bad for my child at first too, but then I read about it and it turns out that I won't be making anyone worse off if I choose to have a child now. Since I won't be making anyone worse off, and I want to have a child now, I am going to do it." At the very least, it seems to me, we would feel *better* about Alma in this case than we would in the 14 Year Old Girl case as it is originally presented. This suggests that at least part of our negative judgment of Alma's action has to do with her beliefs about what she is doing. Since Alma is a 14 year old girl, we do not think that she has investigated her decision in this sort of detail, and this fact affects our reaction to her decision.

Still, there might be lingering doubts about whether or not it would be okay for someone who had researched this issue in great detail to choose to have a child who is less well-off than one they easily could have had instead. What could explain our hesitation about these cases? Dan Moller has suggested that the fact that we are all fallible moral agents who are often mistaken about whether our acts are right or wrong may play an important role in generating this intuition. He says:

Suppose that you are considering performing some act A that you are worried might be wrong. Perhaps someone has presented you with an argument which purports to show that A is morally objectionable. Suppose, further, that after due deliberation you ultimately conclude that the argument fails, as do all the other anti-A arguments you know of. Is this the end of the matter? Should your deliberations be at an end after responsible consideration of the available arguments? The surprising answer is ‘No’ (Moller, 425).

The reason that engaging in moral deliberation is not enough is that we are fallible moral agents that often make mistakes about moral matters. Given that this is the case, it may be bad to perform an action that has a small chance of being very wrong, even though one’s best moral evidence indicates that the action is likely to be entirely permissible. Moller straightforwardly extends this to the Non-Identity Problem. He imagines someone in Alma’s situation – let’s call her Sally to keep the cases separate – engaging in philosophical deliberation about the problem in much the same way that I have in this thesis. Sally knows that there are moral principles that, if true, would explain why it would be wrong for her to conceive now, but judges it slightly more likely that those principles are mistaken. She decides to conceive now because she thinks that her action is probably permissible, and it is more convenient for her to conceive now. If what I have said so far is correct, I suspect that we will feel better about Sally’s decision than we do about Alma’s. If we still feel that Sally’s action is problematic, we need to be able to explain why. Moller’s explanation is as follows:

If nothing else, Sally should take seriously the possibility that in conceiving now she would be doing something deeply wrong; this should be of great concern to her, especially given that conceiving now, by hypothesis, does not do much to promote her interests. Ignoring this risk, as we must do if we are to claim that it does not provide us with at least some reason not to conceive now, seems to give insufficient weight to avoiding wrongdoing (*Ibid*, 438).

The judgment that one has responded inappropriately to conditions of moral risk is not an adequate ground for making a negative deontic assessment of Sally's action, since deontic assessment depends on whether or not her action was wrong in the fact- or evidence-relative senses. However, it does seem like this could contribute to a negative aretaic assessment of her action. It might also help explain why the way in which Alma related to her child would be objectionable. Given her moral and epistemic limitations, Alma ought to have sought to learn about the effects that her action would have on others and taken measures to avoid doing something deeply wrong. The fact that she does not do this provides a ground for judging that the way in which she related to Beatrice was objectionable. Even if it turns out that, luckily for her, her act is not wrong, we can still criticize her decision because of the way in which it was made.

Notice, though, that inappropriately responding to conditions of moral risk is not exactly the same as doing something that is wrong in the belief-, moral-belief, or moral-evidence-relative senses. It thus seems to me that the proper assessment of Alma's action may not be the fairly narrow judgment of blameworthiness, but the slightly broader judgment of moral badness. In the next section I will argue that this is the case. Before turning to this, though, it may be useful to consider a second case in which Alma might reasonably hold the belief that it would not be wrong for her to bring Beatrice into existence. To this end, contrast the 14 Year Old Girl case with that of Sharon Duchesneau and Candy McCullough, a deaf lesbian couple who live near Gallaudet University in Washington DC, the world's first liberal arts university for the deaf. In 2001,

Sharon and Candy chose to have a deaf child by way of artificial insemination by donor. They gave the following reasons for their decision:

- deafness is an identity, not a medical affliction that needs to be fixed
- the desire to have a deaf child is a natural outcome of the pride and self-acceptance many people have of being deaf
- a hearing child would be a blessing, whereas a deaf child would be a special blessing
- they would be able to be better parents to a deaf child than to one who could hear
- the child would grow up to be a valued member of a real and supportive deaf community (Parker, 279).

