

Spring 1-1-2013

Who Should Exist: A Welfare-based Solution to the Non-identity Problem

Duncan Purves

University of Colorado at Boulder, duncan.purves@gmail.com

Follow this and additional works at: https://scholar.colorado.edu/phil_gradetds



Part of the [Ethics and Political Philosophy Commons](#)

Recommended Citation

Purves, Duncan, "Who Should Exist: A Welfare-based Solution to the Non-identity Problem" (2013). *Philosophy Graduate Theses & Dissertations*. 39.

https://scholar.colorado.edu/phil_gradetds/39

This Dissertation is brought to you for free and open access by Philosophy at CU Scholar. It has been accepted for inclusion in Philosophy Graduate Theses & Dissertations by an authorized administrator of CU Scholar. For more information, please contact cuscholaradmin@colorado.edu.

Who Should Exist: A Welfare-based Solution to the Non-identity Problem

by

Duncan Purves

B.A, University of Missouri-Columbia, 2006

A thesis submitted to the

Faculty of the Graduate School of the

University of Colorado in partial fulfillment

of the requirement for the degree of

Doctor of Philosophy

Department of Philosophy

2013

This thesis entitled:
Who Should Exist: A Welfare-based Solution to the Non-identity Problem
written by Duncan Purves
has been approved for the Department of Philosophy

Alastair Norcross

David Boonin

Michael Huemer

Michael Tooley

Date _____

*The final copy of this thesis has been examined by the signatories, and we
Find that both the content and the form meet acceptable presentation standards
Of scholarly work in the above mentioned discipline.*

Abstract

Purves, Duncan (Ph.D., Philosophy)

Who Should Exist: A Welfare-based Solution to the Non-identity Problem

Thesis directed by Associate Professor Alastair Norcross

I propose a solution to what has been dubbed the *Non-identity Problem*. The Non-identity Problem arises for the moral assessment of actions that affect both the quality of life as well as the identities of the people who will come to exist in the future. To see how the problem arises, consider the case of a prospective mother. Because of a peculiar, temporary condition, if this woman conceives a child now, then the child she conceives will suffer from a serious disability. She will name this child Jane. If the woman waits two months to conceive, her condition will subside, and the child she will conceive will not have the disability. She will name this child Jonathan. Jonathan will have a better life than Jane because he will not suffer from the negative effects of the serious disability. It seems that conceiving Jane rather than Jonathan would be wrong. But note that Jane, though she would have a worse life than Jonathan, would have a life worth living. Also, Jane will only exist if the woman conceives now, since Jonathan will be conceived if the woman waits two months. Because Jane would not be worse off, it seems that she has no complaint about being conceived with a disability. We are in this way pressed to conclude that the woman does not act wrongly by conceiving a disabled child now rather than

conceiving a non-disabled child later. The challenge of avoiding this absurd conclusion constitutes the Non-identity Problem.

I propose that we can avoid the absurd conclusion apparently entailed by the Non-identity Problem by appealing to a standing moral reason to promote the good, where the good includes all of the people who will or may come to exist. If we accept the existence of this moral reason, then the fact that Jonathan would have a better life than Jane generates a moral requirement to conceive Jonathan in the absence of other moral considerations. I argue that one can accept my proposal while retaining her other ordinary moral commitments, including the moral relevance of significant personal sacrifice and of constraints against harming.

Acknowledgments

I am indebted to many people for their generous assistance in the course of writing this book. I am most grateful to my advisor Alastair Norcross for his written comments and constant encouragement throughout the writing process. David Boonin also provided extensive written comments on earlier drafts. His challenging feedback contributed greatly to the quality of the final project. I also thank my other committee members, Benjamin Hale, Michael Huemer, and Michael Tooley for their helpful comments on earlier drafts. Others who have supported me throughout the process of writing this book include Chris Heathwood, whose seminar provided the inspiration for what would become my dissertation, Brian Talbot, my diligent editor Joe Fraley, and Ryan Jenkins. Each of the people mentioned here have inspired, advised, and mentored me at some point in the last six years. I thank you all.

Contents

I. Identity Affecting Actions and the Non-Identity Problem	1
I.1 How Can an Action be Identity-Affecting?.....	3
I.2 The Non-identity Problem	6
I.2.1 The Immediate Version	6
I.2.2 The Eventual Version	11
I.3 Criteria for a Solution to the Non-identity Problem	12
I.4 Why We Should Care about the Non-identity Problem	17
I.5 Going Forward: Finding a solution to the Non-identity Problem	19
II. Harm-Based Solutions to the Non-identity Problem	21
II.1 Harm: Comparative, Non-comparative, De Dicto, or Event-based	22
II.1.1 The Comparative Accounts of Harm	25
II.1.2 The Non-comparative Account of Harm.....	26
II.1.3 De Dicto Accounts of Harm.....	26
II.1.4 The Event-based Account of Harm.....	27
II.1.5 Supplemental Criteria for a Successful Harm-based Solution	28
II.1.6 Harm-based Solutions: Some Motivations and Some Shared Features	29
II.2 Rejecting (P1): Comparative Accounts of Harm	30
II.2.1 The Counterfactual Comparative Account of Harm	31
II.2.2 Closest World CCAH.....	32
II.2.3 Contextualist CCAH	34
II.2.4 Counterfactual Comparative Harm-based Solution 1: Rejecting Genetic Essentialism.....	38
II.2.5 Counterfactual Comparative Harm-based Solution 2: Benatar’s Solution.....	41
II.2.6 Counterfactual Comparative Harm-based Solution 3: Roberts’ Solution.....	44

II.3 Rejecting (P2): The Non-comparative Account, De dicto Account, and Event-based Account of Harm.....	46
II.3.1 Objections to the Counterfactual Comparison Account of Harm	47
II.3.1.A Preventative Harm.....	48
II.3.1.B Harms That Fail to Make the Victim Worse-off	51
II.3.1.C Excessive Multiplication of Harms	58
II.3.1.D Death	62
II.3.1.D.a Thomson’s Reply.....	66
II.3.1.D.b The Value of a Life	68
II.3.1.D.c The Timing Challenge	71
II.3.1.D.d The Time of Death’s Harmfulness	75
II.3.1.E Asymmetries in Harming and Benefitting (Hanser and Shiffrin)	80
II.3.2 Alternative Accounts of Harm and Their Solutions: Non-comparative Accounts, De dicto Accounts, and Event-based Accounts of Harm	81
II.3.2.A Non-Comparative Accounts	83
II.3.2.B The De dicto Account of Harm	92
II.3.2.C The Event-Based Account of Harm.....	100
II.3.2.C.a Death and Losing Goods.....	103
II.3.2.C.b Death and Deprivation.....	105
II.3.2.C.c Non-regrettable Disabilities and Collapse	107
III. Non-Harm-Based Solutions to the Non-identity Problem.....	113
III.1 Rejecting P4: Rights-Based Solutions (Wronging Without Harming)	114
III.1.1 The Moral Relevance of Consent	117
III.1.2 Markie and Smolkin’s Solution: Rights-violation without Harm	122

III.1.3 Tooley’s Solution: Future Rights Violations Solution.....	130
III.1.4 Reiman’s Solution: The Appeal to Fairness	139
III.2 The Failure-Based Approach	146
III.3 Rejecting P5: Doing Wrong without Wronging: Non-person-affecting Solutions	146
III.3.1 Kavka’s Solution: Making the World Worse by Creating People with Limited Lives	149
III.3.2 Temporal Indexing and Maximizing Consequentialism.....	156
IV. On the Existence of a Moral Reason to Promote the Good: The Beginnings of a Solution to the Non-identity Problem	159
IV.1 Preliminary Terms and Distinctions	164
IV.1.1 Preliminary Arguments about ‘The Good’	164
IV.1.2 Three Kinds of Goods.....	165
IV.1.3 Against the Impersonal Good	171
IV.1.4 Welfare, Pro tanto, and the Good	179
IV.1.5 Preliminaries about the Modality of Persons.....	182
IV.1.6 Preliminaries Concerning the Moral Relevance of Possible People	185
IV.2 Unrestricted Pro tanto, Restricted Pro tanto, and Solving the Non-identity Problem	191
IV.3 The Good of Welfare	196
IV.4 Pro tanto.....	198
IV.5 The Moral Relevance of Possible People? Choosing between Unrestricted Welfare and Unrestricted Pro tanto	207
V. In Defense of Unrestricted Pro tanto: A Solution to the Non-identity Problem .	210
V.1 Problems for RP	212
V.1.1 Creation Case 1: Happy Lives	212
V.1.2 Creation Case 2: The Problem of Suffering.....	215
V.1.2.A The Comparison Problem	219

V.1.2.B The Asymmetry	222
V.1.3 The No Difference View	227
V.2 Countervailing Reasons	239
V.2.1 Countervailing Reasons and Maximizing Consequentialism	240
V.2.2 Kinds of Countervailing Reason	245
V.2.3 Countervailing Reasons in Ordinary Morality	246
V.2.4 (At Least) Two Kinds of Countervailing Reasons	249
V.3 A Solution to the Non-identity Problem	251
V.4 Objections to the Welfare Argument from Problem Case	256
V.4.1 Rescue Cases	256
V.4.1.A Reply from Inconsistency	257
V.4.1.B Reply from Extreme Differences	259
V.4.1.C Reply from Distinguishing Between the ‘Quality of Life’ and the ‘Moral Value of Life’	263
V.4.2 Creation Cases	264
V.4.2.A Getting a Grip on the Notion of Significant Sacrifice	266
V.4.2.B The Right to Reproductive Autonomy	275
V.4.2.C A Dilemma: How Demanding is Morality?	279
VI. Practical Application	282
VI.1 Environmental Policy	282
VI.2 Preimplantation Genetic Diagnosis	287
VI.3 Reproductive Cloning	293
VI.4 Conclusion	297
References	298

I

Identity-Affecting Actions and the Non-identity Problem

Typically, an act that is morally wrong involves an agent performing an action that affects only individuals who already exist. The wrongness of this kind of action can be straightforwardly explained in “person-affecting” terms. When someone’s actions cause an existing person to suffer in some way, or when a person has a right of hers violated, we can explain the offensive action’s wrongness by appealing to its effects on the individual. The action makes the person worse-off than she otherwise would have been, or it violates her rights in some way, and this wrongs her. This is a “person-affecting” moral explanation of the act’s wrongness. There is, however, a certain category of actions that seems to elude moral assessment using person-affecting moral principles. This category of actions is the category of identity-affecting actions.

Identity-Affecting Action: An action a is identity-affecting if and only if either (i) it causes at least one individual to come into existence who otherwise would not have existed or (ii) it prevents at least one individual from coming into existence who otherwise would have existed.

Within the class of identity-affecting actions there are same-number actions and different-number actions. Same number actions have no effect on the number of people who will

exist. A same-number identity-affecting action, in effect, substitutes the coming-into-existence of one person with the coming-into-existence of another person. It satisfies the follow description:

Same-number Action: An identity-affecting action *a* is a same-number action if and only if, for each individual it causes to exist who otherwise would not have existed, it also prevents the coming-into-existence of an individual who otherwise would have existed.

Different-number identity-affecting actions, by contrast, affect both the identities and the number of people who will exist. For the sake of simplicity I will be primarily concerned with same-number identity-affecting actions (henceforth referred to simply as ‘identity-affecting actions’). The conclusions that I draw about same-number actions will have implications for different-number identity-affecting actions. I will say a bit about these implications in a later chapter. Identity-affecting actions give rise to the Non-identity Problem in cases when the action substitutes someone who is worse-off than the person for whom she is substituted.¹ Before turning to the Non-identity Problem I discuss the conditions that must obtain for an action to affect the identity of who comes to exist.

¹See Derek Parfit, (1976), pp. 100-02, Thomas Schwartz, (1978), pp. 3-13, and Robert M. Adams (1979). Each of these philosophers seems to have ‘discovered’ the problem independently around the same time.

I.1 – How Can an Action be Identity-affecting?

My parents might not have conceived a child when they in fact did. My mother might have left town that day for a month-long, work-related trip. Let's imagine that she went on the trip and that my parents conceived a child a month later. Would this child have been me?

An intuitive answer is that the child would not have been me. My parents could not have simply conceived a child at any time in their lives and ended up with one and the same child. How do we arrive at this intuitive answer? We must accept something like

Time-Dependence Claim: If any particular person had not been conceived within a month of the time when he was in fact conceived, she would never have existed.²

This is Derek Parfit's principle. It is a plausible principle. Although there is no uncontroversial theory of personal identity, the theories would seem to be univocal in supporting the Time-Dependence Claim. According to each major theory of personal identity, the child my parents would have conceived had they conceived a month later than they actually did would *not* have been me. He or she would have existed *instead* of me.³ I suspect that the Time-Dependence Claim is true because another claim about identity is true.

² Parfit, p. 352 (1984).

³ See Parfit, pp. 352-355 for a discussion of these theories, and how they deliver the verdict that the child my parents would have conceived a month later would not in fact have been me. This principle may need to be refined. There are "recycling" cases where the principle seems to entail that two different individuals are numerically identical

The Origin View: Each person has this distinctive necessary property – the property of having grown from the particular egg cell from which that person in fact grew.⁴

There are some good reasons accept the Origin View, because rejecting it has some strange implications. I have a younger brother. Obviously I am not identical to him. Suppose, however, that the egg that produced me had never been fertilized and that my parents instead only conceived a child that came from my brother's egg cell. It seems natural to say that only my brother comes to exist in this imagined scenario. Let's suppose further, contrary to The Origin View, that what is essential about a person is not that she come about from a particular sperm and egg, but that she have a certain kind of life history and that she have the personality traits that she in fact does. If my brother's life history and personality at the possible world in which his sperm and egg are fertilized, but mine are not, sufficiently resembled my life history and personality at the actual world, then on this condition, he would have had my essential properties. According to the history/personality view, then, in that world he would have been me. But then we have seemingly conflicting necessary truths about *de re* identity.⁵ At the actual

because they originate from the same recycled sperm and egg. Perhaps the Origins View may need to include a condition about spatio-temporal location in order to accommodate these cases.

⁴ Ibid, p. 352. Kripke (1980), and Forbes (1980) for discussions of the Origin View.

⁵ I explicitly talk about *de re* identity in order to avoid implying, falsely, that my brother could not have been identical with a particular description that I satisfy in the actual world and that he would have satisfied in a possible world. While I think it is necessarily true that $D \neq C$, it is not necessarily true that $D = \text{my parents' first-born son}$. In the world where I never come to exist but my brother does it would be false that $D = \text{my parents' first-born son}$, but true that $C = \text{my parents' first born son}$.

world Duncan \neq Cameron, but at some possible world Duncan = Cameron. But if my brother and I are non-identical then we are so at each possible world. So either Duncan = Cameron or Duncan \neq Cameron must be false, at pains of contradiction. We should think that originating from a particular egg cell is one essential property of individuals in order to avoid the implication that my brother might have been me.

Perhaps some will be content with the implication that my brother might have been me and also might not have been me. In other words, some might accept that identity claims are not necessary truths. My argument in the previous paragraph, in this case, fails to support the Origin View. One need not suppose the Origin View in order to yield the Time-Dependence Claim, however. Josh Parsons offers an interesting argument for the Time-Dependence Claim that appears to stand apart from any theory of personal identity.⁶ Suppose that my brother and I were in fact born a year apart. My parents could have refrained from conceiving me and waited a year before conceiving my brother. But surely in this case, just as in the actual world in which I come to exist, Duncan \neq Cameron. After all, even if my parents had not actually conceived me, I *could have* existed as Cameron's brother. But whether I am identical to my brother or not cannot depend on whether I come to exist. If we accept the Time-Dependence Claim (and I think we should accept it since it is entailed by almost every theory of identity; the plausible Origin View; as well as Parson's argument), then we should accept that identity-affecting actions are possible in a wide array of cases. Further, I think we should accept that there are actual instances of identifying actions and that they may be fairly common. People frequently make decisions that postpone or hasten the conception of a child by a month or more. Moving to a new city, taking a

⁶ Parsons (2003)

new job, ending a relationship with a significant other, taking a vacation, are all ordinary actions that might alter the timing of conception by a month or more. In a later section I will say more about actual cases in which agents might perform identity-affecting actions.

1.2 – The Non-identity Problem

The Non-identity Problem arises in attempting to morally assess certain same-number identity-affecting actions. These are a subset of identity-affecting actions that determine both the identity and the quality of life of the person who comes to exist. There are at least two versions of the Non-identity Problem. It is worth presenting both versions, because, as we will see, not all purported solutions to the Non-identity Problem are equipped to solve both versions. One version of the problem has an immediate effect on the quality of life of the individual(s) who will come to exist, whereas the other has an eventual effect on the quality of life of the individuals who come to exist.

1.2.1 – The Immediate Version

The simplest version of the Non-identity Problem is the immediate version. For an example of the immediate version of the problem consider:

Betty: Betty takes her newborn baby for a checkup. The doctor says that there is some good news and some bad news. The bad news is that, as things now stand, the

baby is going to develop a disability. The doctor explains that the disability will be significant, meaning something that certainly diminishes one's quality of life in a non-trivial way (e.g., more like blindness than like color-blindness). Though the child will be disabled its life will be well worth living. In other words, it will be better than no life at all. The disability will be irreversible. The child will have it for life. Betty can, however, prevent the disability from ever developing in her baby if she simply gives it a pill once a day for the next two months. The pills will cost nothing for Betty since its cost is covered by her insurance. Betty fully understands both the minor cost to her of buying the pills and giving them to her child, and she understands the grave consequences for her child of not doing so. She decides not to give her child the pills. As a result her child becomes incurably blind.⁷

Wilma: Wilma's doctor tells her that she has a condition that will cause any child she conceives to suffer from a fairly serious disability if she conceives now. It will be incurably blind. The doctor also tells her that this result is not unavoidable. If she takes a pill every day for two months prior to conceiving a child, then the child she conceives will be fully sighted. She decides not to take the pill for two months prior to conceiving, in favor of conceiving immediately. As a result, her child is born incurably blind. She names it Pebbles. Note that had she opted to take the pill for two months the child she would have conceived would have been different from the one she actually conceives, since the sperm and egg that would have merged to create the

⁷ This Case is borrowed from Boonin (2008, pp. 127-128).

child would almost certainly have been distinct from the ones that merged to create her actual child.⁸

Our reactions to these two cases, I will assume, are more or less symmetrical. We react that what Betty does is badly wrong. We also react that what Wilma does is badly wrong. Now consider what will turn out to be an important difference between the cases: Betty's decision either to give or not to give the pill to her child has no effect on the identity of her child, whereas Wilma's decision either to take or not to take the pill determines *which particular child* she has. Wilma's decision, unlike Betty's, is an *identity-affecting* action. The wrong making feature of what Betty does is easy to identify. She acts wrongly because she harms her child by causing it to be worse off than it otherwise would have been. Betty's child would have existed whether or not she gave it the pills, and that child is worse off for her doing so. The same cannot be said for Wilma's child. By choosing unaided pregnancy, Wilma does not cause it to be the case that one and the same child is worse off than it otherwise would have been. After all, there are only two options with respect to Pebbles: blind existence or non-existence. Since it is stipulated that Pebbles has a life worth living, supposing that existing with a life worth living is not worse than non-existence, then (1) Pebbles is not made *worse-off* than she otherwise would have been by being caused to exist. Pebbles is not worse off, because either (i) she is in fact better off than she otherwise would have been (if non-existent individuals have welfare-levels of zero) or (ii) her welfare level would not exist, were she not to exist, and so could not be compared to her welfare level when she exists. On one plausible account of harm, (2) that Pebbles is not made worse-off

⁸ This case is modified from Hanser (2009, pp. 179-180).

than she otherwise would have been means she is not harmed. If we couple this assumption with the claims (3) that an act must harm someone if it is to wrong someone and (4) that an act that wrongs no one is not wrong, then it becomes difficult to see what is wrong with Wilma's action. From initially plausible claims (1)-(4) one can construct the following argument.

Non-identity Problem Argument (NIPA):

P1) Wilma's act does not make Pebbles worse off than she otherwise would have been.

P2) An act harms someone only if it makes them worse off than that person otherwise would have been.

C1) *From P1 and P2:* Wilma's act does not harm Pebbles.

P3) Wilma's act does not harm anyone else.

C2) *From P3 and C1:* Wilma's act harms no one. It is harmless.

P4) If an act does not harm anyone then it does not wrong anyone.

C3) *From C2 and P4:* Wilma's act does not wrong anyone.

P5) If an act does not wrong anyone then it is not wrong.

C4) *From P5 and C3:* Wilma's act is not wrong.⁹

⁹ Boonin (2008, pp. 129-131) explains the Non-identity Problem in a similar way. Obviously there are alternative formulations of the problem, but this formulation neatly sets up the options for solving the problem. The argument presents the problem in such a way that our challenge is clearly to find some justification for rejecting one of the premises.

One instance of the Non-identity Problem is the challenge of retaining our intuition that Wilma does something wrong in the face of a compelling argument concluding that she does not. The challenge, then, is to find a good reason to reject one of the premises of this argument for the conclusion that Wilma does nothing wrong. If it turns out that there are no good reasons for rejecting a premise then perhaps we should abandon our intuition. To solve the Non-identity Problem, then, one must identify a wrong-making property of Wilma's action (thereby rejecting one of the premises of the argument above) without committing oneself to implications that are more counterintuitive than the conclusion of NIPA. Let's now consider the eventual version of the problem.

1.2.2 – The Eventual Version

Suppose that we live in a society that has limited natural resources, and that we are choosing between two social policies. On policy A we will deplete our resources significantly over the next 100 years. On policy B we will conserve our resources over the next 100 years. On policy A the quality of life for those who exist over the next 100 years will be very slightly higher than it would be on policy B. On policy B the quality of life for those who exist over the next 100 years will be very slightly worse than on policy A, but the quality of life of those who come to exist 100 years from now will be very significantly higher than they would be on policy A. We can imagine that life will be very hard for those people who come to exist 100 years from now if we choose A, and that life will be very good for those people who come to exist 100 years from now if we choose B. Now let us suppose that life will be very hard for the people who exist 100 years from now under A, because we have made resources scarce and made the environment

dangerous. Under A, the effects of our depletion include contaminated water ways and radioactive infertile soil. These conditions result in starvation, radiation-related death, and disease-related suffering for many of the people who will exist 100 years from now under A. These conditions, in other words, will result in significant suffering and early death for many people 100 years from now if we choose A. Though life will be very hard for the people who exist 100 years from now, many of them dying prematurely, if we choose A, their lives will still be worth living. For this to be a version of the Non-identity Problem we must also suppose that the people who will exist 100 years from now if we choose policy A will be entirely distinct from the people who will exist 100 years from now if we choose policy B.

Most will judge that it would be seriously wrong for us to choose policy A over policy B, especially since we can do so at a relatively minor cost to those who exist at the time of the choice. We can construct a similar version of the NIPA from this case to the conclusion that we do nothing wrong by choosing A instead of B. Let's call this "NIPA 2."

Non-identity Problem Argument 2 (NIPA2):

P1) Policy A does not make future generations worse off than they otherwise would have been.

P2) An act harms someone only if it makes them worse off than that person otherwise would have been.

C1) *From P1 and P2:* Choosing Policy A does not harm future generations.

P3) Choosing Policy A does not harm anyone else.

C2) *From P3 and C1:* Choosing Policy A harms no one. It is harmless.

P4) If an act does not harm anyone then it does not wrong anyone.

C3) *From C2 and P4:* Choosing Policy A does not wrong anyone.

P5) If an act does not wrong anyone then it is not wrong.

C4) *From P5 and C3:* Choosing Policy A is not wrong.¹⁰

The challenge, again, is to find a premise of the argument that we can reject in order to avoid the conclusion that it is permissible to choose A rather than B. This is the eventual version of the Non-identity Problem.¹¹

1.3 – Criteria for a Solution to the Non-identity Problem

A good solution to the Non-identity Problem ought to solve both the immediate and eventual versions of the problem. It should explain why Wilma acts wrongly by conceiving Pebbles rather than a sighted child, and it should explain why it is wrong to choose policy A over policy B. I now turn to this and other criteria for a successful solution to the Non-identity Problem. There are five premises in NIPA that together entail the conclusion that constitutes the Non-identity Problem. To solve the problem, then, we must reject one of these five premises,

¹⁰ Boonin (2008, pp. 129-131) explains the Non-identity Problem in a similar way. Obviously there are alternative formulations of the problem, but this formulation neatly sets up the options for solving the problem. The argument presents the problem in such a way that our challenge is clearly to find some justification for rejecting one of the premises.

¹¹ The original version of this case is found in Parfit (1984, pp. 361-364).

thereby blocking the inference to the conclusion. We will see that each premise has been rejected by at least one philosopher.

Of course, simply rejecting a premise will not do. Each premise, on its own, seems fairly plausible. The Non-identity Problem wouldn't arise if one of them were obviously false. If one of NIPA's crucial premises was "*P**: *It is always permissible to kill someone if the victim is of a different race than you*" then solving the Non-identity Problem would be easy. The intuitive plausibility of the premises is what makes the problem so tricky. This brings us to our first criterion for a solution to the Non-identity Problem.

Non-*ad hoc*ness: A successful solution must not be *ad hoc*. In other words, it must give intuitively plausible reasons for rejecting a premise that are independent of the fact that doing so solves the Non-identity Problem.

According to the Non-*ad hoc*ness Criterion, these 'independent' reasons to accept a particular premise must, at least at first blush, be strong enough for us to reject some premise. Consider one possible solution to the problem: *We ought to reject P3 because if P3 is false, then NIPA is unsound*. This 'solution' does have one thing going for it. It solves the Non-identity Problem if it is true (more on this shortly). There are two troubles with allowing *ad hoc* solutions like the one described, however. First, doing so leaves us with no reason to reject one premise of NIPA rather than another. We can reject any premise of NIPA on the grounds that doing so blocks the inference to the implausible conclusion, but we want to know *which* solution is the

correct solution. Second, rejecting one of the premises of NIPA comes at a cost, because it does some violence to our moral intuitions. Each of P1-P5 is initially plausible, so, at least initially, we have some reason to retain each premise. In order to choose between rejecting one premise rather than another, thereby minimizing the violence done to our moral intuitions, we should prefer a solution that gives us a *reason* for rejecting a premise aside from the fact that rejecting that premise solves the Non-identity Problem. Whether we choose to reject one premise rather than another will depend on the independent plausibility of the reason(s) we have for rejecting each premise. This leads us to another criterion for a successful solution to the Non-identity Problem.

Intuitive Cost: A successful solution to the Non-identity Problem must not have implications that are more counterintuitive than the conclusion of NIPA.

Let's suppose that some solution to the Non-identity Problem is the most plausible among its competitors. This is not enough for it to be successful. Imagine that the best solution on offer avoids the conclusion that Wilma does nothing wrong by conceiving Pebbles, but it only does so because it entails that all acts of conception are wrong. This would not be a successful solution to the problem, because the intuitive cost of accepting the solution is too great. It avoids the implausible conclusion of NIPA only by forcing us to accept an even less appealing conclusion. Now it is not quite enough for a solution to have low intuitive cost and to be non-*ad*

hoc. It must also be sufficiently general. This brings us to our third criterion for a successful solution to the Non-identity Problem.

Generality: A successful solution to the Non-identity Problem will solve the most versions of the Non-identity Problem among the competing solutions.

Earlier I described both immediate and eventual versions of the Non-identity Problem. Other things being equal, if one of two solutions solves both versions of the problem, then that solution is preferable to one that does not. It is preferable because it allows us to retain more of our moral beliefs than the otherwise equally good solution. Ideally, our solution to the Non-identity Problem will allow us to retain *both* the belief that Wilma does something wrong by not taking the pills *and* the belief that the society acts wrongly by choosing A rather than B.

It is worth noting an important implication of the Generality Criterion. It implies that the *mere fact* that a candidate solution would solve at least one version of the Non-identity Problem counts in its favor, regardless of whether the solution has any independent motivation. The mere fact that a solution solves both the immediate and eventual versions of the problems counts in its favor to some degree. The challenge posed by the Non-identity Problem is, after all, to figure out how to *avoid* a conclusion that we are loath to accept. This means that, for each premise of NIPA, the fact that the premise, conjoined with the other premises of the argument, entails such an implausible conclusion counts against it. We simply do not yet know which premise we ought to reject. Of course, it may turn out that two or more premises warrant rejection on independent

grounds, but for the reasons discussed under the *Non-ad hocness* Criterion, we should prefer rejecting as few of them as possible. The fact that each premise, conjoined with the other premises of the argument, entails the implausible conclusion counts against it means that *any* solution to the Non-identity Problem has at least one thing going for it: it avoids the implausible conclusion. This will be relevant when we consider the class of what I call “non-solutions” to the Non-identity Problem. These are solutions that, rather than attempting to block the inference to the conclusion of NIPA by rejecting one of its premises, attempt to convince us that the conclusion is the product of a sound argument. Because these views do not solve any version of the Non-identity Problem, they will initially be at an intuitive disadvantage compared with views that do solve at least one version of the problem.

One more implication of the Generality Criterion is important to note. Some solution might successfully refute one of the premises of NIPA, but fail to refute some revised version of the premise. If, in its revised version, NIPA still arrives at its problematic conclusion, then the solution in question is not *sufficiently general*. It is therefore unsuccessful. A successful solution must block the inference to the problematic conclusion of NIPA by demonstrating that all revisions of one of P1-P5 are false. Otherwise it constitutes a failed attempt at a solution. It does not solve the Non-identity Problem; it merely demonstrates that the argument for the problem must be revised.

1.4 – Why We Should Care about the Non-identity Problem

Some might think that I have described a problem in ethics that, while interesting theoretically, is really a non-issue in practice. What, if any, bearing could the Non-identity Problem have on our practical ethical decisions in the actual world? I think it will turn out that whether we can solve the Non-identity Problem and which solution we ultimately settle on will have enormous implications for practical ethics. This is because, supposing that we accept the Time-Dependence Claim, and especially if we accept the Origins View, there are many policies that we might adopt and many actions that we might perform *in the actual world* that will affect both the identity and quality of life of the people who will come to exist.