Sharon and Candy seem to believe that deaf members of a deaf community are not worse off than other members of society. But we might specify that they are wrong about this and, in fact, the child they bring into existence will be equally as poorly-off as Beatrice would be. This would make the structure of the two cases similar: both Alma and Sharon and Candy choose to have a less well-off child when they easily could have had a better-off child instead (indeed, Sharon and Candy seem to have gone out of their way to avoid having a non-hearing impaired child). Yet it seems to me that we would feel better about Sharon and Candy's decision than we would about Alma's. The relevant difference seems to be that Sharon and Candy do not believe that they will make their child worse off, and this belief is grounded in their own experience as deaf members of a deaf community.<sup>29</sup> Even if they turn out to be wrong about this, they are not blameworthy, since they do what is right in light of their moral beliefs and

---

<sup>29</sup> Appropriately responding to their philosophical evidence may explain why we would feel better about Sharon and Candy's decision than someone who made the same decision, but had no experience of deafness and did not know what quality of life her deaf child would experience.



evidence. By contrast, even if Beatrice turns out to be better off than any other child Alma could have had, Alma will still be blameworthy, since she did what was wrong in light of her moral beliefs and evidence.

It seems to me that this plays an important role in generating the intuition on which the Non-Identity Problem is based. Almost everyone who is presented with the 14 Year Old Girl case reacts negatively to Alma's decision, and this is something we need to be able to explain. We can accomplish this by noting that, while Alma's decision may not be impermissible since it may not be wrong in the fact- or evidence-relative sense, it would be blameworthy, since it would be wrong in the belief-, moral-belief, or moral-evidence-relative sense. In the next section I will suggest that, while this goes a considerable way to explaining our reaction to Alma's decision, it is in some ways limited, as other factors may play an important role in generating this reaction. I will argue that our reaction to Alma's decision is sensitive to her intentions as well as her beliefs and moral evidence, and that this might change the appropriate aretaic assessment of Alma's action from a judgment of blameworthiness to a judgment of moral badness. In the concluding section I will explain why attending to aretaic assessment is important and how it might affect ongoing debates about the Non-Identity Problem.

## **Motivation and Moral Badness**

In order to get clear on what the appropriate assessment of Alma's decision would be, it may be useful to distinguish three different negative aretaic judgments. Sometimes we judge acts as blameworthy. As we have just seen, this judgment is appropriate when someone performs an act which is wrong in the belief-, moral-belief-, or moral-evidence-relative sense and they do not have an excuse. Other times we assess people's characters as bad or vicious. This judgment is appropriate when someone acts with ill-intentions or out of a vicious character trait. The category of morally bad action falls somewhere between these two: it differs from character assessments because "what is being assessed is not the agent's overall character but rather the quality of the particular piece of decision making that led to the action in question" (Scanlon, 27-8), and it differs from blameworthiness because it is not concerned with the fairly specific question of whether or not an agent did something wrong in the belief-, moral-belief-, or moral-evidence-relative-sense. Instead, judgments of moral badness assess "the way in which an agent went about deciding what to do on a given occasion" (*Ibid*, 3). It seems to me that this judgment best captures our negative reaction to Alma's decision to bring Beatrice into existence: the way in which she would be making her decision is objectionable, so she would be doing something morally bad.

I have already noted one reason why this judgment is preferable over one of blameworthiness: if we believe that there might be reason to object to Sally's action because she failed to respond appropriately to conditions of moral risk, we will have shifted focus away from the question of whether or not she did the right

thing in light of her moral beliefs and evidence, and toward the question of whether or not the way in which she went about deciding what to do was appropriate. Another reason to prefer the judgment of moral badness is that it is possible for an agent to have beliefs according to which their action is permissible, but act for reasons which have nothing to do with those beliefs. For instance, we can imagine that Sally would have chosen to have a child now regardless of how her investigation into the moral status of having a child now turned out, or that the fact that they believed their child would be well-off played no role in Sharon and Candy's decision (perhaps they wanted to have a deaf child to annoy their parents, and would have this child even if they were given conclusive evidence that doing so would be wrong). It seems to me that we would only feel better about Sally and Sharon and Candy's decisions than we do about Alma's to the extent that their beliefs that they wouldn't be harming anyone actually played a role in their decision making process.<sup>30</sup> This suggests that the negative reaction to Alma's decision does not simply have to do with the fact that she may be doing something that is wrong in light of her beliefs and evidence. It may have to do more broadly with the reasons that factored into her decision making process. This means that our assessment of Alma's action is not just sensitive to her beliefs and evidence; it is also sensitive to her intentions. Scanlon says:

When we say that a person did something intentionally, one thing we may mean is simply that it was something that he or she was

---

<sup>30</sup> It seems to me that Bonnie Steinbock's discussion of a case like Sally's is of this sort (Steinbock, 168). This may explain why she does not believe that knowledge that she will not be harming anyone can make Sally's action better.

aware of doing or realized would be a consequence of his or her action... But we also use ‘intention’ in a narrower sense. To ask a person what her intention was in doing a certain thing is to ask her what her aim was in doing it, and what plan guided her action – how she saw the action as promoting her objective. To ask this is in part to ask what her reasons were for acting in such a way – which of the various features of what she realized she was doing were features she took to count in favor of acting in such a way (*Ibid*, 10).

If the argument that I made in the previous chapter is correct, intentions in the narrow sense are not relevant to the deontic assessment of actions, but they are relevant to the aretaic assessment of whether an action was morally good or bad. Given that we would feel differently about Alma’s action if she had different intentions, it seems to me that the most adequate aretaic assessment of Alma’s action is one which references her beliefs, moral evidence, and intentions.

We can thus revisit Wasserman’s claim that “parents’ reasons for having a child with certain characteristics will play a crucial role in determining the acceptability of their decision and the plausibility of any complaint from the child they create” (Wasserman, 136). I disagree with the first part of this statement and agree with the second part. In my view, the acceptability – in the sense that is connected with permissibility or deontic status – of parents’ procreative decisions is determined by the effects their action will have on the child or others. Unless it can be shown that Alma’s action will harm Beatrice or violate an impersonal moral requirement, her action may be acceptable in this sense. However, I agree with Wasserman that the way in which parents’ reason about bringing children into existence will play a crucial role in determining the plausibility of complaints about their decision. Even if Alma’s decision is permissible, others may object to how she made it, which is a matter of her beliefs, moral-evidence and intentions.

## **Implications**

Almost everyone who is presented with the Non-Identity Problem shares a clear and visceral negative reaction to Alma's decision. Insofar as this reaction is not easily dismissed as crazy or misguided, it requires explanation. The account that I have offered can explain both why we have this reaction and why we ought to have it, and it can do this without running into the problems that most attempts to justify a negative deontic judgment of Alma's decision have run into. My account can also play an important role in explaining and unifying our judgments about a number of related cases. Those who argue that there is no adequate way to make a negative deontic judgment of Alma's action are often forced to accept that procreative decisions which look seriously problematic would not be wrong. Appealing an account of aretaic assessment like the one I have offered can help these solutions explain why we feel better about some decisions to bring a less well-off person into existence than we do about others. My view is compatible with nearly every proposed solution to the Non-Identity Problem, and especially helpful for strategies which seem to require us to make significant revisions to our pre-theoretical intuitions.

There is, of course, an important reason that most philosophers have focused on trying to figure out the proper deontic assessment of Alma's action: many difficult questions which depend on the Non-Identity Problem, such as whether or not there is a moral obligation to undergo preconception testing, whether it is morally permissible to bring people with disabilities into existence, and whether it is permissible to genetically select for traits that may make the child

worse off on the whole, seem to hinge, not on the aretaic assessment of the way in which Alma made her decision, but on the deontic assessment of the morally relevant facts of the case. There is a way in which this seems to make my solution secondary and beside the point. However, I think that we ought to be hesitant to treat these as entirely separate questions, since our views about aretaic assessment often affect the conclusions about deontic assessment that we are prepared to accept. To illustrate this, consider one case which is often thought to be problematic for those who accept the implausible conclusion:

*Angela and Betty:* Angela is pregnant. Her doctor discovers that she has a condition that will result in her baby being mildly retarded unless she takes a medication for two months. Angela does not want to take the medication because it is mildly inconvenient. She does not take it and, as predicted, her baby is born mildly retarded.