Consider, for instance, how close the circumstances of the actual world are to those of the society described in the eventual version of the Non-identity Problem. The imagined society must choose between a policy of depleting their limited resources and a policy of conserving those resources. This is, plausibly, the choice facing our society today. Even if there is some doubt that our choice to deplete or conserve our resources *today* will affect the quality of life for future generations, it is beyond doubt that there will come a time when a society will face such a choice. When society faces a choice between depleting and conserving its natural resources that will affect the quality of life of future generations, this choice will be an identity-affecting choice. It will determine *who* comes to exist. To see that this is so we need only consider the impact that a relatively minor shift in environmental policy would have on the daily lives of ordinary people. As part of a policy of conservation, a society's government might impose heavy tolls on people who drive to work. This could have the effect of significantly reducing the number of people who drive to work every day. But people must get to work somehow. As an

alternative, commuters might ride their bike, take the subway, or use the public bus system to get to work. Someone who used to leave work at 8 am and got home at 6 might leave for work earlier and get home later. Changes like these would undoubtedly have an effect on lots of aspects of someone's daily schedule, including the person's sexual schedule. Because such a policy might have a significant effect on the timing of conceptions, it may turn out to be a choice that affects both the quality of life and the identities of the people who come to exist.

The Non-identity Problem and our solution to it could also have moral implications for the use of some contemporary procreative technologies. Preimplantation genetic diagnosis (PGD) and preimplantation genetic screening (PGS) are kinds of technology, the application of which can affect both the identity of the people who come to exist as well as the quality of life of the people who come to exist. PGD refers to a process of testing embryos fertilized using in vitro fertilization for genetic abnormalities before pregnancy. This is typically done when at least one of the parents has as a known genetic abnormality. Embryos that test positive for the abnormality will be rejected in favor of ones that do not. PGS uses similar technology, but it is used by parents neither of whom is known to have a genetic abnormality. If we are sympathetic with the Origins View, then we will likely think that who comes to exist in cases where PGD or PGS is used depends on which egg/sperm combination is selected for pregnancy. This fact means that the choice of embryo will be an identity-affecting choice. Furthermore, assuming that the genetic abnormalities that PGD is designed to test affect the quality of life of the person who has them, then using PGD to determine which embryo is chosen for pregnancy will be an action that affects the quality of life of the person who comes to exist. PGD has all of the relevant features of the cases that give rise to the Non-identity Problem. It would seem, then, that the solution to the Non-identity Problem that we choose will have moral implications for the use of

PGD and PGS in pregnancy. As this technology becomes more widely available, these implications become more pertinent.

The first two kinds of real-world choices for which our solution to the Non-identity Problem might have moral implications involve far-reaching policies about society's conduct. There are also many more pedestrian choices that affect both the identity and the quality of life of the person who will come to exist. For instance, parents might decide to wait a month to conceive so as to move to a neighborhood that is significantly better for raising children. The child they would raise in the new neighborhood would have a better quality of life than, but also be distinct from, the child they would conceive if they stayed where they are. Even very small changes can potentially have massive effects on the timing of conceptions. How late bars remain open on weekends might affect who conceives with whom and when. It may determine whether someone goes home to conceive anyone at all. Coming to exist, it may turn out, is much more precarious than most imagine.

Identity-affecting actions and policies are much more abundant in the actual world than we realize. It turns out that, far from being a mere philosophical oddity that is interesting in theory only, the Non-identity Problem, and our solution to it, can have wide-ranging implications for our actions and policies in the actual world. It is a problem that is interesting in theory and important in practice.

1.5 – Going Forward: Finding a solution to the Non-identity Problem

There are five kinds of solution to the Non-identity Problem corresponding to the five premises of NIPA. I will address solutions that reject P1 and P2 in chapter 2, and I will address solutions that reject P4 in chapter 3. I argue that solutions that reject P1, P2, and P4 are unsuccessful. These solutions fail either the Non-*ad hocness* Criterion, the Intuitive Cost Criterion, or the Generality Criterion for a successful solution. I will only say a bit about P3 here. Some might worry that P3 is false because conceiving Pebbles will be worse for Wilma or for society at large. It is certainly true that children with disabilities sometimes require a greater share of care and resources than their non-disabled counterparts. In the real world this is a valid concern, but for the purposes of this book I can simply stipulate that Wilma's choice to conceive Pebbles will be worse neither for her nor for society at large. Wilma is genuinely indifferent between having Pebbles and having a sighted child, even given the extra burden of raising someone with a disability. She possesses the resources necessary to raise Pebbles without burdening society in any discernible way.

In chapters 4 and 5 I develop my solution to the Non-identity Problem. My solution entails the rejection of P5, the premise which claims that an act is wrong only if it wrongs someone. Wilma's action, I argue, is wrong even though it wrongs no one because it makes the world worse in a way. It makes the world worse because it causes people to exist who are less happy than the people who might otherwise have existed. My solution is independently motivated (not *ad hoc*), it comes with little intuitive cost, and it is sufficiently general: it solves all versions of the Non-identity Problem.

II

Harm-Based Solutions to the Non-identity Problem

Each premise of NIPA, with the exception of (P3) (since it is a stipulation), has been challenged in the philosophical literature. Solutions that reject (P1) or (P2) are what I will call ‘harm-based solutions’. These solutions claim that, contrary to appearances, Wilma has in fact harmed Pebbles either by making her worse off than she otherwise would have been or by some other means. Those that reject (P1) maintain that Pebbles has been comparatively harmed by being created, because she has been caused to be worse-off than she otherwise would have been. According to these solutions, Wilma’s choice to conceive Pebbles makes Pebbles worse off than she otherwise would have been, and this fact explains why Wilma’s choice is wrong. A second kind of solution challenges (P2). These solutions concede that Pebbles has not been comparatively harmed by being created (she is not worse-off than she otherwise would have been), but they argue that she has been harmed in other ways and that these other ways are morally significant. Solutions that challenge (P2) invoke one of three accounts of harm in order to solve the Non-identity Problem: a *de dicto* account of harm; an event-based account of harm; or a non-comparative state-based account of harm. I will refer to them respectively as *de dicto* harm-based solutions, event-based solutions, and non-comparative harm-based solutions. I address harm-based solutions in this chapter. I will address non-harm-based solutions in chapter III. Non-harm-based solutions concede that Pebbles is not harmed in any morally relevant way by being created, but they identify some other wrong-making feature of her mother’s act that makes it wrong. Non-harm-based solutions take issue with (P4) and (P5) of NIPA.

This chapter has the following structure: In section II.1 I will distinguish between four accounts of harm: the comparative, *de dicto*, non-comparative, and event-based accounts. Each of these accounts, provided certain auxiliary claims, seems to be capable of solving the Non-identity Problem. In section II.2 I discuss attempts to solve the Non-identity Problem that reject (P1). These solutions appeal to the comparative account of harm. Section II.3 is divided into two main subsections. In the first half of II.3 I defend the Counterfactual Comparative Account of Harm against several common objections. I conclude that these objections can be overcome. In the second half of II.3 I canvas attempts to solve the Non-identity Problem that appeal to the *de dicto*, the non-comparative, or the event-based accounts of harm. I argue that these solutions to the Non-identity Problem, as well as the accounts of harm on which they are based, are mistaken. I conclude that the Counterfactual Comparative Account of Harm is correct and that there is no satisfactory harm-based solution to the Non-identity Problem presently on offer.

II.1 – Harm: Comparative, Non-comparative, De Dicto, or Event-based

Before canvassing the harm-based solutions to the Non-identity Problem that have been offered in the literature, it will be helpful to become clear about several concepts pertaining to the topic of harm. An account of the *nature of harm*, at least the kind of account that is relevant for purposes of solving the Non-identity Problem, will provide a set of necessary and/or sufficient conditions that must obtain for a person to suffer a harm in a morally relevant sense.¹² The kind of account of harm that is helpful in solving the Non-identity Problem must be an

¹² I use the terms ‘and/or’ because some accounts of harm only purport to provide a sufficient condition for someone’s being harmed (e.g. Harman, 2004).

account of *morally relevant* harm, because it must be employed to *solve* the Non-identity Problem. Solving the problem requires explaining why the identity-affecting actions that give rise to the problem are *wrong* (e.g., Wilma's act of conceiving Pebbles, and the Wealthy Society's choice of Depletion). Only an account of morally relevant harm is suited to that task. From hereon I will use the term "account of harm" as shorthand for "account of morally relevant harm."

It will be helpful, in order to avoid confusion, to draw a distinction early on between an *act* of harming and an *event* of suffering a harm. We normally use the term *harm* to refer to an action. Ordinarily we say that X harmed Y. However, I will take the event of suffering harm to be a more fundamental notion. I will suppose that X harms Y just in case X stands in the right relation to the event *Y's suffering a harm*. When I discuss accounts of harm in this chapter I will thus be presenting them as providing conditions that must obtain for an *event* of suffering a harm to occur. I will not discuss the further question about the relation that an agent must stand in to an event of suffering a harm in order for it to be true that the agent has harmed someone. I assume something significant in the rest of this chapter, however: if it turns out that Pebbles suffers a harm by coming to exist or that the Depletion generation suffers a harm, then I assume that Wilma and the Wealthy society stand in the right relation to their suffering a harm such that it is true that Wilma *harms* Pebbles and that the Wealthy Society *harms* the Depletion generation. Note that when I use the term 'suffer' I am not presupposing that harm always involves some particular mental or physical state. For instance, I leave open the possibility that someone can be harmed without experiencing anguish, physical pain, or emotional distress.

There is a further distinction to be drawn within the category of *accounts of harm*. Some accounts of harm provide a set of *necessary and sufficient* conditions that must obtain for a person to be harmed in a morally relevant sense. These views contend that a person suffers a harm if and only if those conditions obtain. Other accounts provide only a set of *sufficient* conditions. These accounts concede that there may turn out to be conditions, distinct from the ones proposed by the view, that result in a person's being harmed in a morally relevant sense. They simply contend that these conditions are not the whole story. Someone can be harmed in several ways, one of which is the way the view describes.¹³ This distinction between accounts of harm that are sets of necessary and sufficient conditions and those that are sets of sufficient conditions will be relevant when considering objections to particular accounts of harm. Accounts that have a biconditional structure (sets of necessary and sufficient conditions) can be critiqued from either direction of the biconditional. They fail if it can be shown either that the conditions they list are not necessary for a person to suffer a harm or that the conditions they list are not sufficient conditions for a person to suffer a harmed. Accounts that are conditional (sets of sufficient conditions) can be attacked from only one direction. They fail if and only if the conditions they list are shown not to be sufficient for harm. If a biconditional view is shown to be false because it is shown that the conditions it lists are not necessary for harm, the view might be reformulated into a conditional account according to which the conditions it lists are merely sufficient for harm. So long as an account's sufficient conditions remain intact, it is open to the defender of the account to revise it in order to attempt to solve the Non-identity Problem. With the distinction between conditional and biconditional accounts of harm in mind, let's turn to the

¹³ Non-comparative accounts of harm are sometimes described this way, as in Harman (2004).

four broad accounts of harm (comparative, non-comparative, *de dicto*, and event-based) that will be the subject of the rest of this chapter.

II.1.1 – The Comparative Accounts of Harm

Comparative accounts of harm, to determine whether someone suffered a harm, compare how someone *actually fares* in some outcome with how she fares in some relevant alternative. A natural way (but we will see that it is not the only way) to make this comparison is to compare her welfare level in each of those outcomes. For simplicity's sake, and because nothing important hinges on it, I will take the term 'welfare level' to refer to the *net balance* of intrinsic goods and bads that a person is enjoying either at a moment or over some interval of time. These 'intrinsic goods and bads' are the things that make one's life go well in and of themselves. Pleasure and pain, for instance, are popular candidates for intrinsic goods and bads, though many other candidates have been proposed. There are several versions of the comparative account of harm. Each of them proposes a different baseline comparison for determining whether a harm has occurred. On one account, someone suffers a harm in a morally relevant sense when her welfare level is caused to be lower than it was *previously*. According to another account, a person suffers a harm just in case her welfare level is caused to be lower than it *otherwise would have been*. Still another holds that a person suffers a harm just in case her welfare level is caused to be lower than it *ought* to have been. According to comparative accounts of harm then, to determine whether someone has suffered a harm we must compare how someone actually fares either with how she was faring previously, would have fared otherwise, or should have fared. I will discuss these competing comparative accounts in more detail in section II.2.

II.1.2 – The Non-comparative Account of Harm

Roughly, a person suffers a harm, according non-comparative accounts, if she is caused to be in a state that is bad for her. This state might not be worse than some state that she was in, would have been in, or should have enjoyed had the action not occurred. It need only be a state that is *objectively bad from the standpoint of one's interests*. Which particular states are to be included in the category “objectively bad from the standpoint of one's interests” will be discussed in II.3.2 along with other aspects of the non-comparative account.

II.1.3 – De Dicto Accounts of Harm

De dicto accounts of harm share one feature in common with the comparative accounts of harm. Like comparative accounts, they can maintain that someone suffers a harm just in case she is made worse off than she was, would have been, or should have been. Unlike comparative accounts, however, *de dicto* accounts hold that it is sometimes appropriate, when determining whether morally relevant harm has occurred, to compare how some individual actually fares with how some *distinct* individual fared previously, might have fared, or should have fared. How can harm occur in cases where someone is worse off than some distinct person would have been? Answering this question requires examining the *de dicto* component of the account. The *de dicto* view will be discussed in greater detail in II.2.2.3.

II.1.4 – The Event-based Account of Harm

According to the event-based account of harm, someone suffers a harm just in case she undergoes a ‘harming event’ with respect to some ‘basic good’. Someone undergoes a harming event with respect to a basic good if and only if she either loses some quantity of that good or is prevented from receiving some quantity of that good. The event-based account of harm is distinctive insofar as it holds that someone’s suffering a harm is independent of the badness (or worseness) of the state in which it leaves its victim. Though losing a basic good will typically result in someone being in a state that is worse for her than some alternative, it is the *losing* of the good rather than the badness of the resulting state explains *why* the person has suffered a harm.

The comparative, non-comparative, *de dicto*, and event-based accounts of harm will be the topics of the rest of this chapter. Each account has been appealed to in an attempt to solve at least one version of the Non-identity Problem. In the sections that follow I will provide some motivations for harm-based solutions generally, and then discuss the particular accounts of harm and their corresponding solutions. Each of these solutions, I will argue, falls afoul of either the Intuitive Cost Criterion, Sufficient Generality Criterion, or the Non-*ad hoc*ness Criterion for a successful solution to the Non-identity Problem. It is worth taking time to note that there are two ways that a harm-based solution might fail to satisfy these criteria.

II.1.5 – Supplemental Criteria for a Successful Harm-based Solution

As with each solution that I will consider, I will evaluate the harm-based solutions using the criteria for a successful solution to the Non-identity Problem discussed in Chapter 1. There are at least two ways that a harm-based solution might fail to satisfy these criteria. Either (i) the account of harm on which the solution depends fails to satisfy the criteria or (ii) the solution depends on additional premises that are independent of the account of harm to which the solution appeals, and these additional premises make the solution unable to satisfy the criteria. I mention these two ways in order to avoid confusion. I mentioned that I intend to argue both that the Counterfactual Comparative Account of Harm is correct and that it cannot be used to solve the Non-identity Problem. This might seem strange. It might seem that if a solution based on the Counterfactual Comparative Account of Harm fails to satisfy the criteria for a successful solution, then the most obvious explanation for this failure would be that the Counterfactual Comparative Account of Harm *itself* fails to satisfy the criteria. But, as I just noted, there are two ways that a solution might fail to satisfy the criteria. In the case of counterfactual comparative harm-based solutions, I will argue that their failure is due to problems with auxiliary claims that are independent of the Counterfactual Comparative Account of Harm. This allows me to reject the solution while consistently maintaining that the account of harm on which the solution is based is correct. *De dicto*, event-based, and non-comparative harm-based solutions, on the other hand, fail to satisfy the criteria for a successful solution *in virtue of* their dependence on a particular account of harm. That is to say, the *de dicto* and non-comparative accounts of harm *themselves* have problems satisfying the criteria for a successful solution, and this fact explains why the solutions on which they are based fail to satisfy the criteria.

II.1.6 – Harm-based Solutions: Some Motivations and Some Shared Features

With these important distinctions in mind, we can now examine the harm-based solutions to the Non-identity Problem. Harm-based solutions have received a reasonable amount of attention in the literature on the Non-identity Problem.¹⁴ It is tempting to appeal to harm in order to solve the Non-identity Problem for a couple of reasons. First, by appealing to harm, harm-based solutions remain faithful to the ordinary, folk explanation of why it is wrong for Wilma to conceive the blind child or for society to implement a policy of resource depletion. Most people, before being confronted with the peculiar features of identity-affecting actions, will say that Wilma's choice and the Wealthy Society's choice are morally wrong because the choices will *harm* the people who come to exist. Harm-based solutions allow us to retain this initial intuition. In this way harm-based solutions avoid the implication that there is anything *special* about the Non-identity Problem. They pick out a particular person who is affected by the identity-affecting action and say *of her* that she is wrongfully harmed by that action.¹⁵ By picking out a particular person who is the victim of the identity-affecting action, harm-based solutions explain the wrongness of the actions that give rise to the Non-identity Problem in exactly the same way that we explain the wrongness of ordinary harmful actions. Both types of actions are wrong because they harm a particular person. For these reasons, solutions to the Non-identity Problem that appeal to the moral consideration of harm are initially attractive. I now turn to a more detailed

¹⁴ I have in mind here Shiffrin (1999), Harman (2004), and Hanser (2008).

¹⁵ The *de dicto* harm-based solution, unlike other harm-based solutions, actually fails to identify a particular person who is harmed by an identity-affecting action. On this view the harm suffered is harm to a *role* or *office*.

discussion of each of the comparative, non-comparative, *de dicto*, and event-based solutions to the Non-identity Problem.

II.2 – Rejecting (P1): Comparative Accounts of Harm

It might seem obvious that coming to exist with a life worth living is not worse than non-existence. Therefore, it might appear that solutions to the Non-identity Problem that appeal to comparative accounts of harm are non-starters. There are, nonetheless, several ways to challenge (P1) of the argument for the Non-identity Problem by appeal to the comparative account of harm. Before turning to these comparative harm-based solutions I will lay out the comparative account of harm in greater detail. The following is a statement of a schematic version of the comparative account of harm.

Comparative Account of Harm (CAH): A's suffering a harm is A's being worse off in a way than A _____.

This schematic formulation is neutral between several competing comparative accounts of harm. How we fill in the blank will depend on which version of the comparative account of harm we accept. There are at least four variations of the comparative account of harm. The diachronic comparative account compares a person's actual welfare level with her welfare level *prior to the act* to determine if the person has been harmed by an action. On this view, I am

harmed by being kicked in the shins if and only if I am worse off as a result of being kicked than I was *before* I was kicked. Though this view is worth noting, it does not have much promise. It has trouble accounting for the phenomenon of preventative harm. For instance, suppose that my paycheck is intercepted and then cashed by my neighbor. It seems that I am harmed in this case even though I am not worse off than I was *before* my paycheck was intercepted. There is thus no way for the diachronic account to explain why I suffer a harm when my neighbor intercepts my check. Because the diachronic account of harm cannot account for preventative harm, it is unacceptable.

II.2.1 – The Counterfactual Comparative Account of Harm

According to the most popular comparative account, Q suffers a harm just in case Q is worse off in a way than Q *otherwise would have been*.¹⁶ This Counterfactual Comparative Account of Harm (CCAH) holds that P harms Q just in case P stands in the right relation to Q's suffering a harm.¹⁷ I will not take a position on what relation P must stand in to Q's suffering harm in order for it to be true that P harms Q. One obvious candidate is a causal relation. Provisionally, then, we can say that, according to CCAH, P harms Q if and only if P causes Q to suffer a harm, where Q's suffering a harm consists in her being worse off in a way than she otherwise would have been in. We now have a schema for when P harms Q according to CCAH. There is a further question, however, about what is required for someone to be "worse off in a

¹⁶ I include Bradley (2011), Thomson (2011), and Klocksiem (2012) among the recent defenders of the Counterfactual Comparative Account of Harm.

¹⁷ I borrow much of this discussion of the two-step process from Hanser (2008).

way than she *otherwise would have been*.” With what counterfactual state are we to compare a person’s actual welfare level when determining whether she has been harmed? Different versions of CCAH will give different answers to this question.

II.2.2 – Closest World CCAH

On one version of CCAH, Closest World CCAH, Q suffers a harm as a result of an action/event A/E if and only if Q’s state at the actual world is worse than Q’s state at the nearest possible world where A/E does not occur. The nearest possible world might be one in which the action/event simply does not occur or it might be one in which some *other* action/event occurs, bringing about a different outcome for Q.

Though Closest World CCAH is perhaps the most widely endorsed version of CCAH, it faces what some take to be a serious problem. Consider one problematic case:

Bobby Knight: suppose you witness the following scene at Texas Tech University: A member of the Philosophy department, passing Bobby Knight on campus, waves cheerily and says “Hey, Knight.” Bobby Knight, turning as red as his sweater, seizes the hapless philosopher around the neck and chokes her violently, while screaming obscenities. By the time Bobby Knight has been dragged away, the philosopher has suffered a partially crushed windpipe and sustained permanent damage to her voice box, as a result of which she will forever sound like Harvey Fierstein. Has Bobby Knight’s act harmed the philosopher? The intuitive answer is

obvious, and HARM seems to agree. The philosopher is much worse off than she would have been had Bobby Knight not choked her (unless, perhaps, she has always wanted to sound like Harvey Fierstein). But suppose we discover that Bobby Knight has recently been attending anger management classes. Furthermore, they have been highly successful in getting him to control his behavior. When he becomes enraged, he holds himself relatively in check. On this particular occasion (only the third violent outburst of the day), he tried, successfully, to tone down his behavior. In fact, if he hadn't been applying his anger management techniques, he wouldn't have choked the philosopher, but would rather have torn both her arms from her body and beaten her over the head with them. Since it took great effort on Bobby Knight's part to restrain himself as much as he did, it seems that the closest possible world in which he doesn't choke the philosopher is one in which she is even worse off.¹⁸

The apparent trouble for Closest World CCAH is that it seems to entail that Bobby Knight does not harm the philosophy professor. In light of this trouble, we may wish to adopt a different method for choosing the relevant alternative possible world with which someone's state, given an action/event, is to be compared when determining whether that individual has been harmed. On a competing view, the relevant alternative possible world is not always the *closest* possible world, but is rather determined by *conversational context*. Sometimes conversational context will pick out more distant possible worlds as the relevant alternative when

¹⁸ Example taken from Norcross (2005: 165-166).

determining whether someone has suffered a harm. I will refer to this view as Contextualist CCAH.¹⁹

II.2.3 – Contextualist CCAH

According to Contextualist CCAH, conversational context determines which possible world we compare with the actual world when figuring out if someone is worse off than she otherwise would have been as a result of some action/event. In some contexts, the utterance “X harmed Y by doing A,” picks out the closest possible world where X did something other than A. In other contexts, the same utterance, “X harmed Y by doing A”, might pick out a more distant (nomologically speaking) but more conversationally salient possible world where X did something other than A. The closest possible world will often be the most salient world in a given conversational context, but there is no guarantee that this will be so. Bobby Knight is an instance of a kind of case where the closest possible world is not the most salient world given the conversational context. In the case of Bobby Knight, the salient alternative appears to be a world in which Bobby Knight does the morally right thing, neither choking nor ripping the arms off of the helpless philosopher. Moral considerations enter the conversational context because they are salient in our thinking about the case. When selecting the possible world with which to compare an agent’s action, we do not consider only what they would *otherwise* have done but also what they *ought to have done*. There is nothing especially mysterious about the idea that context might determine the relevant alternative possible world with which the actual world is to be compared. The relevant alternative is simply *the possible world with which we compare the actual world*

¹⁹ Norcross (2005) and Bradley (2011) defend similar but distinct versions of Contextualist CCAH.

when thinking about a particular case. Ben Bradley calls the alternative possible world we have in mind when thinking about a particular case the ‘most similar’ world to the actual world. Different possible worlds will be the most similar to the actual world in different conversational contexts, Bradley suggests, because different ‘similarity relations’ will be fixed by different conversational contexts.²⁰ Contextualist CCAH is a view about the correct account of harm (Contextualist CCAH includes CCAH as one of its components), but it is also a view about the meanings of our utterances. In Alastair Norcross’ words:

I say that one act can be correctly described in one conversational context as a harming, and can be correctly described in a different conversational context as a benefiting. The reason why no contradiction is involved is that the claim ‘act A harmed person P’ can express different propositions in different contexts. On my suggested account of harm, to claim that act A harmed person P is to claim that A resulted in P being worse off than s/he would have been if the appropriate alternative to A had been performed. Given the context-relativity of the appropriate alternative, claims about harm (and benefit) have an indexical element. Just as ‘today is a good day to die’ can express different propositions in different contexts of utterance, so can ‘Smith’s act of will-writing harmed his son’.²¹

²⁰ Bradley (2011: 50-54).

²¹ Norcross, (2005: 169)

We can apply this Contextualist lesson to the case of Bobby Knight. When we claim that Bobby Knight harms the philosophy professor, our claim expresses the following proposition: Bobby Knight caused the philosophy professor to be worse off than she otherwise would have been by choking her *rather than leaving her entirely alone*. Of course the world in which Knight leaves the professor entirely alone is not the closest alternative possible world. The *closest* alternative possible world is one in which Knight tears her arms off of her body rather than leaving her entirely alone, but the world in which Knight leaves her alone is the most *salient* or *appropriate* world (in Norcross's words) with which to compare the actual world in normal conversational contexts. In Bradley's terminology, the world in which Knight leaves the professor entirely alone is most *similar* to the actual world, given normal conversational contexts. The Contextualist version of CCAH can be captured by the following:

Counterfactual Comparative Account of Harm (CCA): A's suffering a harm is A's being worse off given o , where o is the actual outcome, than A would have been in $\sim o$, where $\sim o$ is the relevant alternative outcome. Which world is picked out by ' $\sim o$ ' varies depending on conversational context.

I am inclined to accept Contextualist CCAH as an account of harm and as a thesis about the meanings of our utterances. Just as distinct instances of commonplace indexicals ('this', 'you', 'today') that are orthographically identical differ in meaning depending on the context in which they are uttered, so do instances of 'harm' and 'benefit'.

Some will be moved by cases like Bobby Knight to replace Closest World CCAH with Contextualist CCAH. The appeal of Contextualist CCAH extends further than this, however. Incorporating Contextualism allows the proponent of CCAH to offer a neat reply to some of the most compelling objections brought against the view. This strategic appeal of the Contextualist feature of this version of CCAH will be made apparent in later sections.

CCAH combined with a Contextualist component is a promising comparative account of harm, but some defenders of CCAH might be unmoved by the Bobby Knight example and thus unmoved to adopt Contextualism. I wish to avoid alienating defenders of CCAH who are unwilling to adopt a Contextualist component, so, henceforth, when I use the terms ‘Counterfactual Comparative Account of Harm’ or ‘CCA’ I intend them to refer either to Contextualist CCAH or Closest World CCAH. If either version of CCAH is true, we will shortly see, then we cannot solve the Non-identity Problem by appealing to harm. One point worth noting now is that I help myself in places to Contextualist considerations in order to reply to particular objections to CCAH. These replies are unavailable to Closest World CCAH. I turn to a more thorough defense of CCAH in section II.3.1, but I first discuss several harm-based solutions to the Non-identity Problem that are based on CCAH. These solutions contend that (P1) of NIPA is false, despite appearances to the contrary, because Pebbles and the Depletion generation are harmed in that they are made worse off than they otherwise would have been.²² I

²² One might reject (P1) on the grounds that it is incoherent to compare Pebble’s welfare, given that she does not exist, with what her welfare would have been were she to come to exist. The idea behind this worry is that individuals who never exist do not have any properties at all, including a welfare level. I relegate this objection to a footnote because the defender of the argument for the Non-identity Problem can easily sidestep it. (P1) simply states that Wilma does *not* cause Pebbles to be *worse off* than she otherwise would have been. Pebbles might fail to be worse off than she otherwise would have been either (i) because Wilma causes her to be better off than she otherwise would have been or (ii) because such comparisons are incoherent. If the relation *Pebbles being worse off than she otherwise would have been* cannot obtain between her when she exists and her when she does not, then the

conclude that these solutions do not succeed. CCAH cannot be used to solve the Non-identity Problem.

II.2.4 – Counterfactual Comparative Harm-based Solution 1: Rejecting Genetic Essentialism

Recall that an important step in justifying (P1) is the claim that Pebbles is a distinct person from the sighted child (let's call this child "Rocks") that Wilma would have conceived had she waited two months. This claim might seem to depend on some kind of genetic essentialism according to which one of the essential properties of our identity is the particular sperm and egg the fusion of which gave rise to us. One way to deny (P1) then is to deny that our identity is essentially tied to the particular sperm and egg from which we developed. Denying genetic essentialism eliminates an important barrier to maintaining that Pebbles is identical to Rocks even though she develops from a different sperm/egg pair than Rocks would have. Because genetic essentialism is false, the objection proceeds, there is no reason to deny that Pebbles and Rocks are the same person. Since Rocks would have been better off than Pebbles in fact is, and since there is no reason to deny that Pebbles is the same person as Rocks, we can maintain that *Pebbles* is worse off than *she* otherwise would have been by coming to exist now rather than by being conceived later as Rocks.²³

The main worry for this solution to the Non-identity Problem is that it fails to satisfy the Sufficient Generality Criterion for a successful solution. The trouble is that the solution shows

sentence "Pebbles is not worse off than she otherwise would have been" is true because it is simply claiming that such a relation does not obtain.

²³ Wrigley (2011) can be interpreted as defending a solution like this one.

that (P1) is false only if the truth of (P1) depends on the truth of genetic essentialism, but there are cases that seem to demonstrate that it does not.

A variant of the Wilma case seems to indicate that it is implausible to maintain that Pebbles is identical to Rocks even if we reject genetic essentialism.²⁴ In the original case Wilma must decide between conceiving a blind child now (Pebbles) and taking a pill every day for two months in order to conceive a sighted child later (Rocks). But imagine that Wilma must take a pill for twelve months before she will be able to conceive a sighted child. Though it might seem that the sighted child Wilma conceives two months later would be the same person as Pebbles in the original case, this is much less plausible in the twelve month variant of the case. The trouble is that Pebbles and Rocks could exist *at the same time* in the twelve month variation of the case. But clearly if Pebbles and Rocks were to exist at the same time in two different locations they would not be the same person. If Pebbles and Rocks are distinct individuals, given that Wilma conceives both of them, then they must be distinct individuals *even if* Wilma never conceives one of them. Denying this last claim commits us to the position that, in the world where Rocks but not Pebbles is conceived, Rocks is in fact identical to the individual who might have been his sister. The proponent of the view that Pebbles and Rocks are the same person faces an uncomfortable trilemma. She can either (i) maintain that Pebbles and Rocks are numerically non-identical in both the two month version and the twelve month version of the case, (ii) maintain that they are numerically identical in both the twelve month version and the two month version of the case, or (iii) attempt to maintain that Pebbles and Rocks are the same person in the two month version of the case but different people in the twelve month version of the case. Choosing

²⁴ The variation of the Wilma case I am about to offer is derived from a case discussed by Josh Parsons (2003: 148, 150).

(i) is to simply give up her view and to concede that Pebbles is not the same person as Rocks and so not caused to be worse off than she otherwise would have been. Choosing (ii) seems unacceptable as well; it is much harder to believe that Pebbles and Rocks are the same person in the twelve month case (given that the two of them could have existed at the same time in different locations) than to accept that Wilma does nothing wrong in conceiving Pebbles. It is worth pointing out one further counterintuitive implication of (ii). If Pebbles and Rocks are the same person in the twelve month version of the case, then all of the facts about Pebbles must also be facts about Rocks. But this is clearly not the case. Suppose for instance, that only Rocks is conceived in the twelve month version of the case, and that he registers at a community college. In this world, the following sentence is true: There is a world in which Pebbles registers for college (a world, unlike the actual world, in which she comes to exist and registers for college) and Rocks never registers for college. But that sentence must assert a contradiction if Rocks is numerically identical to Pebbles. It would be tantamount to asserting that there is a possible world in which Rocks both registers and does not register for college. But there is clearly no contradiction in asserting that there is a world in which Pebbles registers for college and Rocks never does. It seems, then, that someone who thinks that Pebbles and Rocks are numerically identical in Wilma's case should not choose option (ii). Choosing (iii), on the other hand, seems to result in the solution failing to satisfy the Non-*ad hoc*ness and the Sufficient Generality Criteria for a successful solution. It fails the Non-*ad hoc*ness Criterion because there appears to be no reason to maintain this asymmetry between the two and twelve month cases aside from the fact that doing so would allow us to solve the Non-identity Problem. Choosing (iii) means that the solution fails the Sufficient Generality Criterion, because option (iii) entails that while Wilma does something morally objectionable by causing Pebbles to be worse off than she otherwise

would have been in the two month case, she does not cause Pebbles to be worse off than she otherwise would have been in the twelve month case. But surely the implication that Wilma does nothing wrong by conceiving the blind child in the twelve month case is at least as counterintuitive as the implication that she does nothing wrong by conceiving the blind child in the two month case.

II.2.5 – Counterfactual Comparative Harm-based Solution 2: Benatar’s Solution

David Benatar has argued that coming to exist *always* harms those who come to exist. Everyone, he claims, is in fact worse off than they otherwise would have been by coming to exist. He arrives at this conclusion by arguing that the intrinsic bads that people suffer are *worse* for those who experience them than non-existence, but the intrinsic goods that accompany existence are *not* better for those who experience them than non-existence. Because the bads (e.g., pains) are worse for us than non-existence, and the goods (e.g. pleasures) are not better for us than non-existence, in all actual cases coming to exist is worse for someone than not coming to exist. Let us call this claim Pleasure/Pain Asymmetry²⁵. If the Pleasure/Pain Asymmetry is true, then it lays the foundation for a potential solution the Non-identity Problem when conjoined with the Counterfactual Comparative Account of Harm. The Pleasure/Pain Asymmetry entails that everyone suffers a harm in the counterfactual comparative sense by coming to exist, because it entails that everyone who comes to exist is thereby worse off than she otherwise would have

²⁵ Nothing important hinges on which axiology we choose here. Benatar talks about pleasure and pain as being the things that are intrinsically good and bad respectively for someone. However, his argument does not stand or fall with Hedonism. If the reader prefers, she can substitute desire-satisfaction/frustration or some other unit of welfare.

been.²⁶ It therefore entails that there is something morally bad about Wilma's choice to conceive Pebbles, because Wilma's action results in Pebbles suffering a harm. Of course, Benatar's view also entails that if Wilma were to conceive Rocks, this action would result in Rocks suffering a harm. It seems, then, that the *mere fact* that Wilma thereby harms Pebbles is insufficient to show that Wilma acts wrongly by choosing to conceive Pebbles rather than Rocks, since both options involve someone being harmed. Note, however, that Wilma's choice to conceive Rocks would be *less* harmful for Rocks than her choice to conceive Pebbles would be for Pebbles. If we add that Wilma has an obligation to *minimize* harm, then the fact that she fails to minimize harm by conceiving Pebbles rather than Rocks, combined with the Pleasure/Pain Asymmetry and CCAH does seem to solve the Non-identity Problem.

But is the Pleasure/Pain Asymmetry true? Benatar argues that positing the Pleasure/Pain Asymmetry is the only way to retain other intuitive asymmetrical judgments that we are loath to give up. I will not discuss Benatar's argument for the Pleasure/Pain Asymmetry in this section. I will simply point out that the Asymmetry, conjoined with the Counterfactual Comparative Account of Harm, is committed to two implausible positions. One of these is a moral position, and the other is a metaphysical position. The Asymmetry is metaphysically implausible because it flies in the face of all plausible views regarding the benefits of existence. The Asymmetry is implausible on the moral front, because it entails that all acts of conceiving are morally objectionable.

First consider the metaphysical problem for Benatar's solution. If one holds the view that how someone fares when she does not exist is incomparable to how she fares when she does

²⁶ Benatar (2008: 28-40)

exist (perhaps, as some claim, because someone does not have a welfare level when she does not exist) then one will likely hold that it is not possible to benefit or harm someone comparatively by causing her to exist. If, on the other hand, one holds the view that how well a non-existent individual fares *can* be compared with how she fares when she does exist, then one might maintain that it is possible to benefit and harm comparatively by causing someone to exist. It is difficult to see, however, how one might offer a plausible justification of the view that one can be comparatively harmed by coming to exist but not comparatively benefitted. This is the tricky position that Benatar finds himself in. The Pleasure/Pain Asymmetry carries too much metaphysical baggage to be retained.

The moral problem for Benatar's solution is that the Pleasure/Pain Asymmetry entails that everyone who will experience an instance of pain in her life is worse off than she would have been had she not come to exist. Conjoined with CCAH, the Pleasure/Pain Asymmetry entails that everyone is harmed by being conceived. If the fact that coming to exist harms Pebbles, conjoined with CCAH and the obligation to minimize harm, is sufficient to entail that Wilma acts wrongly by conceiving Pebbles, then the Pleasure/Pain Asymmetry entails that all acts of conception turn out to be seriously morally objectionable. If someone experiences so much as a hangnail during her life, she suffers a harm in a morally relevant sense by coming to exist. This implication seems more implausible than the conclusion of the NIPA that it is supposed to avoid. Benatar's solution fails the Intuitive Cost Criterion for a successful solution and must therefore be rejected.

II.2.6 – Counterfactual Comparative Harm-based Solution 3: Roberts’ Solution

Melinda Roberts offers an argument that, in a limited respect, constitutes a rejection of (P1). She accepts that there is nothing morally problematic about conceiving the worse-off of two children in cases where the probability that the children will be numerically distinct is 1.²⁷ In this respect, Roberts bites the bullet of the Non-identity Problem by accepting a strong person-affecting principle. She thinks the Non-identity Problem only arises in a morally interesting way in cases where there is at least some uncertainty about whether a distinct child will be born if someone does an act A rather than act B. She claims that, when evaluating the moral status of a *potentially* identity-affecting action, we should not look to the actual outcome, but we should instead look at the expected outcomes of an action just prior to its performance. This way we can give a person-affecting solution to what she takes to be the morally interesting Non-identity Problem.²⁸ Given their information, Wilma and Fred have no reason to think that the likelihood that a particular individual will be created if they choose the blind-child option is any higher than if they were to choose the sighted child option. If either action is performed, because existence is so precarious, the probability that Pebbles will exist is, say, (.000001). Since the probability is the same whether the blind-child-creating action is performed or the sighted-child-creating action is performed, she argues that we need only compare the expected outcomes when determining which of Wilma and Fred’s actions are worse *for Pebbles*. When we cancel out the probabilities

²⁷ Roberts (2007). Because Roberts thinks there is nothing morally objectionable about conceiving the worse off of two potential children under such circumstances, her “solution” is not a full solution to what I take to be the central version of the Non-identity Problem. Nonetheless, it is worth considering her limited solution and to see that even this limited solution fails.

²⁸ Roberts (2007).

and compare just the outcomes, it turns out that Wilma and Fred perform the action that is *expected* to be worse for Pebbles, and this is enough to maintain that they act wrongly.

There are two difficulties for Roberts' solution. First, I am perfectly happy to revise the original case so that (i) Fred and Wilma will conceive Pebbles *only if* (all probabilities are 1) they choose the blind-child-creating action and (ii) Fred and Wilma *know* that Pebbles will only exist if the blind-child option is chosen. This circumvents the probability talk while still generating the intuitive problem.²⁹ Because Roberts' solution cannot solve *this* version of the problem it fails the Sufficient Generality criterion for a successful solution. Second, I challenge Roberts' claim that in cases where two mutually exclusive actions yield equal probabilities that one and the same person will exist that we should just compare the outcomes of the two actions when determining their expected harmfulness. It seems that if we want to determine whether an action is likely to make someone worse-off than she otherwise would have been in creation cases, we need to multiply the probability that a person will exist given one action by the probability that *the same* person will exist given the alternative action. This is how you determine the *likelihood* that some action or another will make one and the same person worse-off than she otherwise would have been. When we multiply the probability (.000001% chance) that Pebbles will come to exist as a sighted person by the "loss" in well-being that she suffers by being born blind, however, the likelihood of comparative harm is so small that we do not get so much as a prima facie presumption against Wilma in Fred's action in person-affecting terms. The odds of hitting someone with my car during any given trip to the grocery store are much higher

²⁹ Roberts thinks that cases in which (i) and (ii) obtain there simply is no moral problem. An agent does nothing morally bad by choosing to create the worse off of two possible people. Roberts does not do much in the way of debunking our intuition about what I would call 'genuine' non-identity cases. She simply denies that they generate genuine moral problems.

than the odds of Pebbles coming to exist as a sighted child. Surely I do not need to factor this fact into my decisions about grocery shopping! Roberts ignores this morally relevant consideration by ignoring the probabilities rather than factoring them into the moral calculations. When we factor in the probabilities we do not get so much as a moral presumption against Wilma conceiving Pebbles rather than the sighted child.

II.3 – Rejecting (P2): The Non-comparative Account, De dicto Account, and Event-based Account of Harm

Comparative harm-based solutions to the Non-identity Problem do not have much promise. They fail to satisfy either the Sufficient Generality or the Intuitive Cost Criterion for a successful solution. I now turn to the subcategory of harm-based solutions to the Non-identity Problem I have called “non-comparative,” “*de dicto*,” and “event-based-based solutions.” These solutions share an important feature in common. They depend on accounts of harm that are rivals to CCAH. The non-comparative, event-based, and *de dicto* accounts of harm are motivated by several significant problems for CCAH. This section is thus divided into two major subsections. In section II.3.1 I discuss several recent objections to the comparative account of harm, and I respond to those objections. I argue that the comparative account of harm, properly formulated, can answer the most problematic objections raised against it. This significantly undermines the motivation for turning to one of the rival accounts of harm. In section II.3.2 I canvas CCAH’s rivals: the non-comparative, *de dicto*, and event-based accounts of harm. I also discuss their corresponding solutions to the Non-identity Problem. I argue that, despite their initial appeal, non-comparative, *de dicto*, and event-based solutions are unacceptable because they fail to

satisfy the Intuitive Cost and Sufficient Generality Criteria for a successful solution. The accounts of harm on which they are based face insurmountable problems, and the solutions fail to solve all versions of the Non-identity Problem. My conclusion of section II.3 is that there is no satisfactory harm-based solution to the Non-identity Problem currently on offer.

II.3.1 – Objections to the Counterfactual Comparison Account of Harm

Though CCAH is a popular account of harm, a number of objections to CCAH have been raised against it in the literature. If CCAH is shown to be false, this may prove to be a source of motivation for turning to the alternative accounts of harm that I will discuss in II.3.2. In the first half of section II.3 I address the common objections to CCAH in order to undermine this potential motivation for turning to alternative accounts of harm. According to one objection, CCAH cannot distinguish morally between classic cases of harm and preventative harms. Another objection claims that CCAH cannot account for harms that fail to make someone worse off than she otherwise would have been. Others have objected to CCAH on the grounds that it “excessively multiplies” harms; it entails that someone suffers many harms in cases where, intuitively, she only suffers one.³⁰ Yet another objection contends that CCAH cannot account for the harm of death. According to a final objection, CCAH cannot explain the moral difference between harming and failing to benefit. I respond to these objections in several ways: some objections can be met by making a distinction between harm-in-a-respect and all-things-considered harm; others can be met by appealing to the Contextualist component of CCAH; the harm of death can be accommodated by turning to a “whole lives” version of CCAH according

³⁰ Hanser (2008: 433-444).

to which someone suffers a harm just in case her life as a whole goes worse than it otherwise would have gone; finally I argue that the objections from preventative harm and failures to benefit demand too much of CCAH.

II.3.1.A – Preventative Harm

Some object to CCAH on the grounds that it cannot morally distinguish between a particular type of harm, commonly referred to as ‘preventative harm’, and other types of harms. Someone is preventatively harmed if she is prevented from receiving a benefit. If my neighbor intercepts a check for one million dollars, made out to me, just before it reaches my mailbox I suffer a preventative harm. CCAH can categorize preventative harms as harms. As long as I am worse off in a way than I otherwise would have been by failing to receive the check, I am harmed according to CCAH. The objection from preventative harm, then, is that CCAH does not place preventative harm in a special moral category of harm. This is a problem, some argue, because there seems to be a moral difference between harming by preventing an improvement in one’s welfare and harming by causing a reduction in one’s welfare. Most think it would be morally impermissible, for example, to remove someone’s organs to save two others without first acquiring permission. On the other hand, most think it would be permissible to prevent the same person from receiving an organ transplant without first acquiring permission when by doing so

one could save two others in need of the organs.³¹ CCAH seems incapable of accounting for the moral difference between cases such as these and this, some claim, is a problem for CCAH.³²

The preventative harm objection seems to demand too much of CCAH. A satisfactory account of harm will distinguish harms from non-harms in an intuitively satisfying way. The objection from preventative harm does not give us any reason to think that CCAH fails in this respect. CCAH can, after all, classify preventative harms as harms. The objection from preventative harm demands something further of an account of harm. The objection seems to implicitly demand that an account of harm be able to explain what makes one *action* morally worse than another when the two are on a par with regard to their (counterfactual comparative) harmfulness. This, I say, asks too much. The defender of CCAH can agree that preventative harm is morally less problematic than non-preventative harm but deny that this moral difference must be a function of an act's *harmfulness*. This is because the *way* that a particular outcome is brought about can matter morally. To see that the way an outcome is brought about might matter morally, consider a pair of cases from the literature on the doing/allowing distinction offered by Philippa Foot.³³

Rescue I: We are hurrying in our jeep to save some people—let there be five of them—who are imminently threatened by the ocean tide. We have not a moment to spare, so

³¹ I stipulate that none of the individuals have a special claim on the organs and that the two cases are on par with regard to how much worse off the victims are than they would have been had the harmful act not been performed.

³² Hanser (2008: 428-429).

³³ Foot (1994: 280-289).

when we hear of a single person who also needs rescuing from some other disaster we say regretfully that we cannot rescue him, but must leave him to die.

Rescue II: We are again hurrying to the place where the tide is coming in in order to rescue the party of people, but this time it is relevant that the road is narrow and rocky. In this version the lone individual is trapped on the path. If we are to rescue the five we would have to drive over him. If we stop he will be alright...he is in no danger unless from us.

Saving the five, thereby ensuring the death of the one, in *Rescue I* seems not only permissible but obligatory, whereas saving the five ensuring the death of the one in *Rescue II* seems impermissible. The outcomes are the same so it must be the *way* they are brought about that makes the moral difference between the cases. For similar reasons, the *way* an agent acts when she preventatively harms might be easier to justify than the *way* she acts when she non-preventatively harms. Perhaps an act of non-preventative harm violates the victim's autonomy in a way that an act of preventative harm does not. Diverting an organ from one patient to another before it reaches its intended recipient (preventative harm), for instance, certainly seems less invasive than opening one patient up, taking out her organ, then giving it to someone else (non-preventative harm). Alternatively, individuals might have special claims on their own body parts but not on body parts that would have been theirs someday had no one intervened. If we couple these non-harm-related ways that an act of preventative harm might differ from non-preventative

harm with the plausible view that one has a greater claim against greater violations of autonomy, then one can preserve a moral distinction between preventative and non-preventative harms without maintaining that one kind of harm is worse *qua harm* than the other. This will do to answer the objection from preventative harm.

II.3.1.B – Harms That Fail to Make the Victim Worse-off

Elizabeth Harman contends that a person can be harmed in a morally relevant way even though (i) she does not regret being treated in the harmful way and (ii) she is not worse-off than she otherwise would have been as a result. She offers a pair of cases meant to support her contention. If the cases are indeed instances of morally relevant harm in which (i) and (ii) obtain, then CCAH is false, because it denies that someone can be harmed in a morally relevant sense when she is not made worse off than she otherwise would have been. Let's turn to the examples that Harman offers.

Rape: A woman is raped, becomes pregnant, and ends up raising the child. The woman is remarkably able to separate the trauma of the rape from her attitude to the child, and they have a normal and healthy parent-child relationship. The woman's life is better, due to the value to her of the relationship with her child, than it would have been if she had not been raped, even taking into account the trauma of the rape. This woman loves her child. She does not wish that she had not been raped, because if she had not been raped, then her child would not exist.

Nazi Prisoner: A man was imprisoned in a Nazi concentration camp, where he suffered many harms. But his experience in the camp enriched his character and deepened his understanding of life, such that overall his life was better than it would have been had he not been imprisoned in the camp. He does not wish that the Nazis had not imprisoned him, because he so values what he gained from this experience.³⁴

Harman takes these cases to show that CCAH cannot properly account for all cases of morally relevant harming. The victims in *Rape* and *Nazi Prisoner*, she claims, are clearly impermissibly harmed despite the fact that they are not made worse off than they otherwise would have been.

I suspect that the intuitive force of Harman's examples comes not from the fact that the victims in the cases have been harmed, but rather from the fact that they have had very intrusive things done to them *without their permission*. They have had their rights violated. To see that considerations of rights might be influencing our judgments about the actions in *Nazi Prisoner* and *Rape*, consider two more cases. Let's call them *Not Rape* and *Boot Camp*:

Not Rape: All of the same details from *Rape* except that the "victim" consents to the sex beforehand hoping that she will conceive a child that will make her life on balance better.

³⁴ Harman (2004)

Boot Camp: All of the same details from *Nazi Prisoner* except that the “victim” consents to the treatment in the hopes that he will end up leading an enriched life as a result.³⁵

Our intuitions about *Not Rape* and *Boot Camp* are very different from *Rape* and *Nazi Prisoner* in two important ways. First, most do not think that the victims in *Not Rape* and *Boot Camp* are wronged. Second, I doubt that most have the intuition that the “victims” in *Not Rape* and *Boot Camp* are *harmed*. In fact, it is fairly clear that they are benefitted. But the only difference between the original pair of cases and this revised pair is that *consent* is obtained in the latter but not in the former. But it is implausible that the acquisition of consent by an agent constitutes the difference between an act being harmful or beneficial. Whether I get permission from someone before kicking him in the shins does not determine the act’s harmfulness, though it might very well determine its permissibility. *Boot Camp* and *Not Rape* indicate that our intuitions about the original pair of cases are mistaken. The actions in *Nazi Prisoner* and *Rape* are not harmful, but they are wrong insofar as they are physically intrusive actions performed without someone’s consent. Our judgment to the contrary is mistaken, and this mistake is due to confusion about an act’s being *harmful* on the one hand and *wrong* on the other.

For the reasons just discussed, I think Harman’s examples are unsuccessful in mounting an attack against CCAH. Some, however, might still be uneasy. Surely, they will say, almost everyone will judge that the victims in *Nazi Prisoner* and *Rape* are harmed. If CCAH cannot at least concede this point, then perhaps the view ought to be rejected on the grounds that it is ill-equipped to accommodate a judgment about harm that is shared by virtually everyone.

³⁵ Harman (2004: 99)

Responding to this worry is a good departure point for making the distinction between *harm in a way* or *prima facie* harm and *all-things-considered* harm. Someone suffers a counterfactual comparative harm in a way just in case she is worse off in a way than she otherwise would have been. Someone suffers a counterfactual comparative harm all-things-considered just in case she is worse off all-things-considered than she otherwise would have been. A couple of examples should suffice to illustrate the distinction.³⁶ Imagine you saw these two headlines in the newspaper:

Headline 1: New study shows surgery is harmful!

Headline 2: Scientists develop harmless surgical technique!

Headline 1 would only be interesting if “harmful” was read with the all-things-considered sense in mind. It would not be interesting news if there was a new study that patients were worse off in a way than they otherwise would have been as a result of surgery. All surgery patients are worse off in a way (they suffer temporary pain, weakness, etc.) as a result of surgery, but the surgery nonetheless (normally) leaves patients better off all things considered than they otherwise would have been. A study showing that surgery was worse for patients in the all-things-considered sense (i.e. worse than no surgery at all), however, would indeed be newsworthy. Headline 2, on the other hand, is only interesting if “harmfulness” is read in the harm-in-a-respect sense. As I already noted, surgery is rarely (with the exception of bizarre mishaps) all-

³⁶ This example is borrowed from Ben Bradley (2012).

things-considered harmful. It would be uninteresting to report that this is the case. However, a surgical technique that is not even harmful in a way (has no unpleasant side-effects whatsoever), would be a rather surprising development.

This distinction between harm-in-a-respect and all-things-considered-harm allows the defender of CCAH to avail herself of a reply to the person who is still tempted to describe the acts in *Nazi Prisoner* and *Rape* as harmful. The acts involved in these cases are harmful in a way, but they are beneficial, all things considered.

Note that there are two ways in which someone might be harmed in a respect while being benefitted all things considered. First, someone might be caused to be worse off in some respect for some interval of time while being caused to be better off in another respect during that same interval of time. Consider, for instance, an overbearing dissertation advisor who requires her advisees to read two hundred pages of contemporary articles per week. During each week she harms her advisees in one respect by causing them to suffer severe sleep deprivation. *During that same week*, however, she benefits her advisees in another respect, by enriching their philosophical acumen. There is no interval during which the advisee is worse off, all things considered, than she otherwise would have been, but she is worse off than she would have been along *some dimension* of her well-being during that interval. Second, someone can be harmed in a respect by an action if she is caused to be worse off than she otherwise would have been, all things considered, during some interval of time, but better off than she otherwise would have been, all things considered, if we consider her life as a whole. Suppose that a teenager's parents force him to save the money that he earns during the summer by working as a lifeguard so that he can later invest in a car. In the interval during which he is prohibited from spending the

money he is worse off, all things considered, than he otherwise would have been during that interval. Had he not saved the money during that interval of time, he would have enjoyed going to the movies, going out with his friends, etc. During a later interval, however, the teenager is better off, all things considered (and to a greater degree than he was earlier worse off all things considered), than he otherwise would have been. He enjoys greater freedom of movement thanks to his newly purchased car than he otherwise would have enjoyed during this later interval.

The victims in *Nazi Prisoner* and *Rape* are most likely harmed in a respect in the second rather than the first sense. The victims are likely made worse off than they otherwise would have been, all things considered, for some interval of time, and then turn out to be better off than they otherwise would have been during a subsequent interval of time (it is hard to imagine concentration camps benefitting a prisoner in any respect *at any moment during her imprisonment*). The defender of CCAH can thus admit that there is a sense in which the victims are harmed in *Nazi Prisoner* and *Rape*. They are worse off than they otherwise would have been *in some respect*. While it is open to the defender of CCAH to admit that there is this distinct kind of harm, harm-in-a-respect, and that the victims in *Nazi Prisoner* and *Rape* suffer this kind of harm, it is unclear that harm-in-a-respect is a contender in the debate regarding morally relevant harm. As was pointed out earlier, insofar as we have the intuition that harm-in-a-respect is itself morally relevant (not simply in virtue of its relationship to all-things-considered-harm), this intuition is easily accounted for by considerations of rights and consent. Once consent of a victim is obtained by the agent, we no longer have the intuition that the fact that the victim will be harmed in a way in order to be all-things-considered-benefitted counts against the agent performing the action.

It is also worth noting that, when choosing between (a) suffering harm-in-a-respect while enjoying an all-things-considered benefit and (b) suffering an all-things-considered harm while enjoying a benefit-in-some-respect, it is always prudentially rational to choose to suffer harm-in-some-respect while enjoying the all-things-considered benefit. The reason for this is straightforward. On any plausible account of prudential rationality the optimally rational choice for an agent to make will be the one that makes him fare the best with regard to his welfare, and on any plausible account of individual welfare it is better *for us* to be better off all things considered and worse off in a way than to be worse off all things considered but better off in a way. On the supposition that the correct account of morally relevant harm will correspond with what matters to us, a version of CCAH that counts all-things-considered harm but not harm-in-a-respect as the basic kind of morally relevant harm is better at tracking what matters to us than a version that does the opposite.

Nonetheless, some will be unable to shake the belief that a satisfying account of harm will entail that the prisoner and the rape victim have suffered a harm by their treatment. They will insist that an account of harm must accommodate harm-in-a-respect. CCAH can be modified to accommodate this persistent belief.

CCA'H: A's suffering a harm is A's being worse off *in a way* given *o*, where *o* is the actual outcome, than A would have been in $\sim o$, where $\sim o$ is the relevant alternative outcome. The world picked out by ' $\sim o$ ' varies depending on conversational context.

I will have CCAH' in mind as a point of departure for addressing the remaining objections to the Counterfactual Comparative Account of Harm. Note, however, that my formulation of the view will further evolve over course of addressing these objections.

II.3.1.C – Excessive Multiplication of Harms

Matthew Hanser raises the following objection to CCAH:

Suppose that A shoots B, causing him to become paralyzed from the waist down. A whole series of causally linked events occur here, among which are A's pulling of the trigger, the gun's going off, the bullet's entering B's body, and B's becoming paralyzed. According to the counterfactual comparison account, B is in a distinct harmed state relative to each of these events: he is worse off than he would have been had A not pulled the trigger; he is worse off than he would have been had the gun not gone off; he is worse off than he would have been had the bullet not entered his body; and he is worse off than he would have been had he not become paralyzed. Each of these events, then, comes to him as a separate harm. But I think this is clearly the wrong way to describe the situation. What we have here are not four separate harms, but a single harm—B's becoming paralyzed—with multiple causal antecedents. The counterfactual comparison account

collapses the distinction between events that cause people to suffer harms and the harms that the people are thereby caused to suffer.³⁷

In some ways, CCAH' seems to describe this case correctly. Ordinarily we would say that B is worse off than he would have been (i) had A not pulled the trigger, (ii) had the gun not gone off, and (iii) had the bullet not entered his body. After all, if any of (i)-(iii) had failed to occur then B would not be paralyzed. It is strange to say that each of these events “comes to him as a separate harm,” but the proponent of CCAH' need not maintain any such thing. It is false that B is worse off than he would have been had A not pulled the trigger. A's pulling the trigger, alone, would not have been sufficient for B's being worse off. In isolation, A's pulling the trigger has no effect on B whatsoever. The gun also had to go off, and the bullet had to enter B's body and strike his spinal cord in just the right way.³⁸ This suggests a more specific account of what it is to suffer a harm on CCAH:

CCAH': A's suffering a harm is A's being worse off in a way given $\{o, e\}$, where o is the actual outcome and e is *some actual event consisting of the smallest set of individually necessary conditions that are jointly sufficient* for A's being worse off than

³⁷ Hanser, p. 443 (2008)

³⁸ Note also that A's pulling the trigger is not sufficient for the bullet entering B's body or for it striking B's spinal cord. This avoids the problem of arriving at the conclusion that A's pulling the trigger *is* sufficient for making B worse off by the transitivity of sufficiency. The idea is that if X is sufficient for Y, and Y is sufficient for Z then X is sufficient for Z. By this reasoning, if A's pulling the trigger were sufficient for the bullet entering B's body and the bullet entering B's body were sufficient for the bullet striking B's spinal cord, etc., then A's pulling the trigger would be sufficient for the bullet striking B's spinal cord.

A would have been given $\{\sim o, \sim e\}$ where $\sim o$ is the relevant alternative outcome and $\sim e$ is the relevant alternative event. The world picked out by ‘ $\sim o$ ’ varies depending on conversational context.

On this view someone suffers a harm in relation to some set of conditions that are jointly sufficient for the person being worse off than he would have been had they not occurred. It is reasonable to wonder at this point if we can still felicitously say that A harms B by shooting him. After all, his pulling the trigger is not sufficient, by itself, for B being paralyzed. This is not a problem for CCAH". CCAH" gives conditions that must obtain for someone to *suffer a harm*, but it does not say anything about the relation that A must stand in to B's suffering a harm for it to be true that A *harms* B. I propose the following amendment to CCAH's account of when an actor (who is responsible for an action) harms a victim:

Necessity: A harms B just in case A or one of A's acts is a member of a set of individually necessary conditions that are jointly sufficient for B's suffering a harm.

Necessity allows us to avoid the implication that B suffers multiple harms when he becomes paralyzed while retaining the intuition that A harms B. We do not want to lose sight of the central intuition that A is in large part responsible for B's suffering a harm. After all, if A had not pulled the trigger B would not be paralyzed.

Necessity seems to entail that certain seemingly bizarre claims turn out to be true. For instance, the sentence “the gravitational constant harmed B” is true as long as the gravitational constant is a member of a set of individually necessary conditions that are jointly sufficient for B’s being worse off than he otherwise would have been. But the specific values of the laws of nature are generally part of the set of individually necessary conditions that are jointly sufficient for B’s being worse off than he would have been had they been different. According to Necessity then, the sentence “the gravitational constant harmed B” is almost always true in cases where someone suffers a harm. Perhaps a view that entails that the gravitational constant can harm someone is absurd. Here the Contextualist component of the Counterfactual Comparative Account of Harm plays a role. In most contexts, when we are attributing responsibility for a harm to someone or something, conversational salience rules out causes that would have failed to obtain only in very distant possible worlds. The more appropriate choice for primary harm attribution is typically (i.e. in most conversational contexts) the cause that would have failed to obtain in a relatively close possible world. In the case of A shooting B, the world in which A fails to pull the trigger is a closer world than the one in which A pulls the trigger but the gun fails to fire or the world in which the bullet fails to enter B’s body. It is certainly closer than the world in which the gravitational constant is different. If we incorporate Necessity into CCAH” then it is, strictly speaking, true that the gravitational constant harms B, but the view’s Contextualist component can explain why we do not ordinarily speak that way. Hence, CCAH” and Necessity together explain why B only *suffers* one harm when he is shot while maintaining both that there are multiple causes of B’s suffering a harm and that A is an appropriate target for the attribution of primary responsibility for the harm that B suffers.