Betty wants to get pregnant. She is on medication and knows that, if she gets pregnant while on the medication, her baby will be born mildly retarded. Going off the medication is not a feasible option. Luckily, she only needs to take the medication for two months. She considers waiting until she is off the medication, but does not want wait because it would be mildly inconvenient. She gets pregnant right away and has a baby which, as predicted, is born mildly retarded (cf. Steinbock, 169-70).

Most people agree that both Angela and Betty act objectionably. This has traditionally been seen as providing evidence that we should accept the No Difference View, according to which there is no difference in the moral status of the two actions. This, in turn, has put pressure on those who accept the implausible conclusion, because accepting the implausible conclusion seems to require rejecting the No Difference View. This is because Angela, but not Betty, has harmed her child. By not taking the medication, Angela has caused her child to be worse off than she otherwise would have been. This is not true of Betty

because there is no way that the child Betty brings into existence could have been born without the mild retardation. So, those who accept the implausible conclusion must claim that Angela, but not Betty, has acted objectionably.

My solution draws attention to another dimension along which we might evaluate this case. Insofar as Angela and Betty act with similar beliefs, moral evidence and intentions, their actions will merit similar aretaic assessments. We might accept the No Difference view at this level, but reject it at the deontic level, because whether or not an agent has harmed someone will affect the proper deontic assessment of their action. Introducing this distinction takes pressure off those who accept the implausible conclusion, and forces those who accept the No Difference View to either provide an account of harm which yields equal deontic assessments of the two actions, or explain why the fact that one agent has harmed someone and the other has not does not make a difference to the deontic assessments of their actions.

Thus it seems to me that providing an account of the proper aretaic assessment of Alma's action is important for both explaining and unifying our moral judgments, and for determining the proper deontic assessment of her action. It forces us to reconsider the intuitions we appeal to in arguments about the Non-Identity Problem, and the moral judgments we think they support.

## **Conclusion**

Moral philosophers have been debating about the Non-Identity Problem for over thirty years. I do not purport to have resolved the debate in this thesis. What I hope, rather, is to have contributed to the debate in a way that might help it progress. To this point the debate seems to have been construed almost entirely as a debate about whether or not it would be permissible for Alma to bring Beatrice into existence. By pointing out that this is not the only judgment we might make about her decision, and that the intuitions which generate the Non-Identity Problem are not sensitive to fine distinctions among different types of moral judgment, I hope to have added something of value to the discussion. In at least some cases, the common intuition may not be that Alma has done something wrong, but that she has done something bad. The account I have offered can explain why we have this reaction, and why our negative feelings about Alma's decision are justified.



## Bibliography

- Adams, Robert M. (1979) "Existence, Self-Interest, and the Problem of Evil," *Nous* 13 (1):53-65.
- Benatar, David. (2006) *Better Never to Have Been: The Harm of Coming into Existence* (Oxford University Press).
- Bennett, Jonathan. (1978), "Maximizing Happiness" in R. I. Sikora and B. Barry (eds), *Obligations to Future Generations*, Temple University Press, 61-73.
- (1998) *The Act Itself* (Oxford University Press).
- Boonin, David. (forthcoming) *The Non-Identity Problem* (February 2012 draft).
- Buchanan, Alan, et al. (2000) *From Chance to Choice: Genetics and Justice* (Cambridge University Press).
- Glover, Jonathan (1992) "Future People, Disability, and Screening," in Peter Laslett and James S. Fishkin, eds., *Justice between Age Groups and Generations*, New Haven: Yale University Press, pp. 129-30.
- Hanser, Matthew. (1990) "Harming Future People", *Philosophy & Public Affairs* 19: 47-70.
- (2008) "The Metaphysics of Harm," *Philosophy and Phenomenological Research* LXXVII (2): 421-450.
- Hare, R. M. (1998), "On Possible People," *Bioethics* 2(4): 279-293.
- Harman, Elizabeth. (2004) "Can We Harm and Benefit in Creating?" *Philosophical Perspectives* (18): 89-113.
- (2009) "Review: David Benatar. *Better Never to Have Been: The Harm of Coming into Existence*," *Nous* 43 (4):776-785.
- Heyd, David. (1992) *Genethics: Moral Issues in the Creation of People*. Berkeley: University of California Press.
- Kagan, Shelly and Vallentyne, Peter (1997) "Infinite Value and Infinitely Additive Value Theory," *Journal of Philosophy* 94 (1), pp. 5-26.
- Kamm, F.M. (2008) "Terrorism and Intending Evil", *Philosophy and Public Affairs* 36: 157-86.
- Kavka, Gregory S. (1982) "The Paradox of Future Individuals," *Philosophy and Public Affairs* 11 (2):93-112.