II.3.1.D – Death

Some have argued that CCAH, even the refined CCAH³⁹, runs into problems addressing the harm of death. Consider one rough and ready statement of CCAH offered by Judith Thomson:

CSA: A's suffering a harm is A's being in a state *s* such that A is worse off in a way for being in *s* than he would have been if he hadn't been in *s*.⁴⁰

This formulation of CCAH is state-based. It entails that to determine whether someone has suffered a harm we need to compare how she fares in two states. One of these states is her actual

³⁹ Recall that I am using the term 'comparative account' as shorthand for 'counterfactual comparative account'. Elizabeth Harman (2004) and Seana Shiffrin (1999) defend non-comparative accounts of harm. I will not address the non-comparative version of the state-based account. I will also not be addressing *temporal* versions of the comparative account of harm.

⁴⁰ Thomson (2011, p. 446). Thomson actually rejects CSA in favor of a more complex version because of the following example: 'suppose villain C threw acid in A's eyes, to cause him to be blind. The acid began to affect A's eyes, but before it could complete the process of blinding A, bystander B intervened—he used a neutralizer on A's eyes, thereby causing A to be in a state short of blindness, namely 'has dim vision''. Thomson thinks that this example establishes that 'anyone who thinks that A is suffering a harm in virtue of being in 'has dim vision' must give up the assumption that harming a person is causing him to suffer harm, for B caused A to be in 'has dim vision' but did not harm him'. Ibid, p. 448

I think that CSA can satisfactorily explain why C harms A but B does not. When C throws acid into A's eyes he causes him to be in 'has dim vision'. A is worse off for being in 'has dim vision' than not when C causes him to be in this state, since the alternative, at the moment C throws the acid into his eyes, is being in 'perfectly sighted'. At the moment B comes upon the scene, however, there are only two alternatives left for A, neither of which is 'perfectly sighted'. A will either be in 'has dim vision' or be in 'blind'. 'Perfectly sighted' is not an available alternative given that C has already thrown the acid into A's eyes. B causes A to be in 'has dim vision' and B is in no way worse off for being in this state given that the only alternative, *at the moment* that B uses the neutralizer, is A's being in 'blind'. Thus, when C causes A to be in 'has dim vision' he harms A, but when B causes A to be in 'has dim vision' he does not harm him. One need not modify CSA to get this result.

state, and the other is the relevant possible state that she would have been in had the action/event not occurred. If I hit my head on a cabinet door, thereby suffering a painful headache, CSA entails that I suffer a harm because I am worse off in a way for being in the state of having a painful headache than I otherwise would have been. For someone to suffer a harm in dying, according to CSA, death must be a state in which she is *worse off* in a way than she would have been had she remained alive. At this point the defender of CCAH runs into trouble.

Normally, when someone is worse off in a way in one state than in another state this is because her *welfare level* is lower in one state than in the other. I will suppose that a person's welfare level is the point on some absolute scale of welfare that she occupies at a moment or for some interval of time. Nothing that follows hinges on this supposition, however, so long as the reader agrees that individuals have welfare levels while they are alive and that these levels can be affected for better or worse. I will not analyze the notion of a welfare level any further. The following plausible principle seems to rule out the possibility that the dead have welfare levels:

DW: Someone has welfare properties at a time t only if she is located at t .⁴¹

⁴¹ I borrow this statement of the principle from Ben Bradley's excellent discussion of this issue (2011: 79-92). Bradley offers a compelling argument for the view that the dead *do* have welfare levels. On his view the welfare level of the dead is 0. While I am sympathetic with Bradley's reply to the Epicurean, my reply happily avoids a debate about the ontological status of the welfare properties of the dead. See Holtug (2001) and Bykvist (2007) for defenses of principles that are similar to DW.

I will assume DW because the harm of death poses a problem for CCAH only if DW is true. The truth of DW spells trouble for the defender of the counterfactual comparative account who endorses CSA. If DW is true, then, because the dead are not located, they do not have welfare properties.⁴² CSA states that someone suffers a harm just in case she is in a state that is worse for her in a way than some alternative. But it is plausible that how good a state is for someone depends on her welfare level in that state. If the dead do not have welfare properties, they do not have welfare *levels* that can be compared with their welfare level in a counterfactual state in which they remain alive. For this reason it seems that death does *not* result in a person being in a state that is worse for her in a way than some alternative. Therefore, defenders of the counterfactual comparative accounts who accept CSA seem committed to the counterintuitive Epicurean conclusion that one does not suffer a harm in dying. In Epicurus' words:

...death, the most terrifying of ills, is nothing to us, since so long as we exist, death is not with us; but when death comes, then we do not exist. It does not then concern either the living or the dead, since for the former it is not, and the latter are no more.⁴³

⁴³ Epicurus, 'Letter to Menoeceus,' trans. C. Bailey, *The Stoic and Epicurean Philosophers*, edited and with an introduction by Whitney J. Oates (New York, N.Y.: The Modern Library, 1940: 30-31. Lucretius presents essentially the same argument. See *On the Nature of Things*, trans. H. A. J. Munro, and *The Stoic and Epicurean Philosophers*, p. 131.

Hanser describes the problem for the comparative account as follows:

If the dead fare neither well nor badly, however, then the counterfactual comparison view, as presently formulated, cannot adequately account for the possibility of death's coming as a harm. For it is not generally true that a person who dies would have fared better over some interval of time had the event of his death not occurred. Had his death not occurred his life would have lasted longer, and we may suppose that he would have fared well over this extra interval of time. But this would not have been an interval over which he would have fared better than he actually fared. In actuality he was dead during this interval, and so had no level of welfare level; and in the absence of an actual level of welfare, no comparative judgment can be made.⁴⁴

If we cannot make comparisons between a person's welfare level when she is dead and what her welfare level would have been had she remained alive, then the defender of the counterfactual comparative account who endorses CSA appears to be left unable to explain why death comes to us as a harm. But since death clearly *is* harmful in most cases, we should prefer a

⁴⁴ Hanser, p. 437 (2008). Steven Luper (2007): 239-51, Harry Silverstein (2000): 116-34, John Bigelow, John Campbell and Robert Pargetter, and David Hershenov (2007) are among those who argue for the no-welfare-level view described by Hanser in this passage.

view of harm (perhaps the event-based account) that can accommodate this obvious fact.⁴⁵ The defender of the counterfactual comparative account is in trouble.

II.3.1.D.a– Thomson’s Reply

Judith Thomson offers two rebuttals on behalf of defenders of the counterfactual comparative account to Hanser’s objection from death. First, she contends that:

a [defender of the counterfactual comparative account] might insist that the harming one does in killing the young and thriving is a unique kind of harming for which the usual requirements for harming do not hold. All other harmings (the [defender of the counterfactual comparative account] might go on) cause the victim’s life to be...comparatively bad in a way. But when you harm a young and thriving person by killing him, you don’t cause your victim’s life to be bad in some way, you simply end it, which is a very different affair.⁴⁶

⁴⁵ A plausible account of harm should entail that death is not *always* harmful. It should be able to account for cases in which death would be preferable than continuing to be alive. Perhaps someone who lives in extreme joyless agony is not harmed by death. I discuss this issue in greater detail in section III.

⁴⁶ Thomson p. 455 (2011). Parentheses added.

It is doubtful that many who are sympathetic to Hanser's worry will find this response satisfying. The reply is problematically *ad hoc*. As Hanser complains, 'For Thomson, death is just an unexplained exception to an otherwise correct account of harm'.⁴⁷ But an account of harm that cannot explain why death is harmful is, other things being equal, worse than a view that can.

In giving her second reply to Hanser's criticism, Thomson appeals to evidence from ordinary language to support the view that one can be made worse off by being caused to die. For instance, we do not hesitate to say that it is better to be alive and disabled than dead.⁴⁸ This reply is not *ad hoc*, but it isn't very helpful either. Hanser's challenge to the defender of the counterfactual comparative account is to *explain* why these kinds of claims are true. Appealing to the fact that we make such claims will not do as an explanation of why they are true, especially if one's preferred account of harm seems to entail that such claims are false. In the next section I show that the defender of the counterfactual comparative account can accept DW and still provide an account of the harmfulness of death. Moreover, the defender of the counterfactual comparative account can answer a stronger version of the Epicurean challenge than the one posed by Hanser. The foundation of this answer comes from the commonsense observation that, typically, *lives* are better when they last longer.

⁴⁷ Hanser p. 463 (2011).

⁴⁸ Thomson P. 455-456 (2011). Feinberg (1986) offers a similar argument.

II.3.1.D.b – The Value of a Life

If the defender of the counterfactual comparative account grants DW, then, in order to begin to meet Hanser's Epicurean challenge, she must begin by appealing to Fred Feldman's explanation of the badness of death. Feldman claims that death is bad for the person who dies because it adversely affects the *total value* of her life. The badness of death for someone is determined in the same way that the value for someone of any event or outcome is determined. An outcome or event is bad for someone to the extent that it makes his *life as a whole* go worse than *it* otherwise would have gone.⁴⁹ In this section I amend and refine Feldman's explanation in order to meet the Epicurean challenge, as well as several related challenges, for the defender of the counterfactual comparative account.

On Feldman's view, to determine whether the outcome in which someone dies is worse for him than the relevant one in which he continues to live we need to compare the value of the person's whole life given each outcome. We do not need to compare his welfare *level* after he is dead with what his welfare level would have been were he to have continued to live at the times after his actual death. The truth of DW thus poses no problem for Feldman's view. Death is bad for us because it makes it the case that our lives as a whole involving fewer of the things that make a life go well. It does this by cutting life short. The defender of the counterfactual comparative account should reject CSA and instead use Feldman's insight about the badness of death to revise CCAH in the following way:

⁴⁹ Originally defended by Feldman (1991). He develops the account further in his (1992).

CCA^{'''}: A suffers a harm just in case the total value of A's life l given outcome o , where o is the actual outcome, is lower than the total value of A's life $\sim l$ given $\sim o$, where $\sim o$ is the relevant alternative outcome.⁵⁰

For the sake of simplicity, I have left out several details from CCA^{'''} that were included in CCA['] and CCA^{''}. For instance, I have left out the 'in a way' clause, included to accommodate instances of harm-in-a-respect, as well as the 'smallest set of necessary conditions' clause, included to address the 'multiplication of harms' concerned raised by Hanser. These clauses are intended to be implicit in CCA^{'''}. To determine whether someone suffers a harm in dying, CCA^{'''} asks us to compare the value of her whole life given her actual death with what the value of her whole life would have been had she not died when she did. CCA^{'''} is a plausible revision of the counterfactual comparative account, but it supposes the following principle:

⁵⁰ Parsons (2003) describes a similar account of harm, though he does not discuss its implications for the harm of death. It is worth noting that stating the view as I have might come across as a bit misleading. By saying that S 'is' worse off, it sounds like there is some particular point in time at which S is worse off. This makes it sound like S is worse off in the actual world *after he is dead* than he would have been in the other world. But this need not be the case. When I talk of 'S being worse off' at o (the outcome in which he has died) I am speaking of the person who existed prior to his death and saying, of *him*, that his life was worse than it otherwise would have been. It is true at o , when S dies at time t , that S's life (where S is the person who existed up until the moment of his death) has gone worse than it would have gone at $\sim o$, the closest possible outcome at which S does not die. S's life is worse at o than $\sim o$ because it has contained less of what is intrinsically good at o than it would have at $\sim o$.

VWL: How well things go for someone is wholly determined by the intrinsic goods (bads) that the person enjoys (suffers) during her life or the relations that obtain between those goods during her life.⁵¹

VWL is just a way of stating an implicit commitment of CCAH. CCAH is committed to the view that the only relevant facts in determining whether someone suffers a harm are facts about that person's welfare that obtain during her *life*. Once a person is dead, we have all of the facts that we need to determine whether death came to her as a harm, because we have all of the facts about how things actually went for her. The game is over, and we can tally up the points. This commitment, captured by VWL, seems unproblematic with one possible exception. VWL might seem to be rejected by those who hold that someone can be harmed after her death. There are ways for accounting for posthumous harm that are consistent with VWL. I will only note one of these ways. For instance, she might hold that events that happen after someone's death can affect how her life is going *now*.⁵² Nonetheless, I can help myself to VWL, because it is entailed by

⁵¹ David Velleman (1993) and Doug Portmore (2007) have argued that 'global' features of a life can contribute intrinsic value to the life. For instance, later successes might 'redeem' past sacrifices insofar as they can retroactively make past sacrifices less bad for the person who endured them. On this view the value of someone's life might be determined, not only by the intrinsic goods (bads) that a person enjoys (suffers) during her life, but also by the *relationship* between those goods (bads). This possibility does not pose a problem for VWL or the arguments I present here as long as we think of the good of 'redemption' as a good that obtains *during* someone's life rather than after it.

⁵² One might believe in posthumous harm because she believes that you can suffer intrinsic bads or enjoy intrinsic goods after you die if, say, some of the desires that you have before you die are frustrated or satisfied after you die. It is not obvious that we need to locate these goods and bads at times after a person's death in order to accommodate posthumous harm, however. George Pitcher, for instance, defends a strategy for accounting for posthumous harm

DW, a principle that I am assuming. If a person does not have welfare properties, she cannot enjoy (suffer) any intrinsic goods (bads). This entails that once her life is over (once she is dead) her welfare properties cannot change; things cannot go better or worse for her *after* her death. Therefore, DW entails VWL. CCAH''' allows the defender of the counterfactual comparative account to explain why, typically, death comes to us as a harm: death comes to us as a harm because it typically means that the value of her life as a whole was lower than the value of her life as a whole would have been had she remained alive a while longer. By moving to a 'whole-life' comparison the defender of the counterfactual comparative account avoids comparing the welfare level that the person has when dead with the welfare level the person would have had were she to have remained alive.

II.3.1.D.c – The Timing Challenge

Hanser briefly considers and then dismisses Feldman's solution to the Epicurean challenge to the defender of the counterfactual comparative account. He claims that, by moving from CSA to CCAH''', the defender of the counterfactual comparative account gives up the ability to locate harms and benefits 'temporally'. As a result, the defender of the counterfactual comparative account faces new implausible implications. Here is his complaint:

that avoids the implication that someone can enjoy (suffer) intrinsic goods (bads) after her death (1993: 166-8). On Pitcher's view, any events that change a person's welfare properties that occur after her death would have to cause changes to her welfare properties at times *before* she was dead. This might require backwards causation, a metaphysical commitment that most are loath to accept. Since these changes happen *during* her life, accounting for posthumous harm this way is consistent with VWL. I do not need to address this debate, because DW precludes the possibility that one can suffer intrinsic goods or bads after her death, and I am assuming DW.

Prior to the revision, the counterfactual comparison view located harms and benefits temporally: someone suffered harm if and only if he was in a harmed state for a certain interval of time, namely the interval during which he was faring worse than he would otherwise have fared. Likewise for benefits. On the revised account, however, spending an interval of time in a comparatively bad state is neither necessary nor sufficient for suffering harm. Instead, all harms and benefits are ‘global’: a person suffers harm if and only if his life as a whole would have gone better in some respect had some event not occurred, and he receives a benefit if and only if his life as a whole would have gone worse in some respect had some event not occurred.⁵³

According to CSA, my initial formulation of the counterfactual comparative account, someone suffers a harm just in case she is in a state that is worse for her than the state she otherwise would have been in. Hanser supposes that the ‘temporal location’ of a harm or benefit to someone is the interval of time during which that person is in a state that is worse for her than the state she otherwise would have been in during that interval. Since it is plausible that all states of a person are located in time, CSA entails that someone suffers a harm just in case there is some period of time during which she is in a state in which she is worse off in a way than she otherwise would have been in. In this sense, CSA locates all harms and benefits temporally. CCAH''' gives up this feature by comparing whole lives rather than states of a person in

⁵³ Hanser, (2008: 438-39)

determining whether someone suffers a harm. There are two potential problems for the defender of the counterfactual comparative account if she gives up the temporal location of harms and benefits.

The first problem is that the defender of the counterfactual comparative account might thereby fail to answer other pressing questions raised by the Epicurean challenge. CCAH''' answers the questions: *why is death harmful to the person who dies; and when does the person suffer a harm?* But answering these questions is only part of the Epicurean challenge. The main thrust of the Epicurean worry for the defender of the counterfactual comparative account is the question: *at what times is death harmful for the person who dies?*⁵⁴ If I clumsily hit my head on a cabinet door, I suffer the harm at the moment I hit my head, but that event is *harmful for me* at later times. It is harmful during the following week while I experience a headache when I otherwise would not have been experiencing a headache, and it ceases to be harmful once the headache subsides. Similarly, if I fall into a coma for a month after hitting my head on the cabinet door, the event of hitting my head is harmful *during* the following month when I would have otherwise been conscious and happy. One worry raised by Hanser's remarks is that the move to CCAH''' prohibits the defender of the counterfactual comparative account from locating an interval of time *during which* death is harmful for the person who dies. Hanser claims that by moving to CCAH''' the defender of the counterfactual comparative account must say that death is "timelessly worse. On this analysis," he says, "the harm of death cannot be located temporally at

⁵⁴ Bradley (2010, p. 86). Feit (2002), Lamont (1998), and Silverstein (2000) seem to agree that this is the relevant question posed by Epicurus about the badness of death. Hanser does not explicitly raise this point, but the above passage indicates that he would endorse this version of the Epicurean challenge.

all.”⁵⁵ This worry would be justified if, as Hanser supposes, the question *at what times is death harmful for the person who dies* was tantamount to the question *during what interval of time is someone in a state in which she is worse off in a way than she otherwise would have been*. In this case, CCAH''' could not fully answer the Epicurean challenge. Unlike the interval of time that follows the event of hitting my head on the cabinet door or falling into a coma, the week that follows the event of my death is a week during which I do not have welfare properties. Because I do not have welfare properties, I cannot be in a state during that interval in which I am *worse off* in a way than I otherwise would have been.

Fortunately for the defender of the counterfactual comparative account, conceding that the dead are not in a state in which they are worse off in a way than they otherwise would have been during some interval of time is *not* tantamount to conceding that death is not harmful for someone during intervals of time. CCAH''' can, in its own way, locate intervals of time during which harmful events, including death, are harmful for the person who suffers them.⁵⁶

⁵⁵ Hanser (2008, p. 448)

⁵⁶ Here I may be departing from Feldman. Feldman holds the view that death is bad for the person who dies *at all times*, including the times after the person's death and at the times before the person exists (1991) (1992). Death is bad at all times because the proposition that death is bad for the person who dies *is true* at all times. I am inclined to agree with Neil Feit that Feldman's answer does not address the spirit of the timing question. Feit claims that Feldman answers the question *when is it true that S's death is bad for him*, but he does not answer the question *at which times t is it true that his death is bad for him at t?* Feit (2002, p. 372) I take the Epicurean timing challenge to be asking the latter, not the former question.

II.3.1.D.d – The Times of Death’s Harmfulness

Recall that VWL entails that, once someone’s life is over (i.e. once she is dead), we have all of the facts about how well things went or will go for her. These facts are not going to change after her death.⁵⁷ CCAH''' can thereby determine whether death comes to us as a harm by comparing the total value of someone’s actual life with the total value that her whole life would have contained had she continued to live a while longer. But CCAH''' also permits comparisons between the total value of someone’s actual life with what the value of her life would have been *during intervals of time* after her death. We can ask had someone not died when she did die, how things would have gone during the week, the month, the year, or the rest of the life she would otherwise have had. These comparisons between the values of certain intervals of time during someone’s life are like comparisons between the values of whole lives; both kinds of comparisons can be made without comparing someone’s welfare level when she is dead with what her welfare level would have been had she remained alive.

Suppose that Socrates’ actual life contained 100 total units of value. Now suppose that had Socrates not died of hemlock poisoning he would have escaped prison unharmed and spent five years happily in exile. During these years Socrates would have enjoyed 50 more units of welfare before dying a quick and peaceful death. Since the value of Socrates’ life during this five-year period would have been 50 units, the total value of Socrates *whole* life at the end of the five-year period would have been 150 units of value. Here we have the makings of an answer to the question *at what times is death harmful for the person who dies* that is consistent with

⁵⁷ I am still supposing that genuine posthumous harms are impossible. I claimed earlier that this supposition is unproblematic, because it is entailed by DW, a principle that Hanser assumes.

accepting CCAH''' and DW. At the end of the five-year interval, the total value of Socrates' life had he not died of hemlock poisoning (150 units of value) would have been greater than the total value of his actual life (100 units of value). There is no apparent reason to deny that *this fact*, rather than facts about Socrates' welfare properties at intervals of time after his death, grounds the truth of the claim that Socrates' death by hemlock is harmful for him *during* the five-year interval. This suggests an even more fine-grained explanation of the times of death's harmfulness. Suppose that the 50 extra units of welfare that Socrates would have enjoyed would have been distributed evenly across his five extra years of life. He would have enjoyed 10 of these extra units per year. Suppose we want to know how harmful Socrates' death was for him during the first year after his death. We compare the total value of his actual life (100 units of welfare) with what the total value of his life as it would have been at the end of that first year after his actual death during which he would have otherwise remained alive (110 units of welfare). The defender of the counterfactual comparative account can provide similar answers to the timing question for all intervals of time after Socrates' death without denying DW.⁵⁸

By accepting CCAH''', the defender of the counterfactual comparative account can explain why death is harmful, and she can locate the times of death's harmfulness at times after the person dies. These are the intervals of time at the end of which someone's life as a whole would have gone better than it actually did. Hanser raises a second problem for CCAH''' that is

⁵⁸ Note that my defender of the counterfactual comparative account explanation of the times of death's harmfulness is subsequentist. It locates the times of death's harmfulness at times *after* someone's death. On my explanation, there is no time during someone's actual life during which death is harmful for that person. I do not think this is a problem, for I can maintain that the proposition that someone's death is harmful for her at times after her death is *true* at times during her actual life. This proposition is true at all times, so it is true at all times during her life. Here I think Feldman's explanation of the times of death's badness, noted in footnote 27, gets things right.

related to its supposed failure to temporally locate harms. He claims that CCAH''' cannot explain the harmfulness of events in cases like the following:

Six million dollar man: A man's legs are shattered in an accident. Soon after the accident, he is given 'bionic' replacement legs enabling him to run faster and jump higher than he ever could before.

Hanser claims that, 'although in the long run, so to speak, the event in question (the shattering of his legs) came to him as a benefit, I think we should grant that it at first came to him as a harm'.⁵⁹ CCAH''' seems unable to capture this judgment because the view entails that someone suffers a harm just in case the total value of her whole life is lower than the total value her whole life would have been. The total value of the six million dollar man's whole life is greater as a result of the accident than it would have been, so CCAH''' entails that the accident does not come to him as a harm. In light of this problematic case and my discussion of the times of death's harmfulness above, I suggest the following time-sensitive revision of CCAH''':

CCAH'''+: A suffers a harm just in case, for some interval of time i , the contribution of i to the total value of A's life l given outcome o , where o is the actual outcome, is lower

⁵⁹ Hanser (2008, p. 424).

than the contribution of i to the total value of A's life $\sim l$ given $\sim o$, where $\sim o$ is the relevant alternative outcome.

CCAH'''+ says that someone suffers a harm just in case there is some interval of time during which the total value of her actual life is lower at the end of that interval than the total value of her life otherwise would have been at the end of that same interval. Because the value of the six million dollar man's life is lower than it otherwise would have been at the end of the interval of time after his accident but before he is given bionic legs, CCAH'''+ entails that the accident comes to him as a harm. CCAH'''+ not only yields the result that the six million dollar man suffers a harm by the accident, but it yields the result that death comes to someone as a harm, and it locates the harmfulness for someone of death *at times* after her death. It accomplishes all of this while assuming the truth of DW.⁶⁰ There is another peculiar case, however, that even CCAH'''+ cannot fully accommodate. It is worth mentioning if for no other reason than to further refine the counterfactual comparative account:

The Seven Million Dollar Man: A man's legs are shattered in an accident. Soon after the accident, he is given 'bionic' replacement legs enabling him to run faster and jump higher than he ever could before. During the interval of time before his legs are replaced, the

⁶⁰ Note that this formulation of CCAH can also explain why the rape victim and the Nazi prisoner suffer a harm. Presumably, each of their lives goes worse than it otherwise would have gone *for some interval of time* as a result of the rapist's and the Nazi's treatment.

man experiences excruciating pain, but the pain is accompanied by unparalleled clarity of mind that leads to significant intellectual achievement.

Suppose that intellectual achievement is intrinsically valuable and that pain is intrinsically disvaluable. Suppose also that the value of the seven million dollar man's intellectual achievement outweighs the disvalue of the pain that he experiences during the interval of time after the accident but before the leg replacement. Most might be tempted to say that, even in this case, the man suffers a harm despite the fact that there is no interval of time during which the value of his life is lower than it otherwise would have been during that same interval. This indicates the need for the following further refinement of CCAH'''+:

CCAH'''+: A suffers a harm just in case, for some dimension of welfare d and some interval of time i , the contribution of i to the total d -value of A's life l given outcome o , where o is the actual outcome, is lower than the contribution of i to the total d -value of A's life $\sim l$ given $\sim o$, where $\sim o$ is the relevant alternative outcome.

CCAH'''+ assigns value to intervals of time for someone along particular dimensions of her welfare. It thereby entails that the seven million dollar man suffers a harm during the interval of time after his accident but before his leg replacement, despite the fact that his life contains more *all-things-considered* value at the end of that interval. He suffers a harm because his life

contains less value along one particular dimension of welfare (let us call it ‘experiential welfare’) at the end of the interval of time after his accident but before his leg replacement.⁶¹

I take myself to have shown that the counterfactual comparative account has the tools to explain why and when death comes to us as a harm. I have thereby met the full Epicurean challenge for the view.

II.3.1.E – Asymmetries in Harming and Benefitting (Hanser and Shiffrin)

Seana Shiffrin objects to CCAH on the grounds that it fails to account for the morally relevant differences between (i) actively harming (ii) failing to benefit (iii) failing to prevent harm and (iv) actively removing a benefit. Most have the intuitions that (i) is morally worse than (ii-iv), and (iii) is worse than (ii) and (iv).⁶² CCAH cannot explain why these intuitions are justified. If the victim is just as badly off given any of (i-iv), then each of (i-iv) is an equal harm on CCAH. Shiffrin claims that this is a problem for the view.

I will offer the same reply to this objection that I offered to the objection from preventative harm. As I argued in response to the objection from preventative harms, insisting that the correct account of harm make these moral distinctions asks too much. CCAH need only correctly identify what conditions must obtain for an action to harm an individual. It does not need to account for every morally relevant distinction between *the ways* that harm might occur.

⁶¹ I am not here endorsing a pluralist axiology of welfare. If the 7 million dollar man is to pose any sort of problem for TCSA''+, however, then a pluralist axiology must be presupposed. If, say, experiences of pleasure were to be the only component of a happy life, then it would be impossible to harm someone in the way described in the example. The 7 million dollar man would thus pose no problem for TCSA''+, and TCSA''++ would be a needless revision.

⁶² Shiffrin (1999: 121)

Making such distinctions might involve identifying bad-making properties that some harmful actions possess and that others do not. See section II.3.1A for my reply to the objection from preventative harm.

II.3.2 – Alternative Accounts of Harm and Their Solutions: Non-comparative Accounts, De dicto Accounts, and Event-based Accounts of Harm

The previous section showed that the proponent of CCAH can avail himself of satisfying replies to the most common objections to his view. There is no reason to reject CCAH in favor of its competitors on the basis of these objections. In this section I will raise some problems for the alternative accounts of harm on offer. These are the non-comparative account, the *de dicto* account, and the event-based account of harm. The problems for these accounts give us strong reasons to prefer CCAH.

II.3.2.A – Non-Comparative Accounts

According to non-comparative state-based accounts of harm (NCAH), one way to suffer harm is to be in a state that is bad to be in, even if the state is not worse for you than the one that you otherwise would have been in. More formally, but less informatively, we might state the view as follows:

NCAH: A person suffers harm at a time if he is in a certain kind of non-comparatively bad state at that time.

Seana Shiffrin and Elizabeth Harman have offered similar but distinct versions of NCAH, and I will consider them below.

Shiffrin develops her brand of NCAH primarily as a way to explain the following apparent moral asymmetry:

Harm/Benefit: It is easier to justify (i) harming in order to avoid a greater harm when consent cannot be obtained than it is to justify (ii) harming in order to “bestow pure benefits” when consent cannot be obtained.⁶³

Shiffrin thinks that we can solve the Non-identity Problem by appealing to Harm/Benefit. I will quickly sketch what she takes the morally relevant notions of “harm” and “benefit” to consist in and how these notions are related to a solution to the Non-identity Problem. I will then critically evaluate her solution to the Non-identity Problem.

On Shiffrin’s account:

⁶³ Shiffrin, (1999: 124).

All [harms] have in common that they render agents or a significant or close aspect of their lived experience like that of an endurer as opposed to that of an active agent, genuinely engaged with her circumstances, who selects, or endorses and identifies with, the main components of her life... To be harmed primarily involves the imposition of conditions from which the person undergoing them is reasonably alienated or which are strongly at odds with the conditions she would rationally will.⁶⁴

In short, a state is harmful only if it involves a person being caused to be in a state that, given her values and goals, she (rationally) wishes that she were not in. Shiffrin contrasts her notion of harm with what she calls “pure benefits”. Pure benefits are:

those benefits that are just goods and which are not also removals from or preventions of harm [chasm between one’s will and one’s circumstances]. The central cases of pure benefits involve the enhancement of one’s situation or condition, or the fulfillment of nonessential, but perhaps important, interests. Such enhancement and fulfillment go beyond merely securing the minima that make one’s life more than tolerable and susceptible to active identification. The list of goods mentioned earlier (material enhancement, sensual pleasure, goal-fulfillment, nonessential

⁶⁴ Ibid: 123-124.

knowledge, competitive advantage) seems roughly right. Those items all involve goods, but the absence of them would not create the stark cleavage between one's will and one's experience, life, or circumstances that I suggest characterizes harm.⁶⁵

Shiffrin's account of harm is supposed to solve the Non-identity Problem by entailing that Pebbles suffers a harm in a morally relevant way by being caused to exist without sight. Pebbles has conditions imposed on her, by being caused to exist as a blind person, that she might rationally will had not obtained. If it would be rational for Pebbles to will that she not have been born blind and her parents impose the condition of blindness on her, then, on Shiffrin's view, her parents harm her by imposing such a condition. Because the benefits of existence are only pure benefits, they cannot justify the imposition of the harmful condition.

Shiffrin's distinction between goods, the deprivation of which generates a "stark cleavage between one's will and one's experience", and pure benefits, has some problematic implications that make her solution fail the Intuitive Cost Criterion. If it is necessary to suffer from a "stark cleavage between one's will and one's experience" in order to suffer a harm, then it seems that there are many cases in which, intuitively, a harmful act is performed that would not qualify as a harmful act according to Shiffrin's view. The most obviously problematic cases are those that involve animals or young children. Since neither young children nor animals are capable of rational willing, it would seem that Shiffrin's account excludes them from the set of beings that

⁶⁵ Ibid: 124-125.

are capable of being harmed.⁶⁶ This is a very counterintuitive implication. Another problem for Shiffrin's view is that it seems to entail that individuals who are particularly well-off along some dimension of welfare level or another cannot be harmed by suffering losses along that dimension. If, for example, I intercept my moderately well-off neighbor's paycheck or steal his Lexus I do not generate a "stark cleavage between his will and his experience." Thus, I do not cause him to be in a harmful state on Shiffrin's view even though I cause him to be worse off than he otherwise would have been. This is highly counterintuitive. Similar examples can be concocted for any dimension of welfare that an individual might enjoy in sufficient abundance such that reducing the person's welfare along that dimension would not generate a *stark* cleavage between his will and his experience. Even more curiously, her view seems to entail that by stealing from a wealthy person in order to pay for a blind person's eye surgery, not only do I harm no one, but I actually remove a harm that is already present.

There is one further worry for Shiffrin's view: it appears that Shiffrin's notion of what someone 'rationally wills' smuggles in counterfactual comparisons. It seems that in order for someone to *rationally* will that her circumstances were different, it must be the case that her actual circumstances are less desirable than the circumstances that would have obtained had the actual ones failed to obtain. It would seem that her account of harm does rely on counterfactual comparisons after all. But this raises two problems for Shiffrin's account. The first is that, if her view really does smuggle in counterfactual comparisons between one's actual welfare level and one that she otherwise would have had, then it seems to collapse into CCAH. The second problem is that it seems that it is rational to will that one's circumstances were different than

⁶⁶ Bradley (2009) also points out this trouble for Shiffrin's view.

they actually are only if there is some nomologically possible circumstance that is more desirable than the actual one. In the case of Pebbles, because blindness is a necessary condition for her existence, all of the nomologically possible circumstances in which she is not blind are circumstances in which she does not exist. Thus, assuming that her life is at least not worse than non-existence, it is clearly *irrational* for her to will that she not have been born blind. For Pebbles to will that she not have been born blind is, effectively, for her to wish that some necessary condition for her existence not have obtained.⁶⁷ The same is true in the case of the people who come to exist as a result of the Wealthy Society choosing to implement a policy of Depletion.

Shiffrin might amend her view to avoid these worries. She might claim that what ultimately matters is not the type of good that is caused to be present or absent, but rather whether some good is *gained* or *lost*. On this revised view, it does not matter, when determining if someone has been harmed, whether the good that is gained or lost results in a cleavage between one's rational will and one's experience, but whether a good has been gained or lost period.⁶⁸ This view, which is essentially Matthew Hanser's, suffers from other difficulties that will be discussed later.

In light of these problems for Shiffrin's account, perhaps she is most charitably interpreted as simply endorsing a principle identifying a basic *sufficient* condition for suffering a harm of the sort mentioned at the beginning of the section. Consider again

⁶⁷ I am grateful to Alastair Norcross for helping me to work through these points.

⁶⁸ Of course it is hard to see how this kind of view might allow us to solve the Non-identity Problem as formulated in the case of Wilma and Fred. What good has Pebbles "lost" by existing blindly?

NCAH: A person suffers harm at a time if he is in a certain kind of non-comparatively bad state at that time.

This account of when a person suffers a harm is non-comparative, since, according to it, we need not compare how someone actually fares with how she otherwise would have fared to determine whether she has been harmed. Elizabeth Harman fleshes out NCAH. She gives the following sufficient condition for an act's being harmful – “(1) One harms someone if one causes him pain, mental or physical discomfort, disease, deformity, disability, or death.”⁶⁹

She then offers the following principle to give a unified account of what makes these various states *harmful*:

- (1) An action harms someone if it causes the person to be in a bad state. Bad states are understood as states that are in themselves bad, not bad because they are worse than the state the person would otherwise have been in.⁷⁰

Perhaps *this* version of NCAH can solve the Non-identity Problem. Pebbles is, after all, caused to be in a state that could be considered bad in itself by being caused to exist blind. On

⁶⁹ Harman (2009: 139). This “unifying” principle is not very informative. It might turn out that, when fully fleshed out, Harman’s view looks identical to the problematic view endorsed by Shiffrin.

⁷⁰ Ibid: 139-140.

Harman's view blindness is a state that is harmful to be in, even though it is not worse than non-existence (Pebbles' only other option).

Harman's version of NCAH seems capable of solving both the immediate and eventual versions of the Non-identity Problem. It classifies both the condition of blindness suffered by Pebbles and the conditions of disease and deprivation suffered by the depleted population as bad states, and so it classifies the actions that bring about these conditions as harmful actions. It thus satisfies the sufficient generality criterion. Her view also seems to satisfy the non-*ad hoc*ness criterion. She offers an account of harmful actions that is, at least initially, plausible apart from its ability to solve the Non-identity Problem. For instance, it is intuitively plausible, as her account states, that people are ordinarily harmed when they are caused to be in a disabled state, a state that it is bad to be in. Harman's view, though it is sufficiently general and not *ad hoc*, is unable to satisfy the Intuitive Cost Criterion.

The first trouble for NCAH is that it might seem unable to account for preventative harm. We considered the problem of preventative harms when discussing objections to CCAH. A preventative harm occurs when an agent prevents some individual from receiving a benefit. Preventative harms do not necessarily cause the victim of the harm to be in a non-comparatively bad state (pain, mental/physical discomfort, disease, etc.) When I intercept my neighbor's paycheck, for instance, I preventatively harm him even though I do not cause him to be in a non-comparatively bad state. I do, however, cause him to be *worse-off* than he otherwise would have been. CCAH naturally accommodates the existence of preventative harms while NCAH cannot. Harman would likely balk at this objection, however, for she has only provided a *sufficient* condition for harm in (1). It is open to her to claim that there might be *other* conditions the

obtaining of which are also sufficient for suffering a harm. I think it is reasonable to ask whether this is an unsatisfying reply, but I will let it stand and turn instead to a different worry.⁷¹

The first hint of trouble arises when we scrutinize Harman's description of the bad conditions listed in (1) and described in (2) as "bad in *themselves*". Something that is bad in itself is also sometimes called "intrinsically bad" or "bad apart from its consequences." In the case of states that are bad in themselves, this would mean that they are bad for the person who is in them, insofar as they are in that state, regardless of what results from it. My first problem for Harman's view, then, arises from her characterization of the bad states, the causing of which constitutes a harm, as "bad in themselves." I suspect that the states she lists are not bad in themselves but only sometimes bad because of what they bring about. There is a kind of example that can confirm my suspicion. Compare the following cases:

Cave Dweller: George is a member of a cave-dwelling society. The members of the society live a life of complete darkness. They go about their lives using heightened senses of hearing, touch, and smell. Though eyesight is useless to them, most of them nonetheless have the power of eyesight. George, on the other hand, has developed a rare degenerative disorder that has slowly made him blind. He never notices it, of course, since he has always lived, and always will live, in complete darkness.

⁷¹ Hanser (2008) discusses this problem for NCAH.

Basement Dweller: Jane lives in a basement in New York. Like most people, she depends heavily on her eyesight to navigate the world. She has normal vision until one day a rare degenerative disorder catches up with her, causing her to be blind. As a result, Jane can no longer leave the house by herself, has trouble interacting with others, and generally struggles with day to day activities that used to be routine.

It seems that Jane has, by losing her sight entered a state that is bad for her, but that George, by losing his sight, has not. But if blindness counts as a disability, and disabilities are states that are bad in themselves (regardless of their consequences), then we ought not to have this reaction. Either both of them enter a state that is bad or neither does. I think that the examples illustrate that disability, one of the states listed by Harman as “bad in itself”, does not belong on that list. Blindness is not a bad state for George to be in because it has no *further effect* on his welfare. His life is no *worse* as a result of becoming blind than it would have been had he remained sighted. Disability is a bad state to be in not because it is bad *in itself* but because of what it generally causes (as in the case of Jane). Harman seems to misidentify the sufficient conditions for a person suffering a harm.

Harman might attempt to restrict her list of states that are bad in themselves so that it excludes states of disability. I take my example to have posed a more significant challenge to the non-comparative view, however. I have provided a reason to think that all conditions that are supposedly bad *in themselves* are actually only bad insofar as they make someone’s life worse. Blindness is simply one example. Other examples could demonstrate the same point about other

bad states. Furthermore, even if Harman can respond to my objection simply by eliminating disability from her list of bad states, this would leave her account of harm unable to solve the Non-identity Problem. Presumably, if Pebbles is harmed, this is because she is caused to be in a disabled state. But if Harman concedes that disabled states are bad only insofar as they make someone's life worse than it otherwise would have been, then Pebbles, like George, is not in a bad state. But if she is not in a bad state, then the non-comparative account of harm does not entail that she suffers a harm. Therefore, by removing disability from its list of bad states, the non-comparative harm-based solution is rendered unable to solve the Non-identity Problem.

Another trouble with NCAH as Harman construes it is that it seems to classify certain actions as wrongful harms that are clearly not. Consider the case of a doctor who improves the vision of a completely blind person so that he is now only mostly blind (but still disabled). The doctor, on Harman and Shiffrin's view, has caused the patient to be in a non-comparatively bad state even though he has significantly improved the patient's welfare level along the dimension of eyesight. Harman and Shiffrin seem committed to maintaining that the doctor has harmed the patient when, intuitively, not only has he not harmed him either in-a-respect or all-things-considered; he has benefitted him.⁷² Because both Shiffrin's and Harman's versions of NCAH have decidedly counterintuitive implications for some set of the cases discussed above, the solution on which NCAH is based fails the Intuitive Cost Criterion for a successful solution to the Non-identity Problem. It ought to be rejected on this basis.

⁷² This example is borrowed from Hanser (2009: 187-188).

II.3.2. B – The De dicto Account of Harm

Some have rejected CCAH in favor of a view according to which *de dicto* harm sometimes supplants *de re* harm as the kind that is relevant to the moral assessment of an action.⁷³ Before describing the argument for the *de dicto* view in detail, it is important that I first make clear the distinction between *de re* and *de dicto* harms. This distinction between types of harms has its origin in a distinction made in philosophy of language between reference *de dicto* and reference *de re*. When I utter a declarative sentence involving a definite description such as “the President of the United States has a duty to protect the Constitution”, I might be picking out only the present holder of the office, President of the United States, and saying of *her* that he has a duty to protect the Constitution. This would be the *de re* reading of the sentence. Alternatively, I might be referring to *anyone* who might be picked out by the description “the President of the United States” and saying, of *her*, that she has a duty to protect the Constitution. This would be the *de dicto* reading of the sentence. The *de re* reading rigidly designates the current President whereas the *de dicto* reading picks out different individuals in different contexts.

Just as utterances about the duties of someone who satisfies a definite description can refer in virtue of either *de re* or *de dicto* senses, when we talk of something being good, bad, better or worse for someone who satisfies some definite description, we can have either a *de re* or *de dicto* reading in mind. Consider, for example, two ways in which I might improve the health of my cat: I might buy my actual cat healthier food so that he loses weight. Another way to improve the health of my cat would be to buy a healthier cat and to send my actual obese cat to the pound. Only the first way of improving the health of my cat improves his health *de re*. The

⁷³ Most notably, Caspar Hare (2007). Christian Coons (manuscript) also defends a version of DD.

second way improves the health of my cat *de dicto* (but obviously not *de re*) by causing it to be the case that the office holder of the description “my cat” is healthier than some distinct cat that could have held the office would have been. Casper Hare characterizes the distinction between *de re* and *de dicto* betterness in a technical but helpful way:

De Re Better.—Where S1 and S2 are states of affairs, S1 is *de re* better for —— than S2, when the thing that is actually —— is better off in S1 than in S2.

De Dicto Better.—Where S1 and S2 are states of affairs, S1 is *de dicto* better for —— than S2, when the thing that is —— in S1 is better off in S1 than the thing that is —— in S2 is in S2.⁷⁴

Hare claims that there are special cases in which agents have duties to reduce *de dicto* harm even though doing so makes particular individuals *de re* worse off. He offers the following example to illustrate his point:

The Safety Officer—Tess is a state safety officer, whose job it is to regulate those features of the automobile that protect its occupants in the event of a collision—air bags, crumple zones, and so forth. Noticing that people in her state are not wearing safety belts, she implements some tough new

⁷⁴ Hare, p. 514 (2007). Paraphrased.

regulations and, a year later, is pleased to discover evidence that they have been effective, that the severity of injuries sustained in automobile accidents has been reduced as a result of people belting up. She gives herself a pat on the back.⁷⁵

Note that the implementation of the new safety regulations affects many minor changes in the timing of accidents. The time it takes to put on one's seatbelt may mean the difference between getting hit by a bus and driving to work without incident. Because of these minor changes, the regulations in part determine which people will be involved in accidents. The actual people who find themselves in accidents would not have been so involved if the regulations had not been implemented. Thus, it turns out that the regulations are *de re* worse for the actual accident victims. Hare claims, and I agree, that it would nonetheless be inappropriate to criticize Tess for failing to do her job. She has done her job quite well. She has done her job well, Hare argues, because she has a *de dicto* duty to ensure that accident victims (whoever this description ends up referring to) are as well off as possible. She has met this *de dicto* duty and so she has done her job well. According to Hare, Tess can avail herself of the following reply to anyone who criticizes her for doing what is *de re* worse for accident victims:

“Yes, what I did was *de re* worse for last year's accident victims. But we regulators always know, long before we decide what substantive measures

⁷⁵ Ibid, P. 516

to take, that the substantive measures we will actually take (whether they involve reducing speed limits a little, increasing them a lot, adding more traffic lights or spreading nails all over the roads of America) are going to be *de re* worse for accident victims. My job was not to make things *de re* better for last year's accident victims, but to make it the case that last year's accident victims were, collectively speaking, healthier than those people who would have been last year's accident victims would have been if I had acted otherwise. In brief, my job was to make things *de dicto* better for the health of last year's accident victims. And I did just that.”⁷⁶

Hare thinks that this case shows that sometimes agents can have duties to make things *de dicto* better even when doing so makes things *de re* worse for particular individuals. If Hare is correct that agents sometimes have *de dicto* duties to make things better even when doing so does not make things *de re* better, we now have a neat explanation of what is wrong with what Wilma and Fred do when they conceive Pebbles. They do not make things *de re* worse for their child (remember that Pebbles has a life well worth living). They make things *de dicto* worse for their child by conceiving a child who is worse off than some other child that they could have conceived would have been. Like Tess, Wilma and Fred have a *de dicto* duty. In their case the *de dicto* duty is to promote the health/welfare level of their child, where “their child” is read *de dicto*, and they violate that duty by conceiving Pebbles instead of a sighted child.

⁷⁶ Hare, p. 516-517 (2007).

While I agree with Hare that Tess' duty is to make it the case that last year's accident victims were healthier than those people who would have been last year's accident victims would have been if I had acted otherwise, this fails to establish that Tess' duty is a *de dicto* duty rather than a *de re* duty. Consider the following case:

The Safety Officer 2: Seeking to make this year's accident victims as healthy as possible (*de dicto*), and operating on a limited budget, Tess is struggling to choose between two competing options. She can either install traffic lights at dangerous intersections in Orlando, Florida or she can install those traffic lights at dangerous intersections in Boulder, Colorado. Boulder residents are disproportionately fit relative to average American cities, and Orlando residents are of average health. If Tess reduces the number of traffic accidents in Orlando, thereby failing to reduce the number of traffic accidents in Boulder, she will ensure that this year's accident victims are as healthy as possible in a *de dicto* sense. She thus has an obligation to install the lights at Orlando intersections instead of Boulder intersections.

De re betterness and worseness for accident victims are morally irrelevant for Tess' purposes (we have already supposed that anything she does will be *de re* worse for this year's accident victims).

While I do not think that it would be wrong for Tess' to install traffic lights in Orlando rather than Boulder, she is certainly not morally required to do so. Furthermore, the reasoning

that underlies Tess' *de dicto* requirement is perverse. Surely Tess cannot satisfy her duty by allowing *harm to the healthy rather than the unhealthy*. Yet, if Tess really only has a *de dicto* duty to ensure that accident victims are as healthy as possible, then it is difficult to explain what is perverse about this reasoning. What seems perverse to me is that, not only is Tess required to minimize the severity of accidents, but she is required to ensure that healthier people *are involved* in accidents. This troubling implication of the *de dicto* account of harm means that the solution on which it is based fails to satisfy the Intuitive Cost Criterion for a successful solution to the Non-identity Problem. The view has implications that are more problematic than the conclusion of NIPA.

The defender of the *de dicto* solution might object that I have misconstrued his view. Tess' *de dicto* duty to ensure that accident victims are as healthy as possible, he might contend, is derived from a more fundamental *de dicto* duty to ensure that *whoever* has an accident is involved in a *less serious accident* than otherwise would have occurred. But I think that this explanation of Tess' *de dicto* duty indicates that, fundamentally, Tess does not have a *de dicto* duty at all. On this construal of Tess' *de dicto* duty, there is a way to describe the duty in such a way that it is derived from a more fundamental duty to minimize *de re* harm to accident victims. Let's consider The Safety Officer again. Presumably, the amount of *de re* harm suffered by the actual accident victims of last year, as a result of Tess' regulations, was less than the amount of *de re* harm that distinct accident victims would have suffered if Tess had failed to implement the regulations. In other words, the people who are in accidents as a result of Tess' regulations are worse off than they otherwise would have been *to a lesser extent* than the people who would have been accidents victims would have been had she not implemented the regulations. So it seems that we can explain why Tess has done her job well simply by appealing to the effect of

her safety regulations on the net *de re* harm suffered by accident victims. By implementing the safety regulations, she has minimized net *de re* harm to accident victims, where net *de re* harm is determined by summing the total *de re* harm suffered by actual particular accident victims (calculated by summing the extent to which each victim is worse off than he otherwise would have been) and then subtracting the total projected *de re* harm that distinct particular accident victims *would* have suffered (calculated by summing the extent to which each victim would have been worse off than he otherwise would have been) if the regulations had not been implemented. *This* is the method we must use to determine whether Tess succeeds in satisfying her supposedly *de dicto* duty to ensure that *whoever* has an accident is involved in a less serious accident than otherwise would have occurred. She satisfies this duty just in case she minimizes net *de re* harm. *De re* harm turns out to be the fundamental consideration, even in cases of supposedly *de dicto* duties. But if the satisfaction of quintessential *de dicto* duties is a function of calculating net *de re* harm, then we cannot solve the Non-identity Problem by claiming that Wilma violates her *de dicto* duty to her child. Wilma's choice to conceive Pebbles rather than Rocks does not result in greater net *de re* harm than her choice to conceive Rocks rather than Pebbles would have resulted in. Neither child suffers *de re* harm.

My explanation might leave out one important detail. Had Tess refrained from imposing new safety regulations she would not have *done* something to make any particular accident victim worse off than he otherwise would have been. She would have merely *allowed* particular accident victims to be worse off than they otherwise would have been. By implementing the regulations, however, she *does* something that makes particular accident victims worse off than they otherwise would have been in order to minimize net *de re* harm. In many cases, however, it seems morally worse to do harm than to merely allow it. In II.3.1.A I considered a pair of cases

that illustrated this point. While it is permissible to allow one to die in order to save five, it is not obviously permissible to kill one in order to save five. This seems to show that it is harder to justify doing harm than it is to justify allowing harm. But how, then, do we explain why this normally morally relevant consideration (doing versus allowing harm) seems to be irrelevant in our moral assessment of Tess' action? There are two ways to address this question: a hard-line way and a moderate way. The hard-line way is to claim that the doing/allowing distinction is morally irrelevant and cases like Safety Officer are evidence for its irrelevance. I don't think that this hard-line reply is necessary, however. The moderate reply will suffice. According to the moderate reply, there are some circumstances where the doing/allowing distinction does not matter. Tess finds herself in one of those circumstances. Perhaps the doing/allowing distinction is morally irrelevant in virtue of Tess' occupying a special role in society. Maybe causal proximity matters; Tess' policies are a remote cause of the accidents but a direct cause of the reduction in harm to the victims of accidents. In any case, it is apparent that Tess is responsible for minimizing net *de re* harm to accident victims, and it does not matter if the harm is harm that she causes or merely allows.⁷⁷

Hare's *Safety Officer* case fails to establish that *de dicto* harm is morally relevant apart from its relationship to *de re* harm. Hence, an appeal to *de dicto* harm fails to satisfy the Non-*ad hocness* Criterion for a successful solution to the Non-identity Problem. Because the *de dicto* solution, insofar as it capable of satisfying the Intuitive Cost criterion, collapses into the counterfactual comparative *de re* account of harm, there is no reason to prefer the *de dicto*

⁷⁷ The intend/foresee distinction might very well be morally relevant to the permissibility of Tess' regulations. A policy on which particular individuals are intentionally caused to be involved in accidents in order to benefit other drivers might seem morally worse than a policy on which such occurrences are merely foreseen.

solution over the other candidates. The only reason to accept it is that doing so would block the inference from NIPA's premises to its conclusion. But this reason is shared by all candidate solutions.

II.3.2.C – The Event-Based Account of Harm

Matthew Hanser defends what he has called an “event-based” account of harm. I will call it “the event-based account”. On the event-based account, a person's suffering a harm does not involve her being worse off than she otherwise would have been. Nor does it necessarily involve her being in a state that is bad for her to be in. According to the event-based account, to suffer a harm is simply to be subject to an event of harming. An event's status as a harm is not derivable from either the badness of the resulting state or the worseness (relative to some available alternative) of the resulting state. How, then, does one determine whether a particular event is a harming event? In Hanser's words:

- (i) Someone suffers a [non-derivative] harm with respect to a certain basic good if and only if he loses some quantity of that good. Someone receives a [non-derivative] benefit with respect to a certain basic good if and only if he acquires some quantity of that good. Every event of the schematic type S's losing some quantity of basic good G is someone's suffering of a non-derivative harm; every event of the schematic type S's acquiring some

quantity of basic good G is someone's receiving of a non-derivative benefit.⁷⁸

Hanser is careful to note that when he uses the term "good" he is not referring to a state that it is good (better) to be in. Rather, he is referring to a thing that it is "good to have". "Basic goods" are a subset of the general category "goods". These goods are "those the possession of which makes possible the achievement of a wide variety of the potential components of a reasonably happy life."⁷⁹ Hanser includes basic powers and capabilities in the category of basic goods. Eyesight, for example, is a basic good. When you lose a basic good you suffer a harm.

Not all harms and benefits, on Hanser's view, consist in the loss or acquisition of basic goods. There are also preventative harms and benefits. Hanser describes this type of harm in the following passage:

(ii) Someone suffers a [derivative/preventative] harm with respect to a certain basic good if and only if he is prevented from receiving a level-n benefit with respect to that good. Someone receives a level-[derivative/preventative] benefit with respect to a certain basic good if and

⁷⁸ Hanser (2008: 441). Parentheses added.

⁷⁹ Ibid: 440.

only if he is prevented from suffering a level-n harm with respect to that good.⁸⁰

According to the event-based account, all harming (benefitting) events are either direct, consisting in some loss (gain) of some good, or preventative, consisting in the prevention of a loss (gain) of some good. There is one important thing to note regarding Hanser's view. On his view, there can be a harming event that does not result in someone being in either a bad state or a state that is worse, from the standpoint of her interests, than some alternative (Hanser's example is the harm of death). An event's status as a harm in no way depends on the badness of some resulting state for the victim. For example, becoming blind, Hanser claims, is a harm regardless of the badness of the state in which one finds oneself (with regard to welfare) as a result. Perhaps it is non-comparatively bad and comparatively worse to be in a state of blindness (when the alternative was sightedness), but, according to Hanser, these facts do not explain why becoming blind is a harm. It is a harm because it constitutes *the loss of a good*.

The success of the event-based account of harm, supposedly, does not depend on any particular view about which goods are *basic*. Hanser indicates, however, that the basic goods should be understood as instrumentally rather than intrinsically valuable. He says, "the basic goods are those the possession of which makes *possible* the achievement of a wide variety of the potential components of a reasonably happy life."⁸¹ This strongly implies that the basic goods are not valuable in and of themselves, but, rather in virtue of their relationship to the components of

⁸⁰ Ibid: 442. Parentheses added

⁸¹ Ibid, p. 440. Emphasis added.

a life that make it go well. Experiencing pleasure or having one's desires satisfied, for example, might count as components of a happy life. If, say, experiences of pleasure were the only components of a happy life, then a basic good such as eyesight would be valuable insofar as it made possible the achievement of experiences of pleasure.

I shall now turn to some problems for the event-based account. In this section I argue that the event-baser's treatment of death gives rise to more problems than it solves. That the event-based account has such problems should be surprising, because its treatment of death is taken to be one of its virtues. I also argue that the event-based account cannot adequately handle the phenomenon of "non-regrettable" disabilities.

II.3.2.C.a – Death and Losing Goods

Recall that DW, a principle that I have been assuming in this chapter, states that someone has welfare properties at a time t only if she is located at t . Hanser appeals to DW when he criticizes CCAH's treatment of death, but he thereby leaves himself open to be hoisted by his own petard. According to the event-based account, the harm of death consists *entirely* in the loss of certain basic goods (powers). But DW entails that when someone dies she ceases have any welfare properties. Normally, if a quantity n of some good g has been lost in some event e there is a person s who exists at a time t with a quantity n of the good and who also exists at time $t+1$ but who now possesses a smaller quantity $n-x$ of that good. That is the natural way to determine that an event of losing a good has occurred in ordinary cases; we compare two quantities of the good. One is the quantity that the person had in the past and the other is the quantity she has

now. It is not clear, however, how to go about applying this method of determining that a loss of a good has occurred in the case of death. If, as Hanser claims and as I have supposed, a person who dies ceases to have any welfare properties at all, then she does not have zero goods when she is dead. She does not *exist* to have zero of anything. But then, after a person dies, she does not possess some smaller quantity $n-x$ of some good that we can compare to the quantity n of the good that she possessed prior to her death. Without two quantities to compare, the event-baser cannot determine whether a person has suffered a loss of a good when she dies.

A solution to this problem for the event-baser would be to contend that a person suffers a loss of a good just in case she *no longer has quantity n* of the good, where n is any positive quantity. Since it is true that a person ceases to have quantity n of any goods when she dies, it would be true that she suffers a loss of those goods when she dies. This simple reply will not suffice. Though it is true when she dies that she no longer has quantity n of the good, this does not mean that she has quantity 0 of the good. She no longer has quantity n of the good because she no longer has a welfare level or any other properties at all. Her welfare properties are *undefined*. This is entailed by DW, a principle endorsed by Hanser. Furthermore, even in ordinary cases, someone's no longer having some quantity n of a good is not, by itself, sufficient for it to be the case that someone has suffered a loss of that good. For instance, a person who had quantity n of some good but who is subsequently benefited by receiving an additional quantity x of that good also ceases to have quantity n (she now has $n+x$), but she does not suffer a loss. It is thus unclear how the defender of the event-based account is to go about determining that a loss of a good has occurred when a person dies. But if it is unclear how the event-baser is to go about determining that a loss of a good has occurred when a person dies, then it is unclear how the event-baser is to explain death's harmfulness. Hanser claims that the event-baser "need take no

position on the question whether events are analyzable in terms of states.”⁸² To avoid the problem raised here for the event-based account, however, he does need to take a position on this matter.

II.3.2.C.b – Death and Deprivation

A second problem for the event-based account is that it fails to accommodate the intuition that death is harmful, not simply because it constitutes the loss of certain basic goods, but because it deprives its victim of many years of enjoyable life. The whole harmfulness of death, according to the event-baser, is accounted for in a moment: the moment the basic goods are lost. But death is harmful, not just because it deprives us of the goods we had *prior* to death, but because of the *future* goods we would have enjoyed had our lives not been cut short.

The event-baser might appeal to derivative harm to accommodate the deprivation intuition. Recall the event-baser’s account of derivative harm: Someone suffers a [derivative] harm with respect to a certain basic good if and only if he is prevented from receiving a [non-derivative benefit] with respect to that good.

Perhaps death, in cutting our lives short, prevents us from receiving a non-derivative benefit, where a non-derivative benefit is the receiving of some quantity of a good. This might satisfy the intuition that death is harmful partly because of what it deprives us of. The question then is *what good are we prevented from receiving when we die?* Recall that, according to the event-based account, the relevant goods are certain powers and capabilities such as eyesight. If

⁸² Ibid, p. 443

these are the goods the prevention of the receiving of which makes death harmful to the person who dies, then it seems that death will harm us in this way only in a limited range of cases. When many people die they already possess all of the relevant powers and capabilities that they will ever have, so they are not prevented from receiving more of them when they die. CCAH is able to explain why death comes to these people as a harm, but the event-based account is not.