- Kumar, Rahul. (2003) "Who Can Be Wronged?" *Philosophy & Public Affairs* 31 (2): 99-118.
- McMahan, Jeff. (2013). "Causing People to Exist and Saving People's Lives," *Journal of Ethics* 17 (1-2): 5-35.
- McNaughton, D. and Rawlings, P. (1998) "On Defending Deontology," *Ratio* 11 (1): 37-54.
- Parfit, Derek. (1976) "On Doing the Best for Our Children," in Michael D. Bayles (ed.), *Ethics and Population*. Cambridge, MA: Schenkman, pp. 100-115.
- (1982) "Future Generations: Further Problems," *Philosophy and Public Affairs* 11 (2): 113-172.
- (1984) *Reasons and Persons*. Oxford: Clarendon Press of Oxford University Press.
- (2012) *On What Matters* (Oxford University Press).
- Parker, Michael. (2007) "The Best Possible Child," *Journal of Medical Ethics* 33: 279-83.
- Roberts, M. A. (1998) *Child Versus Childmaker: Future Persons and Present Duties in Ethics and the Law*. Lanham, Maryland: Rowman & Littlefield.
- (2002) "A New Way of Doing the Best That We Can: Person-Based Consequentialism and the Equality Problem," *Ethics* 112 (2), pp. 315-350
- (2009) "The Nonidentity Problem," *The Stanford Encyclopedia of Philosophy*.
- Roberts, M. A. and Wasserman, D. T. (2009) *Harming Future Persons*. (Springer).
- Ross, W.D. (1930) *The Right and the Good* (Oxford University Press).
- Savulescu, Julian. (2001) "Procreative Beneficence: Why We Should Select the Best Children," *Bioethics* 15 (5-6):413-426.
- Scanlon, T. M. (2008) *Moral Dimensions*. (Belknap Press of Harvard University Press).
- Schwartz, Thomas. (1978) "Obligations To Posterity," in R. I. Sikora and B. Barry (eds), *Obligations to Future Generations*, Temple University Press, 3-13.
- Singer, Peter. (1999) *Practical Ethics*, 2<sup>nd</sup> ed. Cambridge University Press.
- Shiffrin, Seana. (1999) "Wrongful Life, Procreative Responsibility, and the Significance of Harm," *Legal Theory* 5: 117-148.

- Smith, Holly. (2010a) "Subjective Rightness," *Social Philosophy and Policy* 27 (2):64-110.
- (2010b) "The Moral Clout of Reasonable Beliefs," In Mark Timmons (ed.), *Oxford Studies in Normative Ethics, Vol. I*. Oxford University Press, 1-25.
- Smolkin, Doran. (1999) "Toward a Rights-Based Solution to the Non-Identity Problem," *Journal of Social Philosophy* 30 (1):194-208.
- Steinbock, Bonnie. (2009) "Wrongful Life and Procreative Decisions," in *Harming Future Persons*. Eds. M. A. Roberts and D.T. Wasserman.
- Thomson, Judith Jarvis. (1986) *Rights, Restitution, and Risk: Essays in Moral Theory*, Harvard University Press.
- (1991) "Self-Defense," *Philosophy and Public Affairs* 20 (4):283-310.
- Velleman, J. David. (2008) "Persons in Prospect," *Philosophy and Public Affairs* 36 (3):221-288.
- Wasserman, David. (2005) "The Nonidentity Problem, Disability, and the Role Morality of Prospective Parents," *Ethics* 116: 132-152.
- Woodward, James. (1986) "The Nonidentity Problem." *Ethics* 96 (4): 804-831.