Everyone who dies is of course prevented from the *further enjoyment* of her powers and capabilities. But the event-baser cannot appeal to this kind of consideration in accommodating the deprivation intuition without including, as a component of the view, the problematic state-based comparisons described in section I. To see that this is the case, suppose that the event-based account holds that death is a preventative harm because it is an event of preventing the person who dies from further enjoying her powers and capabilities. The event of dying prevents a person from the further enjoyment of her powers and capabilities, and it thereby *deprives* her of these goods. By making this move the event-baser can accommodate the deprivation intuition by appealing to preventative harm. The trouble for this move is that death prevents someone from further enjoying her powers and capabilities only if she loses those powers and capabilities when she dies. But to determine that a loss of a good has occurred we must compare two quantities of the good. One is the quantity before the losing event occurs and the other is the quantity after the losing event occurs. But as I explained in the previous section, because someone who is dead does not have any welfare properties, there does not seem to be a *quantity* of any good after a person's death to compare with the quantity enjoyed by the person before her death. Thus the

event-baser seems to be left without the tools that he needs to accommodate the unshakeable intuition that death is harmful, in part, in virtue of the goods of which it deprives us.⁸³

I might be accused of begging the question against the event-baser by appealing to deprivation in critiquing his account. After all, the event-baser will likely hold that the harmfulness of death simply cannot be accounted for by appealing to deprivation, so I cannot undermine his account by appealing to the intuition that death is harmful, in part, because of what it deprives us of. I do not think any questions have been begged, however. My claim about deprivation and death was a claim about ordinary *judgments* about the harmfulness of death. I did not appeal to an implication of some theory about the harmfulness of death. Death seems to be harmful in part because of what it deprives us of, and a theory that is able to accommodate that view is, in at least one respect, preferable to one that cannot. Moreover, I have already defended a version of CCAH (CCAH''++) that can explain why death is harmful, in part, because of what it deprives us of. Therefore, the comparativist seems to fare better with regard to the deprivation intuition.

II.3.2.C.c – Non-regrettable Disabilities and Collapse

A separate and final problem for the event-based account targets the most fundamental difference between the event-baser and the comparativist. That fundamental difference is that,

⁸³ Perhaps Hanser can account for this intuition if he takes something like Marquis's view according to which a future-like-ours is at least in some sense something that you already have. That future would then be a good that you lose via the event of your death. Though this is an interesting reply, there seems to be an obvious problem with it. A future is not a good that you have in the same way that your power of eyesight is a good that you have. For one thing, there are goods in the future (in Hanser's sense of the term) that we will not have in the present. It would be strange, for instance, to say that a child has the power to walk now because she will have that power in the future.

according to the event-baser, an event's status as a harm is entirely divorced from the badness (or worseness) of the state in which it leaves its victim. Hanser concedes that the *severity* of an event of harming in some sense depends on the badness (or worseness) of the state in which it leaves its victim, but he denies that whether that event *qualifies* as a harm in any way depends on the badness of that state. He says "when losing a good does put one into a bad state, the duration of that state affects the seriousness of the harm one suffers in losing the good."⁸⁴ The loss of sight might result in someone being in a state that is comparatively worse than the state she otherwise would have been in (supposing the alternative was sightedness), but, according to Hanser, these facts do not explain why becoming blind is a harm. Losing one's sight is a harm solely in virtue of the fact that it constitutes *the loss of a good*.

But this is backwards, and an example involving what I call a "non-regrettable disability" demonstrates just how backwards it is. My aim in providing this example is to show that the intuition that someone suffers a harm when she loses a basic good depends on a more fundamental intuition that the harmed person is worse off than she would have been had she not lost the good. Insofar as we are inclined to have the intuition that the loss of a basic good (eyesight for example) is harmful, this is due to the difficulty of imagining someone becoming blind without also imagining that she has ended up in a state that is worse for her than some alternative (being sighted). This becomes obvious when we consider cases where the two properties are pried apart. When we consider a case in which a person is caused to lose a basic good (eyesight, for example), the loss of which clearly does *not* result in the person being worse-off than she otherwise would have been, it is clear that the mere loss of eyesight *itself* does not

⁸⁴ Hanser 2011, p. 467

constitute a harm. In II.3.2.A we considered a case where someone suffers a “disability” that leaves him no worse off than he otherwise would have been. Let us reconsider this example of a non-regrettable disability:

Cave Dweller: George is a member of a cave-dwelling society. The members of this society live a life of complete darkness. They go about their lives using heightened senses of hearing, touch, and smell. Though the capacity to see is useless to them, most of them nonetheless have the power of eyesight. George, on the other hand, has developed a rare degenerative disorder that has slowly made him blind. He never notices it, of course, since he always has lived, and always will live, in complete darkness.

It is clear that George, in losing his eyesight, does not suffer a harm. CCAH is consonant with this intuition; George is not worse off than he otherwise would have been by becoming blind, so he does not suffer a harm. Blindness does not affect the value of his life whatsoever. The event-based account of harm, on the other hand, must say that George has suffered a harm in losing his eyesight because he has lost a basic good.

The event-baser might contend that powers and capacities, such as eyesight, are only basic goods in *certain contexts*. In some societies, for example, having dyslexia is a major disadvantage, whereas in other societies (those without written language, for example) dyslexia might not even count as a disability. In *Cave Dweller*, the capacity for eyesight is not a basic

good, because it is *useless* in that particular context. Because George loses a useless capacity, he does not lose a basic good, and so he is not harmed. Hanser indicates that he would endorse this reply when he says, “the basic goods are those the possession of which makes *possible* the achievement of a wide variety of the potential components of a reasonably happy life.”⁸⁵ Since George’s eyesight has no relationship to his ability to achieve components of a happy life, his eyesight does not count as a basic good.

This reply requires establishing some non-arbitrary way to distinguish between contexts in which a capacity makes possible the achievement of the components of a happy life, and contexts in which the same capacity does not. Distinguishing between these contexts spells yet more trouble for the event-baser. The most obvious, plausible way to determine whether some capacity makes possible the components of a happy life in some particular context is to establish whether someone would be *worse-off* without it than with it in that context. This means that the event-baser should offer the following explanation of why George is not harmed by the loss of his eyesight: George is not harmed by the loss of his eyesight, because eyesight is not a basic good for him in the context of the cave-dwelling society. Eyesight is not a basic good for him in the context of the cave-dwelling society, because George is not better off for possessing the power of eyesight than he would be in its absence (it does not help him to achieve the components of a happy life). This is just a roundabout way of saying that George is not harmed by the loss of his eyesight, because he is not worse off for losing his eyesight than he otherwise would have been. An event of losing a good such as eyesight turns out to constitute a harm only if it results in someone being worse off than she otherwise would have been. In accommodating

⁸⁵ Hanser 2008, p. 440

the intuition that George is not harmed by the loss of his eyesight, the event-based account collapses into CCAH. Because the event-based account gets the wrong result in cases like the *Cave Dweller* case and in the case of death, the solution that depends on the event-based account fails to satisfy the Intuitive Cost Criterion for a successful solution to the Non-identity Problem.

A final trouble for the event-based account is that it also fails to satisfy the Sufficient Generality Criterion for a successful solution to the Non-identity Problem. Because Hanser's account requires that a person suffer a *loss* of some good in order to be harmed, and Pebbles never suffers a loss of some good by coming to exist without sight, his view cannot explain what is wrong with causing her to exist rather than an alternative sighted child. She is not wrongfully harmed, because she never loses a basic good. She 'starts off' blind. Even if we were to accept Hanser's account as a solution to the eventual version of the Non-identity Problem, we would still need to find another solution to account for the wrongness of bringing Pebbles into existence rather than a sighted child. If we can find a solution that solves both versions of the Non-identity Problem, then we ought to prefer it.

Conclusion

In this chapter I have canvassed the most promising harm-based solutions to the Non-identity Problem. I argued that the Counterfactual Comparative Account of Harm cannot be used to solve the Non-identity Problem. I also argued that alternative accounts of harm also fail to solve the Non-identity Problem. Despite the general appeal of harm-based solutions, each falls afoul of either the Non-*ad hoc*ness, Sufficient Generality, or Intuitive Cost Criterion. If we are to

find a solution to the Non-identity Problem, it must appeal to considerations other than harm. I turn to the category of non-harm-based solutions in chapter III.

III

Non-Harm-Based Solutions to the Non-identity Problem

In Chapter II I considered solutions to the Non-identity Problem that appealed to the moral consideration of harm. I found that each of these solutions falls afoul of the *Non-ad hocness*, the Intuitive Cost, or the Sufficient Generality Criterion for a successful solution. Benatar's comparative harm-based solution, for instance, comes at much too great an intuitive cost to be worth adopting. His solution would force us to accept the view that we are all worse-off for having come to exist. Each of the non-comparative, event-based, and *de dicto* harm-based solutions also fail the Intuitive Cost, the *Non-ad hocness*, or the Sufficient Generality Criterion. The event-based account of harm, for instance, seems unable to explain why death comes to us as a harm. This renders the event-based account of harm unsuccessful, along with any solution to the Non-identity Problem that is based on it. I concluded that there is no successful harm-based solution to the Non-identity Problem on offer. In this chapter I will consider an array of solutions that appeal to moral considerations other than harm in addressing the Non-identity Problem. The category of "non-harm-based solutions" is rather broad. Within this category there are two subcategories. One subcategory includes solutions that reject P4 of NIPA. I will refer to this first subcategory as the set of "rights-based solutions." Each of the solutions within this category appeals, in some way, to considerations of rights. The other subcategory includes solutions that reject P5. I will call this second subcategory of solutions "non-person-affecting solutions." Each of the solutions in this category appeals to the general claim that an act can be morally wrong

even when it fails to wrong anyone. (P4) says: if an act does not harm anyone then it does not wrong anyone. (P5) says: if an act does not wrong anyone then it is not wrong.

In the first main subsection of this chapter I will consider rights based solutions to the Non-identity Problem. In the second main subsection of the chapter I will consider non-person-affecting solutions.

III.1 – Rejecting P4: Rights-Based Solutions (Wronging Without Harming)

In defending one of the first rights-based solutions to the Non-identity Problem, James Woodward accepts (or at least remains agnostic about) the Counterfactual Comparative Account of Harm. He rejects, on the other hand, the claim that one must be harmed in this sense in order to be wronged. Woodward only considers the Counterfactual Comparative Account of harm in making his assertion, but we need not reconsider the alternative accounts of harm; I showed in the previous chapter that the Counterfactual Comparative Account of Harm is preferable to those alternatives. Woodward does not discuss Wilma's case in particular, but his view clearly has implications for that case. Remember that the challenge of the Non-identity Problem, as I have formulated it, is to explain why Wilma does something wrong by conceiving Pebbles.

Woodward's main contention is that someone can be wronged by having a right violated without being harmed (in a counterfactual comparative sense). He appeals to some very clever cases as evidence for his view. Consider one example:

Promise: You promise your friend that you will perform a conjunctive action AB.

You could have either made this promise or not. You knew when you made it that if you promised to do AB you would only do A, but that your friend would be better off as a result of you doing A than she would have been had you not promised to do AB and thereby failed to do either A or B.

About this case Woodward suggests that the obvious intuition is that you ought not to have made the promise to your friend, even though by doing so you make her better off than she otherwise would have been. He thinks that the reason you ought not to have made the promise is that by making the promise you ensure that you will lie to her later, and lying will violate her rights even though it will also leave her better off, all-things considered. Now consider one last case meant to illustrate a similar point.

Discrimination: Suppose that Smith, who is black, attempts to buy a ticket on a certain airline flight and that the airline refuses to sell it to him because it discriminates racially. Shortly after, that very flight crashes, killing all aboard.

There is a clear sense in which the airline's action has the result that Smith is better off than he otherwise would be, and if selling or not selling Smith the ticket are the only relevant actions which the airline can perform, not selling leaves him better off

than any other possible action the airline might have performed. Nonetheless, it seems quite natural to say that the airline's action wrongs Smith.⁸⁶

Each of these cases presents us with a victim who is wronged by an action that makes her better off, all things considered, than she otherwise would have been. If we agree with Woodward about these judgments then it might seem that we ought to reject P4. Indeed I think that he shows that P4, strictly speaking, is false. Someone can be wronged without being harmed. If P4 is false then perhaps we have a solution to the Non-identity Problem. It would leave open the possibility that the members of the depleted society, who live difficult but non-terrible lives, have their rights violated even though they are not worse off than they otherwise would have been living in a resource-depleted world. Or perhaps Pebbles, we can contend, has her rights violated by being caused to exist in her blind state even though she is not worse off than she otherwise would have been for having the condition. The trouble is that P4 can be modified in a suitable way such that NIPA still arrives at its troubling conclusion for both the immediate and the eventual versions of the Non-identity Problem. A plausible modification of P4 will allow NIPA to generate the Non-identity Problem even if we reject P4. I will develop this modification of P4 in the next two sections.

⁸⁶ Woodward (1986: 810-811).

III.1.1 The Moral Relevance of Consent

Note first that there is an important disanalogy between the cases to which Woodward appeals and the Non-identity cases about which Woodward's cases are supposed to reveal a moral insight. This disanalogy will prove problematic for Woodward's solution. The difference is that in *Promise* and *Discrimination*, there is someone who exists who could have either consented or not consented to the treatment that violates his or her rights. Furthermore, the act which violates the victim's rights goes against the *express wishes* of the victim in a circumstance where the agent has no authority to flout the victim's express wishes. It is possible, then, that these factors explain why the victims in the examples are wronged. In Non-identity cases, on the other hand, there is no one who exists to either consent or to express some preference regarding the treatment that would supposedly violate her rights. Therefore, the fact that the agent failed to obtain consent from the victim or that the agent acted against the victim's express wishes *cannot* explain the wrongness of actions in Non-identity cases. To avoid these conflating factors we need to provide a kind of case that is suitably analogous to Non-identity cases. This will be a case in which something can be done to an existing person who has not expressed an explicit preference regarding her treatment and from whom consent *cannot* be obtained. To mirror the case of Wilma, perhaps we can imagine a case involving a person who is comatose who through some operation can be caused to enjoy a fully conscious life again, but at the cost of suffering from a significant disability for the rest of her life. This kind of case is closely analogous to the case of Wilma, because, like someone who never exists, someone who is comatose does not have the capacity for hopes, experiences, or consent *without some action* first being performed by an agent. Consider:

Comatose: A patient lies in a coma at the hospital. A recent technological development allows a doctor to return the patient to a fully conscious state. The operation will have the side-effect of rendering the patient incurably blind. Her disability will make her life more difficult than it would have been were she to be sighted, but her life will be well worth living. The doctor performs the operation, returning her to her conscious state, happy but blind.

If Pebbles has her rights wrongfully violated when she is brought into existence without sight, then surely the patient in *Comatose* has her rights violated by the doctor when he performs the operation. Hence, according to Woodward's solution, the doctor acts wrongly by returning the patient to her conscious state because he also causes her to be blind. But most will react that the doctor does not act wrongly by returning the patient to her conscious state. Of course, if the doctor could have done the same operation without the cost of the patient's eyesight, then the doctor ought to have done so, but since he could not there is no objection to him performing the operation. I suspect that most react that the doctor does not wrong the patient by performing the operation because the doctor does not so much as *prima facie* harm the patient, and this gives us confidence that the patient *would have* consented to the operation were she able to. First, drawing on the distinction I made between harm, all things considered, and harm in a respect from the previous chapter, we see that the doctor does not cause the patient to be worse off than she otherwise would have been for *any interval of time* after the operation. The patient is better off than she otherwise would have been in the long run, but she is also better off relative to each

interval of time after the operation. Second, this fact gives us confidence that the patient, had she been rational and able to express a preference or to consent to the operation, surely would have preferred that the operation be performed and consented to the operation. But like the patient in *Comatose*, we can suppose that there is no interval of time after she comes to exist during which Pebbles is worse off than she otherwise would have been had she not been conceived. This indicates that if Pebbles was rational and able to consent, surely she would have consented to Wilma's bringing her into existence. Neither Pebbles nor the patient is worse off than she otherwise would have been for any interval of time as a result of Wilma's or the doctor's actions. Because both Pebbles and the patient would have consented were they fully rational and able to consent, neither has her rights been violated. Even if Woodward has shown that an existing person can be wronged by having her rights violated in cases where she is not harmed (e.g. *Promise*, and *Discrimination*), the case of *Comatose* shows that he fails to establish that a person who does not yet exist can be wronged by being caused to exist with a life that is, in some respect, deficient. We can thus modify P4 so that NIPA still goes through.

P4*) A person from whom consent cannot be obtained is wronged by an action only if *either* she is thereby made worse off than she otherwise would have been, all things considered, for some interval of time, *or* she would not have consented to the action were she rational and able to consent.

Since we are supposing at this point that Pebbles is not worse off, all-things-considered, for any interval of time after she comes to exist, we can safely assume that Pebbles, had she been able to consent, would have consented to Wilma's choice. P4* thus allows NIPA to reach the unpalatable conclusion that Wilma does not act wrongly by conceiving Pebbles. P4* generates no objection to Wilma's choice because (i) Wilma could not obtain consent from Pebbles prior to conceiving her, (ii) coming to exist is not worse for Pebbles, all things considered, for any interval of time, and (iii) we can be confident that Pebbles would have consented to coming to exist if she were rational and able to do so.

We might be hesitant to accept P4* because of what it says about the following case.

Nazi Guard: A Nazi guard is walking along the perimeter of a concentration camp and finds a man comatose on the ground just outside. The guard knows that if he is added to the prisoner population the comatose man will suffer many wrongs, and have his rights violated many times. He also knows, however, that these wrongs would not leave the man worse off than if he had remained unconscious, because they would enrich his character and deepen his understanding of life, such that overall his life would be better than it would have been had he not been restored to a conscious state while imprisoned in the camp. Ultimately, the man would not wish that the Nazis had not imprisoned him, because he would so value what he gained from this experience. The guard tosses the man into the concentration camp who is then subjected to many serious wrongs.

P4* fails to raise an objection to the guard's conduct only if (i) consent could not in fact be obtained by the victim prior to the guard's action, (ii) the victim is **not** worse off than he otherwise would have been, all things considered, for some interval of time after he is thrown into the concentration camp, and (iii) the victim **would have** consented to his treatment if he were rational and able to consent. But neither (ii) nor (iii) obtain in the case of the Nazi guard. The man *is* worse off than he otherwise would have been, all things considered, for some period of time, as a result of the guard's conduct (assuming that his life would have been better had he not been in the prison camp during the period of time that he was in fact in the camp), *and* we can be confident that the man would *not* have consented to the guard's conduct were he rational and able to consent (perhaps because the benefits of his treatment in the concentration camp could not be predicted). The Nazi guard thus poses no problem for P4*. Therefore, P4* is an acceptable revision of P4, and, like P4, it allows NIPA to reach its troubling conclusion. Though the cases that Woodward provides show that a person can be wronged even in cases where she is not harmed, the cases fail to show either that Pebbles is wronged even though she is not harmed, and this is what must be shown if we are to solve the Non-identity Problem by rejecting P4.

Woodward's rights-based solution also has counterintuitive implications when applied to a kind of Non-identity case that is different from the Wilma case in one crucial respect. In this kind of case, a parent is *only* able to conceive a disabled child. Even if she waits to conceive, her child will have a disability. The choice for the parent in these cases is between producing no additional person and an additional person with a disability. Most have the intuition that it is permissible to conceive a disabled child under these circumstances, but a rights-based view like Woodward's seems to entail that there is a strong *prima facie* objection to conceiving the blind child even if conceiving a blind child is the parent's only option. This is an unacceptably

counterintuitive implication, especially given that we have seen no reason for thinking that Pebbles' right is violated by Wilma in the first place.

Discrimination, Promise, and Nazi Prisoner seemed to show that a person can be wronged without being harmed all-things-considered just in case consent *can be* but is *not* obtained by the victim. But Pebbles' consent cannot be obtained by Wilma, so our judgments about these examples fail to show that Wilma acts wrongly by conceiving Pebbles. In order to accommodate these examples, and in light of this disanalogy, I provided P4*. P4* is a revision of P4 of NIPA that applies specifically to cases where consent cannot be obtained by the supposed victim. These examples do not pose a problem for the truth of P4*, and replacing P4 with P4* still allows NIPA to reach its unpalatable conclusion. I am tempted to add the following further condition to P4*: a person must *exist* in order to be wronged. In each of the cases that purport to demonstrate that one can have a right violated or be wronged without being harmed, there is someone who exists who is the victim of the rights violation. But Pebbles does not *exist* when the act of conception is performed that supposedly violates her rights. Since it seems plausible that someone must have rights in order to have her rights violated, and that someone who does not exist does not have rights, it will be difficult for a defender of the rights-based solution to reach the conclusion that Wilma acts wrongly by violating Pebbles' rights.

III.1.2 Markie and Smolkin's Solution: Rights-violation without Harm

Doran Smolkin tries to show that there is no barrier to claiming that Pebbles might have a legitimate complaint against her parents for being born with a disability that prevents her from

“being able to live well,” despite the fact that she is no worse off than she otherwise would have been for having been created with a disability.⁸⁷ Like Woodward, Smolkin concocts cases in which a person can have her rights violated, and thus have a legitimate complaint against someone, despite the fact that the violation fails to make her worse off, all things considered, than she otherwise would have been. He asks us to consider, for example, a case of a doctor who lies to her patient about the gravity of his condition in order to prevent the patient from suffering the harm associated with anticipating death. Many think the doctor has wronged the patient by violating his rights (perhaps the right to full disclosure on the part of one’s doctor), despite the fact that the patient is better off as a result of the lie. Now consider the case of a non-voter whose ability to vote is unjustly taken away. She never would have exercised her ability to vote, but most think that she is nonetheless wronged by being deprived of the opportunity. Her rights have been violated, and she has a legitimate complaint, even though she is no worse off than she otherwise would have been. Smolkin believes that these examples give sufficient reason to “deny that there is anything especially problematic about claiming that the future persons in Non-identity Problem cases may have legitimate complaints that acts necessary for their coming into being with lives worth living, though handicapped, wronged them.”⁸⁸

This last inference is a bit quick. Smolkin is correct to point out that there is nothing problematic about claiming that an act that fails to harm someone might nonetheless wrong him. We have already seen several such cases in discussing Woodward’s proposal. However, this alone is insufficient to show that there is *nothing* problematic about claiming that Pebbles is

⁸⁷ Smolkin (1994) and (1999).

⁸⁸ Smolkin (1999: 202)

wronged despite the fact that she is not harmed. In both of the cases to which Smolkin appeals, there is some existing person who is either *deceived* or has an opportunity *taken away* from her. Furthermore, in the cases to which Smolkin appeals, someone *exists* and, in virtue of existing, has a right at the time that is violated by some action that fails to make her worse off. Neither of these conditions obtains in Non-identity cases like Wilma's. Pebbles is not deceived, and she does not have any opportunities taken away from her as a result of Wilma's choice. Most importantly, Pebbles does not yet exist when the act is performed that causes her to exist with a life that is "in some serious sense, defective."⁸⁹ One might attempt to fix this disanalogy by attributing actual rights to merely potential people, but doing so seems to raise more problems than it solves. For example, by attributing actual rights to potential people, it becomes difficult to avoid the implication that the sighted child that Wilma could have conceived is also wronged when she conceives Pebbles.⁹⁰

There is a further problem of specifying the nature of the right that Pebbles has that is violated when she is caused to exist with a life that is, in some serious sense, defective. Remember that Pebbles' alternative is non-existence. If Pebbles has a right not to exist with a defective life, then it must be the case that she has the right not to exist with a defective life when the alternative is *not to exist at all*. But if this is the nature of Pebbles' right, then many of our medical practices are much more morally problematic than we supposed. Consider a person who is involved in a terrible car accident. The unconscious woman is laid on the operating table. She can be saved, but only by amputating both of her legs. If Pebbles' life is in some serious sense

⁸⁹ I borrow the phrase from Smolkin.

⁹⁰ Boonin (2008: 140)

defective due to her blindness, then presumably a life without the use of one's legs is also defective.⁹¹ The alternative in this case is allowing her life to end. If individuals have a right not to exist with a life that is, in some serious sense, defective when the alternative is non-existence, then it would seem that the patient is wronged if her legs are amputated in order to prevent her death.⁹² Surely this implausible implication gives us strong reason to reject Smolkin's proposal.

Smolkin might respond that his solution only relies on the idea that someone has a right not to *be brought into existence* with a life that is in some serious sense defective. Since the patient is not brought into existence with a defective life by having her legs amputated, *this* right is not violated, so it is not wrong to perform the surgery. Pebbles, on the other hand, is brought into existence with a life that is defective, so her right not to be brought into existence with a life that is defective *is* violated by Wilma's action. This distinction is unacceptably arbitrary and so will not save Smolkin's proposal. It seems that what is problematic, if anything, about *bringing* someone into existence with a defective life is the fact that it results in someone *living a defective life*. Causing someone to *continue* to exist with a defective life also results in someone living a defective life. Thus, if one has a right before she exists not to be brought into existence with a defective life it would be surprising, once she exists, if she did not have a corresponding right not to continue to live a defective life. The defective life is what matters, not the coming into existence.

⁹¹ If a life with no legs is not limited enough for the reader to judge it as seriously defective, then she is free to imagine a worse fate that elicits the judgment. Remember that we are not to imagine some condition so bad that it render's one's life not worth living.

⁹² If the special duties entailed by the doctor/patient relationship complicate the picture, then imagine that a surgically skilled bystander comes upon the victim at the scene of the accident. For this example to work it must be the case that existing individuals have *at least* all of the rights that merely potential people. This seems like a relatively uncontroversial assumption.

Peter Markie offers a rights-based solution to the Non-identity Problem by drawing an analogy between cases in which an existing child is caused to have a significant handicap, and cases like Wilma and Fred in which an additional child is caused to exist with a significant handicap. He thinks that the wrongness of the action in the latter type of case can be explained in terms of the wrongness of the action in the former type of case. Wilma acts wrongly in creating Pebbles because, though she does not harm her, she wrongs her by causing her to have a significant handicap, just in the way she would wrong an existing child of hers by causing her to have a significant handicap. Below is his explanation:

In Postconception/Wrongful Disability, the mother's negligent act occurs during pregnancy, when she needlessly takes a drug she knows is likely to cause a severe disability in her child, which it in fact does. Her child's opportunities are substantially limited by his disability. In this case, the immorality of the mother's conduct has a straight-forward explanation. She harms her child. She violates his rights in a way that makes him worse off than he would have been otherwise. Exactly which of her son's rights is violated is open to interpretation, but a plausible candidate is the right he has once he exists as a person not to be caused to have a restricted life, a life seriously deficient in one or more of the major respects that make human lives worth living... The mother's conduct in Wrongful Conception/Wrongful Disability [Wilma and Fred cases] has exactly the same consequences for her child as her conduct in Postconception/Wrongful

Disability. In each, she causes him to be severely disabled but with a life worth living. Since her conduct has the same consequences for him in each case, and she clearly violates his rights, once he exists as a person, in Postconception/Wrongful Disability, she also violates his rights in Wrongful Conception/Wrongful Disability. She wrongs him after all.⁹³

The line of argument in this passage appeals to something like the following principle.

Consequence: If two actions have the same actual consequences for an individual then if one act violates the individual's rights (wrongs him) then the other act violates the individual's rights (wrongs him).

The problem for Markie's proposal is that Consequence is false. There are two plausible but distinct principles each of which we must reject if we are to accept Consequence:

Alternate Possibilities (AP): Whether an action wrongs A depends on how the value of the outcome *o* it produces (for A) compares with the value of alternative outcomes {~*o*} of distinct actions that might have been performed.

⁹³ Markie (2005: 296-297).

Alternate Value Possibilities (AVP): Whether an action wrongs A depends on the all-things-considered-value of the resulting outcome o for A. The all-things-considered value (for A) of an outcome o is determined by the difference between how well A's life goes given o compared with how well A's life would go in $\sim o$, where $\sim o$ is the relevant available alternative outcome.⁹⁴

AP and AVP are two ways of stating the same idea. If either of them is true then Consequence is false. They are clearly true. Determining whether an action wrongs a person clearly depends, not only on the actual consequences of the action, but on the available alternatives. Consider two cases:

Exploding Car: A is stuck in a car that is about to explode. B rushes to his aid, and in hurriedly unclasping his seatbelt, breaks his arm. If B had not done so, then A would have been killed by the explosion.

⁹⁴ The idea here is that when we are calculating the value of an outcome we should not just sum up the intrinsic value of the goods that obtain in the outcome. We should sum up the basic goods and then subtract the sum of the intrinsic value of the goods that would have obtained had some *relevant* alternative outcome occurred instead. Once this operation has been performed you have the all-things-considered value of the actual outcome.

Jammed Seatbelt: A is stuck in a car that is parked on the side of the street. B rushes to his aid, and, in undoing his seatbelt, breaks his arm. If B had not done so, then A would have been stuck in his car for another 2 minutes until he was able to undo the seatbelt.

It is clear to most that A is wronged in *Jammed Seatbelt* but not in *Exploding Car* despite the fact that the actual “consequences” for A (in Markie’s sense)⁹⁵ are the same in each case. A is alive with a broken arm. Because the actual outcome for A is the same in each case, the *alternative outcomes* available to the agent must help determine whether the action in each case wrongs A. In *Exploding Car* the alternative outcome is much worse than the actual outcome, whereas in *Jammed Seatbelt* the alternative outcome is much better. This is evidence that we are committed to either AP or AVP. It would be absurd to claim that A is wronged in both cases simply because in each case A *ends up alive with a broken arm*. The success of Markie’s solution hinges on the truth of Consequence, but Consequence is clearly false. Though the actual consequences for the Conception/Wrongful Disability child are the same as the actual consequences for the Postconception/Wrongful Disability child, this is only *part* of the story in determining whether one or both children are wronged. Once we take into account the relevance of counterfactuals in determining whether someone has been harmed, Markie’s solution loses its appeal. There is no longer any reason to think that the fact that the negligent, already pregnant mother wrongs her child has any bearing on whether Wilma wrongs Pebbles.

⁹⁵ If we have a broad enough understanding of the notion of an action’s “consequences”, then we can include among those consequences the counterfactual outcomes that might have obtained had the actual outcome not obtained. Of course, on this understanding of “consequences,” the consequences for the child in Conception/Wrongful Disability are *not* the same as the consequences for the child in Postconception/Wrongful Disability.

III.1.3 Tooley's Solution: Future Rights Violations Solution

Michael Tooley has argued that the wrongness of Wilma and Fred's act of conceiving Pebbles is analogous to the wrongness of making a promise that you know you will be unable to keep. When you make a promise you know you will be unable to keep you do something wrong, because you cause it to be the case that the person will later have *more* rights violated than she otherwise would have. Similarly, in the case of Pebbles, by causing her to exist blind, Wilma and Fred cause it to be the case that Pebbles will suffer certain rights violations in the future. Perhaps she will endure hardships that she has a right not to endure.⁹⁶ It is easy to misconstrue this view as an absolutist one on which the *mere fact* that Pebbles will have a right violated is sufficient to make Wilma's act of conceiving her impermissible. If this were the view, it would be easy to dismiss since it would seem to entail that virtually all acts of procreation are impermissible. If conceiving Pebbles is wrong because it causes it to be the case that she will endure future rights violations, then almost *all* possible acts of conception will turn out to be wrong. After all, everyone has a right violated at some point in one's life, and that right would not have been violated if the person had not been conceived.

A proponent of this solution might respond that in certain cases, an agent can act wrongly by ensuring that *the agent herself* will later violate someone's rights. For instance, it might be wrong to make a promise that one knows she will not keep. Earlier I considered a nice example of this kind of rights violation from Woodward:

⁹⁶ Tooley, (1985: 270-273) offers this solution. Woodward offers a similar view in his (1986).

Promise: You promise your friend that you will perform a conjunctive action AB.

You could have either made this promise or not. You knew when you made it that if you promised to do AB you would only do A, but that your friend would be better off as a result of you doing A than she would have been had you not promised to do AB and thereby failed to do either A or B.

In this case you are uniquely responsible for the rights violation that your friend later suffers. There is no intervening agent to whom we can attribute moral responsibility, so you are the only candidate morally responsible agent. This might be sufficient to render your *act of promising* wrong in *itself*. Let's assume for a moment that Pebbles, when she exists, has a right to fully functional eyesight. According to the view under consideration, Wilma will be uniquely responsible for the fact that Pebbles' right to a life with full sight is violated if she is conceived. On the other hand, the fact that Pebbles will have her right violated by being lied to by *someone else*, at some point in the future, does not count against Wilma's choice to conceive her. In that case, Wilma is not uniquely responsible for Pebbles' right being violated. There is an intervening agent to whom we can attribute primary moral responsibility. The trouble with this reply is that, presumably, many parents will violate their children's rights in some minimal way *at some point* (perhaps by eaves-dropping on a phone-conversation with a boyfriend) in the future. Surely these kinds of violations alone do not make it impermissible to conceive those children.⁹⁷

Perhaps this is not a very charitable interpretation of the "future rights violations" solution. A more plausible interpretation would incorporate insights from the Alternate

⁹⁷ In conversation, Michael Tooley has described his view this way.

Possibilities or Alternate Value Possibilities principles discussed in II.1.3. The result would look something like the following:

Minimizing Rights Violations (MRV): An action is right in case its resulting outcome involves the fewest rights violations of the available alternatives, other things being equal.

Unlike the first formulation of the future rights violation solution, MRV allows the proponent of the future rights violations solution to avoid the implication that it is almost always wrong to create a person. According to MRV, the fact that an act will bring about a rights violation *counts against* performing an action, but that fact alone is not sufficient to render it impermissible. We must also compare the extent and number of rights violations that result from that action with the alternatives. The fact that creating a person would result in that person suffering a rights violation does not always mean that it would be wrong to perform the act. If that person would suffer the fewest rights violations of all of the people who could have been created, then the act is permissible, according to MRV. The ‘other things being equal clause’ allows the proponent of the view to include other values in the set of morally relevant considerations for determining the permissibility of an act of creating someone. For instance, the defender of MRV can include welfare among the morally relevant values. Imagine that you can produce child A or child B. Child A will suffer more rights violations than child B, but he will be happier over the course of his life. On the sophisticated version of the future rights violations

view it may turn out that is permissible to create A rather than B. This version of the view also entails that Wilma acts wrongly in creating Pebbles; the fact that Pebbles suffers more rights violations than Rocks would have suffered morally counts against creating Pebbles rather than Rocks. Because there are no other factors (such as welfare) that override the consideration of rights, MRV entails that Wilma acts wrongly by creating Pebbles rather than Rocks.

This version of the future rights violations view has a curious manner of explaining why Wilma acts wrongly by conceiving Pebbles. We can see this curiosity if we contrast the rights of Pebbles that will be violated if she is conceived with the rights of the sighted child that would *not* be violated if he were to be conceived. The right of Pebbles that will be violated must be something like a right to a *particular quality of life*. As I pointed out in the previous chapter using the case of George the cave dweller, there is nothing bad about being blind *per se*. As such it would seem strained to insist that we have rights *against being blind* or a corresponding right *to sight*. Blindness is bad only insofar as it is related in the appropriate ways to *how well someone's life goes*. Let's assume, then, that Pebbles and the sighted child each would have an equal right⁹⁸ to a particular *quality of life* were either of them to come to exist. The formulation of the future rights violation solution under consideration must hold that this right is violated to a greater extent in the case where Pebbles exists than it is in the case where the sighted child exists. *This* fact explains why Wilma's choice to conceive Pebbles is wrong.

This last feature of the future rights violation solution means that the solution faces a problem that is similar to one faced by Shiffrin's non-comparative harm-based solution,

⁹⁸ If one thinks that the strength of one's claim to a particular level of welfare depends in part on what she *deserves* then I will hold this fixed. The two children would be equally *deserving* of a particular quality of life.

discussed in *II.3.2.A*. Remember that on Shiffrin's view an act is harmful if it causes someone to suffer a cleavage between what she rationally wills and her circumstances. The trouble with this account was that it seems difficult to believe that a person could 'rationally will' that she exist under circumstances that are nomologically impossible. Doing so would be tantamount to wishing for her non-existence. But Shiffrin's solution to the Non-identity Problem is successful only if one could rationally will that one exist under nomologically *impossible* circumstances, given the way the world has unfolded. But, I claimed, wishing for such circumstances is clearly *not* rational, so Shiffrin's solution fails. The future rights violations solution under consideration must also deny that nomological possibility constrains what we have moral rights to.

Just as it seems that for someone to *rationally* will that her circumstances be different, there must be some available world in which they are different (and better) than her actual circumstances, it seems that when a person has an unsatisfied right to a particular quality of life, there must be nomologically possible (and counterfactually close) circumstances in which things go better for her than in the actual circumstances. However, given that Wilma has the temporary condition that will cause whatever child she has to go blind, blindness is a necessary condition for Pebbles' existence. All of the circumstances in which Pebbles is not blind, given the history of the world up to the point of Wilma's decision, are circumstances in which Pebbles does not exist. Thus, assuming that Pebbles' life is at least not worse than non-existence, it is hard to accept the implication that she has a *right* go unsatisfied in virtue of the fact that she was born blind. In ordinary cases it seems that when someone's right is violated she has a legitimate complaint against the person who made it the case that her right was violated. But what *complaint* does Pebbles have against Wilma? She can't complain that Wilma ought to have conceived *Pebbles* as a sighted child. Wilma couldn't have. In order for it to be the case that

Pebbles has a right to a higher level of welfare than the level of welfare she actually has, it must also be the case that she has a right that is violated by the obtaining of certain necessary conditions of her existing with a life worth living. This is implausible.

We might worry about the further implications entailed by my claim that Pebbles cannot have a right violated by the obtaining of certain necessary conditions for her existing with a life worth living. Consider a case where someone has a miserable life that is the only life she could have had. Does this person have a right to a higher quality of life even though there is no outcome in which she exists and has a life worth living? Given that a non-miserable life is nomologically impossible for her, I do not think this person has a right to a life worth living, but this does not mean that she does not have any right violated by being caused to exist. There is, after all, an outcome in which this person does not live a life of terrible suffering. It is the outcome in which she does not exist, and it is the outcome that is *preferable* from the standpoint of the person's interests to the actual outcome. So it seems plausible to think that this person has a right *not* to live a life of terrible suffering and that *this* right is violated when she is caused to exist.

There is a second problem for the future rights violation solution if we think that rights are tied to duties in a particular way. Normally we think that if someone has a right to something there must be an agent who has a duty to satisfy the right. If I have a right against having my computer stolen, then someone (everyone) has a duty, say, not to take it from me. If I have a right to healthcare then someone has a duty to provide me with it. But if Pebbles has a right to sightedness, then we must reject this intuitive view about the relationship between rights and duties. In her case, no one is able to satisfy her right to sightedness, because (i) her blindness is

incurable and (ii) there is no nomologically possible world in which *she* comes to exist as a sighted person. But it seems that if it is nomologically impossible for anyone to satisfy her claim then, supposing that ought implies can, no one has a duty to satisfy it. This forces the proponent of the future rights violations solution to reject the following principle about the relationship between rights and duties:

Right-Duty 1: If S has a right to y, then there is some actual agent A such that A has a duty to satisfy the right to y.

This formulation of the principle might seem problematic. There are cases where we seem to have rights to something even though no one has a duty to satisfy the right. For instance, in a world where I buy my computer from someone and then everyone except for me immediately blinks out of existence, I still have a right to my computer. No one else exists, so no one has a duty not to take my computer from me. Nevertheless, I still have a right to the computer. Perhaps the principle should be reformulated to accommodate this kind of case:

Right-Duty 2: If S has a right to y, then, in a world with at least one an agent A, A has a duty to satisfy y.

This statement of the view retains the intuitive idea that rights entail duties but it accommodates cases where there are no agents. It also entails that Pebbles does not have a right to being sighted. Given that Pebbles exists, no agent can make it the case that she is not blind. If ought implies can, then no agent has a duty to cause her to be sighted, given the way that events have unfolded. Thus, if Pebbles has a right to being sighted, this must be a right that no one who exists has a duty to satisfy (since they can't). Thus, even on the second, weaker version of the principle, Pebbles does *not* have a right to sight.

The defender of the future rights violations position might reply by claiming that Wilma has a duty to satisfy Pebbles' right to sightedness even though, once Pebbles exists, there is no way for Wilma (or anyone else) to satisfy this duty. Perhaps there are other cases that show that one can have a duty that she cannot fulfill without rejecting the doctrine that ought implies can. Imagine, for instance, that a mugger assaults a person. The victim suffers financial, physical and emotional damage to the tune of fifty thousand dollars, and so is owed that much compensation by the mugger. But the victim has also developed a victim complex such that any compensation that he is paid by the mugger will only make him all-things-considered worse-off by reinforcing his feeling that he is a victim. It seems that the mugger has a duty to compensate the victim even though it is now impossible for him to do so. The mugger's duty doesn't vanish simply because he cannot now fulfill it. This explanation of the mugger case does not seem to violate the doctrine of ought implies can, because the mugger *could have* avoided incurring the obligation in the first place by refraining from mugging the victim. Similarly, we might think that Wilma still has a duty to fulfill Pebbles' right to sightedness even though it is nomologically impossible for her to do so. Like the mugger, Wilma could have avoided doing something in the past that caused her to have a duty now that she cannot now fulfill.

There is an important difference between the mugger case and Wilma, however. The mugger has a duty that he can now not fulfill because he previously made his victim *worse off than he otherwise would have been*. But Wilma does not make Pebbles worse off than she otherwise would have been. Hence the mugger case demonstrates only that an agent can have a duty to someone that she cannot later fulfill in cases where the agent (i) has previously caused someone to be worse off than she otherwise would have been and (ii) cannot now compensate the victim for that harm. This does *not* show that an agent can have a duty to someone that she cannot now fulfill in cases where the agent has *not* at any point made anyone worse off than she otherwise would have been. But this is what must be shown in order to provide a reason for thinking that Wilma has a duty to satisfy Pebbles' right to sightedness even though she cannot now satisfy that right.

Yet a further trouble arises for the future rights violations solution. The plausibility of the view is motivated by cases in which it appears to be wrong to make a promise *now* because one thereby ensures that she will be unable to keep the promise in the future. According to the future rights violation solution, the best explanation of the wrongness of false promises is that, by making a false promise, one ensures that someone will suffer rights violations in the future. According to a plausible alternative explanation of the wrongness of false promises, however, making a promise to someone that one knows one cannot and will not keep violates that person's rights *then and there*. This explanation seems just as plausible as the first, but it will not support the solution offered here without some questionable amendments. First of all, if the rights violation occurs at the time the promise is made then the person who is wronged must *have a*

right at the time the violation occurs.⁹⁹ But Pebbles does not exist when Wilma chooses to conceive her, so if Wilma violates her right then, Pebbles must have the relevant rights *before* she exists. But, as was noted when considering Markie's rights-based solution in section III.1.3, it is implausible to attribute actual rights to people who do not yet exist. Doing so embroils us in metaphysical and moral problems. I have tried to show that the future rights violations solution has highly problematic implications. The solution is successful only if we can have rights to goods that no one could give us. This seems to misconstrue the relationship between rights and duties. The solution also seems to flout the view that rights entail duties. Even if these troubling implications were avoided, I have shown that there is a second way of explaining the wrongness of false promises that undercuts the initial motivation for the future rights violations solution.

III.1.4 Reiman's Solution: The Appeal to Fairness

Jeffrey Reiman argues that Pebbles has her rights violated by being conceived with a disability and that this rights violation is grounded in the hypothetical choices of individuals in a Rawlsian Original Position.

Because they do not know their generation, the parties in the original position, in effect, represent all and only those people who, from this moment on, will ever exist: people who are currently living, and future people who do not yet exist but who one day will... I shall assume that

⁹⁹ Boonin (2008: 141-142) offers this objection to the Future Rights Violations Solution.

parties in the original position would agree on a principle providing that currently living people have a duty to do what they reasonably can, subject to their needs, rights, and other moral duties, to ensure that future people are capable of normal functioning and live lives with normal life expectancies and normal morbidity rates. If this is current people's duty to future people, then future people have a correlative right to those efforts on the part of currently living people....If we ask whether conceiving a child with a disability when we could have conceived a different normal child later would be acceptable, the answer would have to be that we had failed in our moral duty to future people. We had not given one of them what future people have a right to expect. Thus, in Preconception Wrongful Disability, bringing a person into existence with a disability would not only be wrong, but a wrong to the future person thus conceived, a violation of that person's right to his forebears' efforts to ensure normal functioning.¹⁰⁰

On Reiman's account, the choices of the members of the Original Position guarantee that Wilma wrongs Pebbles when she conceives her. Wilma wrongs Pebbles by violating her right to "her forebears' efforts to ensure normal functioning." Strangely, this explanation of the wrongness of Wilma and Fred's action seems to be unavailable given Reiman's description of the Original Position. According to him, the OP represents "all and only those people who, from this moment on, **will** ever exist: people who are currently living, and future people who do not

¹⁰⁰ Reiman (2007: 78-82).

yet exist but who one day **will**.” He then proceeds to assume that the individuals in the OP would agree “on a principle providing that currently living people have a duty to do what they reasonably can, subject to their needs, rights, and other moral duties, to ensure that future people are capable of normal functioning and live lives with normal life expectancies and normal morbidity rates.” Let’s call this principle to which the people in the OP would all agree “Future”. Reiman thinks that Future can be derived from the OP and can explain why it is wrong to conceive a child with a disability when one could have conceived a numerically distinct normal child later. Future, as stated, has an ambiguous description. “Future people”, as it appears in the principle, might be interpreted in either a broad way or a narrow way:

Broad Future (BF): Currently living people have a duty to do what they reasonably can, subject to their needs, rights, and other moral duties, to ensure that future people are capable of normal functioning and live lives with normal life expectancies and normal morbidity rates. “Future People” refers to *all possible future people*. The duty to *all possible future people* might require currently living people to alter their behavior so as to ensure that those who come to exist are better off than *alternative possible people* would have been who might have come to exist.

Narrow Future (NF): Currently living people have a duty to do what they reasonably can, subject to their needs, rights, and other moral duties, to ensure that future people are capable of normal functioning and live lives with normal life expectancies and normal

morbidity rates. “Future people” refers to the *people who will actually come to exist*. As such, the principle only requires that currently existing people do what they can to ensure that those people who actually come to exist, once they exist, are capable of normal functioning, etc.¹⁰¹

NF cannot solve the Non-identity Problem. On NF, either Pebbles or Rocks is in the OP, but not both. Let’s suppose that Pebbles is one of the people who *will actually* come to exist, perhaps because Wilma has already made her choice. Narrow Future states that currently living people have a duty to do what they can to ensure that Pebbles has a life of normal functioning, but it does not entail that Wilma has an obligation to conceive Rocks *rather than* Pebbles (since we are assuming that the choice has already been made). Moreover, Wilma cannot do anything to ensure that *Pebbles* is capable of normal functioning before or after she comes to exist (since a necessary condition for her existence is that she exist blindly). She can maybe take measures to ensure that Pebbles’ blindness has only a minor negative impact on the quality of Pebbles’ life, but she cannot ensure that Pebbles is sighted. As long as she takes these measures, she is not in violation of NF. There is nothing to be done either before or after Pebbles exists that will ensure that she is capable of normal functioning, so NF does not entail that Wilma’s *choice* to conceive Pebbles is wrong. But to solve the Non-identity Problem we must explain why Wilma’s *choice* to conceive Pebbles is wrong. If the OP is to solve the Non-identity Problem then it must be the

¹⁰¹ Notice that these two interpretations of Future are eerily similar to the distinction between *de re* and *de dicto* duties described by Hare. Since Reiman needs Broad to be the correct interpretation of Future to solve the NIP, and since Broad is the *de dicto* reading of Future, his solution likely falls prey to all of the same problems as Hare’s account.

case that the people in the OP would agree to BF rather than NF. Let's turn to BF and ask – would the OPers agree to it?

When considering the prospects of BF for solving the Non-identity Problem it is important first to note that Reiman stipulates that the OPers represent all and only those people who *will* ever exist. This raises an immediate puzzle. If the identities of the OPers are fixed, then how can they recommend a policy that might change *who comes to exist*? Perhaps we can reconcile this apparent tension by picturing the OP as follows: as things stand, there is one group of individuals who occupy the OP. These people are *possible people* insofar as they are the ones who will come to exist in the absence of any intervention. The people in the OP might not exist, however, if some intervention occurs. So, there must be some distinct set of possible people who are not currently members of the OP but who may become members of the OP if the course of nature is disrupted by some intervention. On this description of the OP, there are many possible people, only some of whom are members of the OP, and it is possible for the current OPers to choose a policy that would replace them with some group of possible people *waiting in the wings*, as it were. Of course, it would be surprising if all and only individuals who are currently members of the OP were *also* all and only the people who would enjoy lives with normal functioning, normal life expectancies, and normal morbidity rates. It seems, then, that if the current OPers chose BF this would almost certainly preclude the existence of some of them. But if the OPers were to agree to principles that preclude their existence, then, according to Reiman's own description of the OP, they would no longer be members of the OP, and their choices would be irrelevant (since, according to Reiman, those who will never exist “do not have interests”). I suspect that Reiman would want to avoid this complication by holding the identities of the members of the OP fixed (to those people who *will exist*, whoever they turn out to be). But then

we cannot employ the OP to select principles governing identity-affecting actions, adherence to which would preclude the OPers' very existence. Hence it seems that the members of the OP can only choose procreative principles that are restricted so as to range over actions that do not determine who will exist. In other words, they must choose NF. But, we have seen, NF seems incapable of explaining why it is wrong to conceive a child with a disability when one could have conceived a different normal child later. Identity-affecting actions are outside of its domain. It is perfectly permissible, according to NF, to create the worse-off of two possible people and then ensure that *that particular person* is as well-off as possible. NF only requires that we do what we reasonably can to guarantee a certain quality of life for those people who *will* exist, once they exist. It does not require that we ensure that the people who will exist are better off than the alternative people who might have existed.

Reiman responds to this worry by challenging the supposition that the future person, represented in the OP, must face a choice between existing or not. He does so by providing the following, curious, claim about the identities of the future people represented in the OP:

From a future person's standpoint, it makes sense to think that it is in his or her interest to be born with certain properties rather than others, but it is not in his or her interest to be born this particular one rather than that particular one *independent of differences in properties*.¹⁰²

¹⁰² Reiman (2007: 84)

But if being “this particular one rather than that particular one” just is to be some numerically distinct person rather than some other numerically distinct person, then it seems as if Reiman is claiming that one and the same future OP person might turn out to be *either* the disabled child *or* the normal child. The future person, represented in the OP, is numerically identical to whichever one comes to exist. But, if this is Reiman’s claim, then we get an egregious violation of transitivity of identity. This implication can be brought out by considering Wilma’s choice between conceiving Pebbles and Rocks. First suppose that Pebbles is not numerically identical with Rocks. This is not a problematic supposition, and I have already defended it in chapter I. Furthermore, Reiman must agree with this supposition, for if it were mistaken, then we would not require the use of his OP to solve the Non-identity Problem. We could simply solve it by maintaining that Pebbles is made worse off by being caused to exist blind rather than sighted. With this supposition in mind recall that, according to Reiman, the future person in the OP is identical to both Pebbles and Rocks (since, as Reiman says, he might “be born this particular or that particular”). So we get Pebbles=Future Person, and Future Person=Rocks, yet we must deny that Pebbles=Rocks. This is a violation of transitivity of identity, so we must reject Reiman’s supposition that a future person can be identical to (the same particular as) both Pebbles and Rocks. The future person is either identical to one or to the other but not both. But if the future person can be identical to either Pebbles or Rocks, but not both, then the OP cannot be used to choose principles governing identity-affecting actions.

III.2 The Failure-Based Approach

In this chapter so far I have shown that harm-based solutions and rights-based solutions to the Non-identity Problem are unsuccessful. Some might at this point wonder if all candidate solutions are bound to fail. If all of the candidate solutions fail, then we seem to be left with no choice but to accept the conclusion of NIPA. The Failure-based approach takes a negative approach to biting the bullet of the Non-identity Problem. David Boonin employs the Failure-based approach by raising objections to the most popular solutions to the Non-identity Problem, and then offering an error-theory to explain why we mistakenly believe that Wilma acts wrongly by conceiving Pebbles.¹⁰³ He concludes that NIPA is a sound argument for a true conclusion. Because this approach is successful only if my proposal in the following two chapters fails, I will not consider it in greater detail.

III.3 Rejecting P5: Doing Wrong without Wronging: Non-person-affecting Solutions

Setting the failure-based approach aside, I turn to some candidate solutions that begin to resemble my proposal in chapters IV and V. Derek Parfit, one of the first to explain the Non-identity Problem and to gesture at a solution to it, was sympathetic to the idea that an agent might act wrongly in Non-identity cases by failing to produce the more valuable of the outcomes available to her.¹⁰⁴ In other words, Parfit has seemed sympathetic to solutions that reject P5 of

¹⁰³ Boonin (2008)

¹⁰⁴ I will discuss Parfit's views on the Non-identity Problem in more detail in a later section. I would like to think that the solution to the Non-identity Problem that I will propose might be what Parfit had hoped for when he challenged the philosophical community to come up with what he called "Theory X"; a fully satisfactory solution to

NIPA. Parfit proposes that a principle he calls “The Same Number Quality Claim” or Principle Q might suffice to solve the Non-identity Problem.

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.”

We might use this principle to solve the Non-identity Problem by pointing out that Wilma chooses the worse of the available outcomes by conceiving Pebbles rather than Rocks. When formulated as a principle about wrongness, Parfit’s principle Q does seem to entail that Wilma acts wrongly by conceiving Pebbles rather than Rocks:

Q+: If in either of two outcomes the same number of people would ever live, it would be morally wrong to choose the outcome in which those who live are worse off, or have a lower quality of life, than those who would have lived.¹⁰⁵

Brock has defended a similar principle to Q+ that he calls “Principle N”:¹⁰⁶

the Non-identity Problem that avoids the metaphysical problems with person-affecting solutions as well as the demandingness problems of full-blown Consequentialist solutions.

¹⁰⁵ Boonin (manuscript) restates Q in roughly this way so that it explicitly solves the Non-identity Problem.

N: Individuals are morally required not to let any possible child or other dependent person for whose welfare they are responsible experience serious suffering or limited opportunity if they can act so that, without imposing substantial burdens or costs on themselves or others, any alternative possible child or other dependent person for whose welfare they would be responsible will not experience serious suffering or limited opportunity.

Q+ or N *would* solve the Non-identity Problem if accepted. Wilma certainly violates both principles by conceiving Pebbles rather than Rocks. I do not think that Q+ or N are, in and of themselves, satisfactory proposals however. Q+ and N satisfy the Intuitive Cost and Sufficient Generality Criteria for a successful solution; they do not have obviously problematic implications, and they solve both the immediate and eventual versions of the problem. But they satisfy the first two criteria *only* because they are *ad hoc* principles designed *specifically* to solve the Non-identity Problem. They do not provide independent *justification* of our judgments about the Non-identity Problem. Rather, they appear to be mere assertions of those judgments. There is no reason to accept Q+ or N other than the fact that doing so enables us to solve the Non-identity Problem. This means that Q+ or N fail to satisfy the Non-*ad hocness* Criterion for a successful solution.

¹⁰⁶ Brock (1995). For a slightly modified version of N, see also Buchanan, Brock, Daniels and Wikler (2000: 249). See also Brock (2005: 79-87). Glover (2001: 444), Peters (2009: 325), and Steinbock (2009: 172, 174; 2011: 90) seem to offer similar proposals.

The solution that I will propose in the following two chapters entails some version of Q+ or N, but, unlike these principles, it is not *ad hoc*. My proposal is *independently* plausible, even if we ignore its implications about the Non-identity Problem. Like Q+ and N, my solution falls into the category of “value-based solutions,” so, before defending my proposal, it is especially important that I explain some competing proposals that also fall into this category and that I make clear what I take to be inadequate about them.

III.3.1 Kavka’s Solution: Making the World Worse by Creating People with Limited Lives

Gregory Kavka can be interpreted as suggesting that Wilma makes the world all-things-considered worse by choosing to create Pebbles, and that this fact explains the wrongness of her action. In places, he seems to argue that there is intrinsic value (disvalue) in certain states of the world and that this value is not reducible to the welfare (lack of welfare) of the individuals contained therein:

One approach to evaluating the desirability [intrinsic value] of states of society seems especially promising, in the present context. Let us introduce the notion of a restricted life, a life that is significantly deficient in one or more of the major respects that generally make human lives valuable and worth living. Thus, the life of a slave is restricted in this sense, owing to the slave's lack of liberty. Clearly, the lives of persons significantly handicapped, either mentally or physically, from birth and of those struck down in the prime of life by illnesses caused by radioactivity

are also restricted. So, for a variety of reasons, are the lives of many living in a very overcrowded world.¹⁰⁷

If we accept that “restricted lives” can contribute intrinsic undesirability or disvalue to a state of affairs, then we might go on to adopt the following principle:

Restricted Life: ...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.¹⁰⁸

Kavka maintains that something like Restricted Life allows us to get the right verdicts in Non-identity cases without entailing that we have duties to produce the happiest child possible.¹⁰⁹ Here is his (rather long) explanation:

Then we can derive the intuitively correct verdicts in the cases involving the slave child, nuclear power, the pleasure pill, and controlled growth. For, in each case, the morally wrong act can be condemned on the grounds that it foreseeably leads to (states of the world containing) restricted lives. At the same time, the stated principle is consistent with our belief that prospective

¹⁰⁷ Kavka, p. 105 (1982)

¹⁰⁸ Ibid. p. 106

¹⁰⁹ He worries that this is a problematic implication of any view that appeals to the moral reason to promote welfare in order to solve the Non-identity Problem. I will address this worry as it pertains to my view below.

parents have no obligation to produce the "best" child they can. And it gains plausibility from its affinity with a frequently invoked canon of social justice: that high priority should be given to providing all with the means to live at least a minimally decent life and engage in the major activities of human and community living. Our principle says to prevent restricted lives, this canon says to prevent lives from being restricted. Both seem to derive from the common intuition that there is some-thing seriously wrong with people living restricted lives, which makes it incumbent upon others to stop this from happening if they can. Perhaps this intuition depends, in turn, on a fundamental belief in the dignity of human beings, conjoined with the observation that, given human nature, **many people living restricted lives are likely to be treated by others, or to treat themselves, with less than the full respect that they deserve.** In any case, it is an intuition that many of us share and that may serve as a provisional grounding for our principle concerning restricted lives.¹¹⁰

The above paragraph leaves it ambiguous what exactly is supposed to ground the Restricted Life principle. According to Kavka's original statement of the principle, Restricted Life seems to ground Wilma's moral obligation not to conceive Pebbles in the intrinsic disvalue that obtains in certain states of affairs involving individuals with severely restricted lives. The bolded sentence in the above paragraph, however, seems to indicate that something like the notion of *respect for persons* is doing the principle's moral work. I will consider both the value-

¹¹⁰ Kavka, p. 105-106 (1982)

based and the respect-based justifications of the principle. On each justification, Restricted Life turns out to be unacceptable.

Let's consider the value-based interpretation of Restricted Life first. On this interpretation, the existence of individuals with restricted lives is intrinsically disvaluable. In other words, they in some way contribute negative value to states of affairs. There are two ways of construing this view. On one way of construing it, individuals with restricted lives make the world intrinsically worse all-things-considered; the world would be intrinsically better if they didn't exist. This construal is highly counterintuitive and, dare I say, offensive. Surely it is at least not worse for a world to contain some people living happy, but nonetheless restricted, lives than it is for a world to contain no one at all. On a second, more moderate, construal of the value-based interpretation of Restricted Life, individuals with restricted lives only make the world intrinsically worse *qua* living restricted lives. They do not, however, always make the world *all-things-considered* worse. Their lives might contain other features that are highly valuable such as an overall positive level of well-being. These other features of their lives might outweigh the disvalue attached to them living restricted lives. Thus the value-based interpretation can avoid the highly counterintuitive implication that a world in which no one exists is better than a world full of people living happy but restricted lives.

Upon further consideration, however, even this second construal of the value-based interpretation of Restricted Life has implications that almost no one is willing to accept. If individuals with restricted lives make the world intrinsically worse insofar as they are living restricted lives, then the intrinsic disvalue of a person's restricted life must, unless it is to be rendered negligible, outweigh the positive value of, say, a person's welfare. In other words, we must accept that there are cases in which, even though a person would have a life worth living,

and even though no one else would come to exist in his stead, that it would be wrong to create that person because his existence would make the world intrinsically worse all-things-considered. This is very hard to swallow. If we have the choice between bringing someone into existence with life that is restricted but worth living and creating no one at all, then it is at least permissible to bring the person into existence with a restricted but worthwhile life. Suppose, for instance, that Wilma's choice is to conceive a blind child or to conceive no child at all. Surely most will agree that Wilma, in this scenario, is morally permitted to conceive a blind child, given that it is her only option.

A further problem with even the second, moderate construal of the value-based interpretation of the Restricted Life principle is that it seems to confuse instrumental with intrinsic value. Most will agree that an individual living a restricted life is *instrumentally* disvaluable insofar as it makes that person worse off than she otherwise would have been, but this does not imply that restricted lives are *intrinsically* disvaluable. To show that the moderate construal of the value-based interpretation of the Restricted Life Principle is conflating instrumental with intrinsic value in this way we must turn to a variant of a pair of cases that I considered in Chapter II.

Cave Dweller: George is a member of a cave-dwelling society. The members live a life of complete darkness. They go about their lives using heightened senses of hearing, touch, and smell. Though eyesight is useless to them, most of them nonetheless have the power of eyesight. George, on the other hand, has developed a rare degenerative disorder that has slowly made him blind. He never notices it, of course, since he always has lived and always will live, in complete darkness.

Cave Dweller 2: Jane is a member of a cave-dwelling society. The members live a life of complete darkness. They go about their lives using heightened senses of hearing, touch, and smell. Though eyesight is useless to them, most of them nonetheless have the power of eyesight. Jane has and always will have perfect vision. She will never use it, of course, since she always has lived, and always will live, in complete darkness.

According to the moderate construal of the value-based interpretation of the Restricted Life Principle, the world is worse if George exists than if Jane exists simply in virtue of the fact that George is blind. I suspect that this will seem incorrect to most. George's blindness in no way detracts from the quality of his life, since he has and always will go about his life in complete darkness. So how can the *mere fact* that he is blind mean that the world is worse than if he were sighted? Kavka might reply that George is not living what is properly described as a "restricted life." But if we recall Kavka's description of a restricted life as one that is "significantly deficient in one or more of the major respects that generally make human lives valuable and worth living," then I think we see that George does qualify as having a restricted life. His life is deficient with regard to his eyesight, and eyesight is one of the powers that *generally* contribute to a life that is valuable and worth living. It just so happens that eyesight is unimportant for George. Maybe Kavka would say that blindness does not cause someone to lead a restricted life in George's society because it does not generally make lives valuable and worth living *in that society*. This reply is similar to the one I offered on behalf of Hanser's event-based view in Chapter II. Like Hanser's reply, this reply seems to give up the game by conceding that the disvalue of restricted lives is *derived* from the effect of that restricted life on the welfare of the person who lives it. If

George's life is not restricted because of his blindness this must be due to the fact that it has no bearing on his welfare. It now seems like the "undesirability" of restricted lives is merely derived from the fact that they generally cause individuals to be worse off than they otherwise would have been. But this way of describing what is bad about restricted lives indicates that they are not intrinsically undesirable at all. A restricted life might be *instrumentally* less desirable than a non-restricted life insofar as a restricted life contains less of what is intrinsically valuable (welfare, for instance) than a non-restricted life would have contained. But this does not mean that the mere presence of restricted lives contributes intrinsic disvalue to a state of affairs.

On a second interpretation of the grounding of Restricted Life, it is wrong to create individuals with restricted lives because those individuals are unlikely to enjoy the *respect that they deserve*. I take this interpretation of Kavka's proposal to be equivalent to one of the versions of the rights-based solutions to the Non-identity Problem discussed in section III.1. It thereby inherits all of the problems that I raised for these solutions.

III.3.2 Temporal Indexing and Maximizing Consequentialism

Neil Levy proposes a solution to the Non-identity Problem that rejects P5 by appealing to temporal indexing in conjunction with Maximizing Consequentialist principles. The main component of the temporal indexing part of the view is the claim that we can adopt either a “forward-looking” or “backward-looking” attitude toward some action that was *expected* to have terrible consequences at the time the action was performed. When we adopt the forward-looking perspective we condemn the action, and when we adopt the backward-looking perspective we are grateful for it. The Non-identity Problem only arises, he claims, because we swing back and forth from one perspective to another. Levy offers the following analogy with the rationality of apology to illustrate the difference between forward-looking and backward-looking perspectives:

When we say that the slavery was wrong, when we apologize for it, our condemnation is temporally indexed. It makes essential reference to a particular, long past, moment in time, the time at which the choice was made to adopt the institution of slavery. Our condemnation therefore assumes and adopts this past temporal perspective, the perspective from which those responsible for slavery ought themselves to have grasped its wrongness. That our existence is causally dependent on these injustices is not relevant to our condemnation, or to our apology, for when we make it, the backward-looking perspective from our own present, the perspective we adopt when we are grateful for our own existence, is temporarily placed in parentheses. Just as, when we say that buying the lottery ticket was

irrational, we judge the action in the forward-looking terms which were, or ought to have been, available to a past agent, so when we condemn slavery or apologize for it, our judgment is temporally indexed.¹¹¹

Just as temporal indexing can explain why it can be both rational for one and the same person to be sorry for slavery *and* to be grateful that she exists as a result of it, Levy thinks that it can explain the paradox of the Non-identity Problem as well. He says:

The wrongness of the act is best understood from the forward-looking perspective of the agent, of Paula, for instance. Since her action is not the best of the alternatives available to her, she acts wrongly. The air of paradox arises when we assume the backward-looking perspective of Paul, whose existence is causally dependent upon Paula's act. Since he prefers existence to non-existence, we think he cannot condemn her act. We assume his perspective and conclude that no one's interests were harmed. But Paul, like us, can consistently condemn the act, without committing himself to the proposition that he would be better off dead. His condemnation, like ours, is temporally indexed, to the forward-looking perspective of Paula, whereas his joy at living is indexed to his present.¹¹²

¹¹¹ Levy (2002: 366).

¹¹² Ibid: 366

The curious thing about Levy's explanation here is that temporal indexing need not be doing any work whatsoever in explaining our (or Paul's) attitude toward Paula's action. Levy says, "since her action is not the best of the alternatives available to her, she acts wrongly."¹¹³ This is true regardless of the perspective we adopt. It is true if we adopt the backward-looking perspective and it is true if we adopt the forward-looking perspective. The further observation that Paul is happy to exist should do nothing to confuse our judgment that Paula fails to choose the best among her available alternatives. Ultimately, it would seem that Levy's solution to the Non-identity Problem amounts to a Maximizing Consequentialist (MC) solution. Unfortunately, he leaves it woefully under-developed. What we get in the above passage is as fleshed out as his account gets. If Levy is indeed opting for a MC solution, then his account will be open to the standard objections to MC. Because my solution to the Non-identity Problem, like Levy's, appeals to some sort of duty to produce more valuable rather than less valuable outcomes, I will take the failure of his account as a jumping-off point for my own in the following chapter.

¹¹³ Ibid: 366.

IV

On the Existence of a Moral Reason to Promote the Good: The Beginnings of a Solution to the Non-identity Problem

Chapters II and III discussed harm-based and non-harm-based solutions to the Non-identity Problem. I take the central lesson of those chapters to be that there are serious worries for attempts to solve the Non-identity Problem by appealing to harm or rights. In light of these worries, a successful solution to the Non-identity Problem must identify some wrong-making feature of Wilma's choice to conceive Pebbles and the wealthy society's choosing Depletion without claiming that they wrong *some particular individual*. The project of this chapter is to identify a non-harm-based and non-rights-based wrong-making feature of the actions that give rise to the Non-identity Problem.

I have rejected the categories of solutions to the Non-identity Problem that deny premises (P1-P4) of NIPA. This means that if the solution that I am to propose in this chapter and the next is to be successful, it must succeed by way of rejecting (P5): if an act does not wrong a particular individual then it is not wrong. In order to successfully reject (P5) we must identify a wrong-making property of actions other than the property of wronging someone. In the previous chapter I considered a couple of solutions that reject (P5) (Kavka, Parfit and Levi), and I found them to be unacceptable. Some have doubted that we might find a property "that successfully entails that Wilma's act is morally wrong without also generating further implications that are at least as

problematic as the claim that her act is not wrong.”¹¹⁴ This pessimistic outlook is likely due to a tendency of philosophers to suppose that positing such a property would entail a moral theory that looks disturbingly like Maximizing Act Consequentialism:

MC: An action is permissible if and only if it produces the best outcome of the actions available to the agent.¹¹⁵

MC has well known counterintuitive implications in at least two kinds of cases. The first kind of case exposes MC’s overly demanding implications. For instance, if it would be better for you to watch a movie at home and donate some money to UNICEF rather than spend that money on going out to the movies, then MC entails that you are morally obligated to stay at home. This seems like an overly demanding implication. A second kind of case exposes MC’s lack of respect for what has been called by some ‘the separateness of persons’.¹¹⁶ For instance, the view entails that one can be morally required to kill one person in order to prevent the death of two others. Even more implausibly, the view entails that one might be morally required to kill one person in order to prevent a sufficiently large number of headaches.¹¹⁷ Defenders of MC are not

¹¹⁴ Boonin (2008): 142.

¹¹⁵ Neil Levy, ‘The Apology Paradox and the Non-identity Problem,’ *The Philosophical Quarterly* 52 (July 2002): 358-368 suggests something similar.

¹¹⁶ Rawls (1971) famously endorses the moral relevance of the separateness of persons in opposition to Utilitarianism.

¹¹⁷ Norcross (1997) partially responds to this criticism of MC, arguing that we are committed, by non-Consequentialist reasoning, to the view that some number of headaches is morally worse than one death.

without strategies for responding to these kinds of objection, but these two kinds of cases have moved many people to reject MC. Of course, this means that most will reject any solution to the Non-identity Problem that depends on MC, on the grounds that the solution would fail to satisfy the Intuitive Cost Criterion for a successful solution to the Non-identity Problem.

In this chapter and in Chapter V I develop a moral theory that entails the rejection of P5 of NIPA. This theory is not overly demanding and it satisfies the intuition that we ought to respect the separateness of persons. I argue that we can avoid the conclusion of NIPA, concluding instead that Wilma and the Wealthy Society both act wrongly, by appealing to what I consider to be moderate moral principles that virtually everyone will accept. My solution to the Non-identity Problem is inspired, but not strictly supported by, a pair of claims defended extensively by Shelly Kagan. These claims concern the nature of the good and the nature of the reasons we have to promote it.¹¹⁸ I will call the first claim ‘Welfare’.

Welfare: The welfare of individuals is one of the things that constitute the good.

I will refer to the second claim as ‘Pro tanto’.

Pro tanto: There is a moral reason to promote the good, and this reason generates a moral requirement in the absence of sufficient countervailing reasons.

¹¹⁸ Kagan (1989)

Pro tanto is a departure from MC because the principle's 'sufficient countervailing reasons' clause allows space for moral considerations other than promotion of the good. Specifically, the clause leaves open the possibility that significant personal sacrifice and the separateness of persons are morally relevant considerations. The content of countervailing reasons clause and the importance of countervailing reasons for the success of my view will be emphasized in Chapter V. For now I will settle on the following analysis of 'countervailing reason':

CV: A reason r is a countervailing reason just in case it counts against some action a that is supported by a reason r' ; r must arise from a different kind of moral consideration than r' .

I will discuss the kinds of considerations that might qualify as countervailing reasons in the next chapter. For now it will do to note that Maximizing Consequentialists cannot recognize countervailing reasons, defined as they are in CV, as being morally relevant. All of one's moral reasons, according to Maximizing Consequentialism, are reasons to promote the good, and these reasons are univocal in requiring agents to do 'the best'.¹¹⁹ Pro tanto leaves open the possibility that there are reasons that count against agents always doing the best with regard to promoting

¹¹⁹ Pluralistic Consequentialists might object at this point that their view leaves room for countervailing reasons because they recognize different elements of the good. So, that option A promotes happiness is a reason to do A, but that option B (inconsistent with A) promotes friendship is a reason to do B instead. These are not genuine countervailing reasons, however, for they both derive from a more general concern with promoting the good. They are simply different avenues to accomplishing that same goal.

the good. In Chapter V I argue that a suitable interpretation of Welfare and Pro tanto, conjoined with supplementary premises about Wilma's and the Wealthy Society's particular moral circumstances, yields a solution to the Non-identity Problem that is consistent with the major components of ordinary morality. My aim in this chapter is more modest. Here I aim to carve out the conceptual landscape and to argue that the simple versions of Welfare and Pro tanto, described above, are true.

The argument contained in this chapter has the following structure: in IV.1 I discuss some preliminary terms and distinctions that will be important for the discussion in this chapter. In section IV.2 I argue that Pro tanto and Welfare are indispensable components of morality. For my solution to the Non-identity Problem to be successful I must first defend these principles in their basic and modally neutral form. If Welfare and Pro tanto are false, then, *a fortiori*, the formulations of the principles that I employ to solve the Non-identity Problem are false, and my solution fails. In section IV.3 I distinguish between two competing formulations of Pro tanto and Welfare. I show that only what I call the 'unrestricted' formulations (in contrast with competing 'restricted' formulations) of Pro tanto and Welfare can be used to solve the Non-identity Problem. I argue in the next chapter that we should accept the unrestricted formulations of Welfare and Pro tanto rather than the restricted formulations, but here my aim is simply to lay out the ground work. I have already introduced some terms, the meaning and referents of which may be unfamiliar. Before proceeding into the more complicated territory of IV.2 and IV.3, then, I discuss some of the terms and distinctions contained in those sections.

IV.1 Preliminary Terms and Distinctions

The debates contained in this chapter and the next will turn on arguments that employ the terms ‘the good’, ‘promoting the good’, ‘possible people’, ‘merely possible people’, ‘actual people’, ‘countervailing reasons’, and ‘sufficient countervailing reasons’. It is not immediately clear what we should take the referents of these terms to be (indeed it is not immediately clear that all of them have referents), and the distinctions between them are sometimes subtle and easily misunderstood.

IV.1.1 Preliminary Arguments about ‘The Good’

I defend my preferred versions of Welfare and Pro tanto in Chapter V, and this defense constitutes part of my proposal for a solution to the Non-identity Problem. Both Welfare and Pro tanto appeal to *the good* as a central moral consideration, so before I defend my preferred versions of the principles, it is important that there is no question about what *the good* refers to. By the end of this section it should be clear that the conception of the good needed for my purposes as it appears in Welfare and Pro tanto is unproblematic. Jeff McMahan distinguishes between three kinds of value (disvalue): comparative person-affecting value, non-comparative person-affecting value, and impersonal value.¹²⁰ Though he uses the term ‘value’, I take the subject matter of his discussion to be the same as mine. I will thus continue to use the term ‘good’, and will simply note that I consider it to be intersubstitutable with his term ‘value’. It is worth noting that this discussion about the good will not assume any theory of welfare. I will

¹²⁰ Much of the discussion in this section is based on McMahan (2009).

sometimes use pleasure as an example of something that is good for someone, but my discussion is meant to be equally amenable to hedonist, preferentist, and objective list theories of welfare.

Figure 1 is a helpful starting point for exhibiting the explanatory upshots of the theoretical differences between the three kinds of goods (bads) that I will discuss.

Figure 1




Figure 1 represents outcomes A, B, and C. Population x enjoys welfare level 90 in A and C, and welfare level 100 in B. Population y enjoys welfare level 120 in C but does not exist in A or B. Suppose, as seems plausible, that B is a better outcome than A.

IV.1.2 Three Kinds of Goods

First we can distinguish between two kinds of ‘person-affecting [goods].’ Person-affecting goods are those that, in and of themselves, make individuals’ lives go well or better. One kind of person-affecting good is what McMahan calls a ‘comparative person-affecting

[good]'. A comparative person-affecting good is a good the obtaining of which is *better* for someone than its absence would be. When I am caused to enjoy, say, one more unit of pleasure, my experiencing this pleasure is a comparative person-affecting good for me, because it is better for me to have it than not to have it. A comparative person-affecting bad is worse for someone than its absence would be. When I am caused to suffer, say, one additional unit of pain, this is a comparative person-affecting bad, because it is worse for me to suffer this additional pain than not to suffer it. The 10 additional units of welfare enjoyed by population x in outcome B in Figure 1 is a comparative person-affecting good because it is better for the members of y than its absence would be (represented in outcomes A and C). Comparative person-affecting good is inextricably tied to the *welfare* of individuals. When someone enjoys a comparative person-affecting good, her welfare level is in some respect higher than it otherwise would have been. We can talk of someone's *lifetime* welfare level, determined by summing all of the person-affecting goods that she enjoys and then subtracting all of the person-affecting bads that she suffers over her whole life. We can also talk of someone's welfare level *at a time* or *during an interval*, determined by summing all of the person-affecting goods and then subtracting the person-affecting bads that obtain *at that time* or during that interval.¹²¹

A second kind of person-affecting good is what McMahan calls a 'non-comparative person-affecting [good]'. This good is also related to the welfare of individuals. A non-comparative person-affecting good (bad) is good (bad) for a person to have in an *absolute* sense; its presence makes her life go *well (badly)* to some degree, but its presence is *not better (worse)* for that person than its absence would have been. The goods associated with being caused to

¹²¹ This analysis is more or less the same as Ben Bradley's analysis (2009: 18).

exist with a life worth living might be non-comparative person-affecting goods, because it might turn out that the goods associated with a life worth living are not *better* for the person that enjoys them than non-existence.¹²² We might think this is the case if, e.g., a person's welfare level when she exists cannot be coherently compared with what her welfare level would have been had she not come to exist.¹²³ In that case, the goods that accompany a life worth living would be non-comparative person-affecting goods, but not comparative person-affecting goods. In Figure 1, supposing that a life worth living is not better than non-existence, outcome C is non-comparatively good for population y. y has a positive welfare level in C (a life worth living), but y does not exist in A and B. Thus the benefits y enjoys in C, though they are good for y, are not *better* for y than their absence would have been.¹²⁴

A third type of good is what McMahan calls an 'impersonal good.' According to McMahan, 'Value [good] is impersonal when it is neither good or bad nor better or worse *for* anyone.'¹²⁵ According to McMahan, "for all or at least most individual-affecting values, there are corresponding impersonal values. An individual's suffering is, in itself and apart from its effects, bad in individual-affecting terms. But it is also impersonally bad. Its presence makes the world

¹²² I will take a 'life worth living' to refer to a life in which the non-instrumental goods outweigh the non-instrumental bads.

¹²³ The view that such comparisons are incoherent is contentious. I discuss the coherence of this kind of comparison in more detail in Chapter V. For now I note that Holtug (2001) argues that there is no difficulty in making comparisons between someone's welfare when she exists with what her welfare would have been had she never come to exist. Bykvist (2007) and recently Herstein (2013) argue that such comparisons are indeed incoherent. It is not, they conclude *better* to exist than not.

¹²⁴ The contribution of non-comparative person-affecting value to the good will be the subject of a later section.

¹²⁵ McMahan, (2009, p. 50).

worse.”¹²⁶ Though the impersonal good might often correspond to some person-affecting good of either the comparative or non-comparative variety, not all impersonal goods will correspond to person-affecting goods. Beauty in art or equality are not person-affecting goods (their presence might not be good for or better for anyone), but they might nonetheless correspond to an impersonal good. The world might be better when they are present and worse when they are absent. McMahan’s remarks here seem to indicate that he thinks there is an ontologically distinct entity – *the impersonal good* – that corresponds with (perhaps because it supervenes on) but is distinct from the person-affecting goods and non-person-affecting goods to which it corresponds. I am only going to consider person-affecting goods in the remainder of this chapter, since those goods will prove important for solving the Non-identity Problem.

Let us suppose for the moment that McMahan is correct that person-affecting goods correspond with but are ontologically distinct from impersonal goods. This leaves us with several options regarding what we should say about the relationship between outcomes A, B, and C in Figure 1:

Option 1: B is better than A and the additional *comparative* person-affecting good in B alone explains this fact.

¹²⁶ Ibid, p. 51. I take it that what McMahan calls ‘impersonal value’ is what philosophers like Philippa Foot have criticized under the label ‘goodness *simpliciter*’.

This explanation does not entail that ‘B is better than A’ means the same thing as ‘B is better for x than A’, it simply entails that, other things being equal, the truth of the proposition expressed by one sentence entails the truth of the proposition expressed by the other. Other things being equal, it is true that B is better than A just in case it is true that B is better for x than A. On this picture, the additional impersonal good contained in B that supposedly corresponds to the additional comparative person-affecting good plays no role in the explanation of B’s being better than A. If we accept Option 1 then we are pressed to reject the existence of impersonal goods, because they seem to have no explanatory role in determining the relative value of outcomes. Other things being equal, it seems that we should deny the existence of explanatorily useless entities. I will turn to this possibility shortly. Of course this explanation gives us no guidance in determining C’s ranking relative to either A or B. *Non-comparative person-affecting good* obtains in C, because y only exists in C, but it is not yet clear how non-comparative goods are to be weighed against comparative person-affecting goods. How then do we rank C relative to A and B? Answering this question is, in large part, the project of this chapter and the next.

Option 2: B is better than A because more impersonal good obtains in B. This impersonal good corresponds with the comparative person-affecting goods in B, but the impersonal goods *themselves*, rather than the comparative person-affecting good to which they correspond, explain why B is better than A.

What this explanation implies for the relationship between C and the other outcomes depends on the nature of the relationship between non-comparative person-affecting goods and the impersonal good. Option 2 will be appealing to those who wish to accommodate the possibility that purely impersonal goods like beauty, proportional justice, or equality contribute value to an outcome. These goods might not be better for anyone if they obtained so they can only have impersonal value.

Option 3: the gain in person-affecting value *and* the gain in impersonal value that corresponds to it explain why B is better than A.

The idea is that B is better than A in two respects: it is better in terms of person-affecting value, and it is better in terms of impersonal value. It seems that this explanation problematically double counts person affecting comparative good. It counts its person-affecting value *and* it counts its impersonal value in determining the relative value of an outcome. One way to put the worry is this: if I know that B is better for X than A, I already know the degree to which that outcome is better in virtue of its being better for X. I do not also need to know information about how much impersonal good corresponds to the extra person-affecting value enjoyed by X in B to know that B is better than A.

Recalling McMahan's remarks above gives guidance as to which possibility he prefers. He says "...An individual's suffering is, in itself and apart from its effects, bad in individual-

affecting terms. But it is also impersonally bad. Its presence makes the world worse.”¹²⁷ The last two sentences of this quotation seem to indicate that McMahan prefers Option 2. It is the increase in impersonal good, rather than the increase in person-affecting good to which it corresponds, that explains why B is a *better world* than A. McMahan seems to think that a world’s goodness or badness is a function of the impersonal goodness or badness that obtains rather than the person-affecting goodness or badness to which the impersonal good corresponds.

IV.1.3 Against the Impersonal Good

McMahan’s conceptual distinctions between goods are helpful, but I am wary of accepting the view, suggested by his remarks, that the impersonal good *corresponds* with, but is nonetheless ontologically distinct from, the person-affecting good (either comparative or non-comparative). I suspect that we can do away with positing the existence of an ontologically distinct, impersonal good for two reasons. The first reason is that the impersonal good, if it did exist, would be explanatorily useless. I discuss three motivations for positing the existence of the impersonal good, and I argue each of them can be satisfied without appeal to the impersonal good. The second reason to reject the impersonal good is that positing its existence needlessly complicates the moral landscape, leaving its proponents open to the objection that they misidentify the source of moral value.

The impersonal good, as McMahan construes it, might seem to play three explanatory roles. First, positing the existence of the impersonal good might seem to be the best way to

¹²⁷ Ibid: 51.

justify certain judgments we make about one outcome being *better* ‘all-things-considered’ or ‘*simpliciter*’ than another in cases where neither person-affecting nor non-person-affecting value obtains. McMahan says:

Impersonal value may also make essential reference to well-being and yet not correspond to any individual-affecting value. Because impersonal value is comparative, it is impersonally better, other things being equal, if an individual whose life would not be worth living never exists than if she does exist. Similarly, it is impersonally worse if an individual whose life would be worth living never exists than if she does exist. In such cases, there may be no one for whom an individual’s never existing is better or worse, or even non-comparatively good or bad.¹²⁸

McMahan worries that, without impersonal goodness or badness we are left without the resources to explain why it is appropriate to judge that it is *worse* when someone who would have had a life worth living never comes to exist or that it is *better* when someone who would have had a miserable life never comes to exist. Such states are not, after all, better (worse) or good (bad) for the person who never exists, he assumes, so it seems that their betterness or worseness cannot be grounded in person-affecting goods. As justification of these judgments, however, we need not be committed to the existence of the impersonal good. The truth of the

¹²⁸ McMahan, p. 51 (2009)

claim that it *is better* when a miserable person never comes to exist can be explained by the fact that her existence would have been *non-comparatively* bad, to some degree *n*, *for her*. The degree to which existence would have been bad for the miserable person is the degree to which the world in which the miserable person exists would be ‘worse’ than the world in which she does not. The fact that the world in which the miserable person never comes to exist *lacks* the negative non-comparative person-affecting value that accompanies her miserable existence explains *in person-affecting terms* why that world is ‘better’ than the one in which the miserable person comes to exist. The non-comparative person-affecting bad, not the impersonal good, plays the important role in this explanation. A similar person-affecting explanation can be offered for the claim that it is worse when a happy person fails to come to exist. This response implies that non-comparative person-affecting good partly determines the ranking of outcomes according to their moral worth.

This reply also undermines McMahan’s second reason for positing the existence of the impersonal good, that only the impersonal good (bad) can explain why we have *moral reasons* to perform certain actions rather than others. Specifically, positing the existence of the impersonal good might be the only way to explain why we ought to avoid bringing an additional miserable person into existence. McMahan, who supposes that existing with a miserable life can be *bad* but not *worse* for someone than non-existence, contends that the fact that existence can be bad but not worse for this person entails that if there is a reason not to cause a miserable person to exist then this reason must derive from impersonal rather than person-affecting considerations:

Some philosophers think this [the reason not to create a miserable person] is an individual-affecting reason, one that does *not* derive from the effects the act would have on pre-existing people. But this is a mistake. It is true, of course, that *if* one causes such an individual to exist, there will *then* be someone for whom one's act was *bad*, though not *worse*. Yet at the time of one's choosing between acting and not acting, **there is no one whose interests would be affected by one's choice**. If one were to act on the reason not to cause an individual to exist, **there would never be anyone for whom that would be better**. And if one were to act *against* that reason, **there would never be anyone for whom one's act was worse**. The fact that acting against the reason would be bad in non-comparative individual-affecting terms seems insufficient to ground an individual-affecting reason not to cause a miserable person to exist.¹²⁹

Here McMahan seems to suppose that if an action will not be worse (better) for someone, then there cannot be a person-affecting reason that counts against (in favor of) performing that action. He does not do much to justify this claim, however, other than to say that non-comparative person-affecting considerations 'seem insufficient' to ground such a reason. However, the fact that some act would cause a person to exist with a life that is non-comparatively *bad* for her seems at least as capable of grounding a reason not to perform the act as the fact that the existence of such person would be impersonally *worse*. It is at least not

¹²⁹ McMahan, p. 52 (2009).

immediately obvious how the fact that the miserable person would not otherwise exist poses a problem for the ability of non-comparative person-affecting considerations to ground a person-affecting moral reason of this sort (unless we beg the question by assuming that *only* comparative person-affecting considerations can ground a reason of this sort). Similarly, in the case of comparative person-affecting goods, the fact that some action would result in someone enjoying a benefit that is better for her to have seems just as capable of grounding a reason to perform the action as the fact that the benefit to her would be impersonally better. I see no reason to prefer the impersonal explanation over the person-affecting explanation here. If the impersonal good is not needed to justify any of the judgments that most will want to make about the kinds of cases that McMahan appeals to, then I see no positive reason for positing its existence.

A third reason to posit the existence of the impersonal good is that doing so might seem to be the only way to accommodate the non-instrumental value (value that is had by something apart from its further effects) of ‘impersonal’ ideals such as equality, beauty, and proportional justice.¹³⁰ These ideals, if they have non-instrumental value, must have value that is not reducible to their effects on the welfare of individuals; they contribute value to an outcome even if their obtaining is not better *for* anyone. For instance, someone who holds that beauty is a non-instrumentally valuable ideal might claim that a world with an abundance of beauty, but in which everyone is slightly worse off, is better than a world in which there is very little beauty, but in

¹³⁰ Kagan (1988) prefers the term ‘final value’ to refer to the kind of value I discuss here on the grounds that ‘non-instrumental value’ is misleading. ‘Non-instrumental value’ appears to describe a kind of value that has its value solely in virtue of its non-instrumental properties. I avoid this confusion by using the term ‘non-instrumental value’ to refer to value that is worth pursuing regardless of its further effects. This kind of value is neutral between extrinsic and non-instrumental value.

which everyone is slightly better off. Denying the existence of the impersonal good may seem to implausibly rule out the possibility that goods such as beauty contribute non-instrumental value to outcomes apart from their effects on individual welfare. I think this objection to rejecting the existence of the impersonal good has its source in an ambiguity concerning the term ‘impersonal’. McMahan uses ‘impersonal good’ to refer to a kind of good that corresponds to person-affecting goods which contributes value specifically to an outcome or a world. The kind of impersonal goods under present consideration are not of this sort. Beauty does not correspond to some person-affecting good. It is rather a *competing* good, which stands apart from person-affecting goods. The same is true in the case of equality and proportional justice. It would be better then to call beauty, equality, and proportional justice ‘non-person-affecting’ goods. One way to think about the distinction between impersonal and non-person-affecting goods is to notice that non-person-affecting goods might correspond to an *impersonal* good of the sort that McMahan thinks corresponds to person-affecting goods. If non-person-affecting goods were identical to the impersonal good, then correspondence would seem to be the wrong kind of relation to posit between them.

With this distinction between impersonal goods and non-person-affecting goods in mind, note that I have at no point suggested that non-person-affecting goods do not exist or that they fail to contribute value to outcomes. In fact, though I will not defend this view here, I find it plausible that at least some of these goods do in fact contribute non-instrumental value to outcomes. It seems plausible, for instance, that it is better when people get what they deserve. I can think of one reason that accepting the existence of non-person-affecting goods might seem to require positing the existence of the impersonal good. Suppose that both person-affecting and non-person-affecting goods contribute value to outcomes; an outcome is better when there is an

increase in either aggregate welfare or aggregate beauty, and these goods can be weighed against one another. On this supposition, there may seem to be a concern about how to *compare* these person-affecting and non-person-affecting goods that appear to be different *kinds* of good. Appealing to the impersonal good might appear to be the only way to unify these goods under one kind of value, rendering them comparable. If all non-instrumentally valuable goods corresponded to the impersonal good, and the *kind* of impersonal good were to remain the same across person-affecting and non-person-affecting goods, then we would have a neat way to compare these seemingly independent goods when determining the value of outcomes. I believe this is the strongest motivation for positing the existence of the impersonal good. However, I do not think that positing the existence of the impersonal good is the only way to compare person-affecting and non-person-affecting goods when determining the value of an outcome. In fact, I think that Pro tanto provides the framework for comparing these seemingly independent goods. According to Pro tanto, there is a moral reason to promote the good. If non-person-affecting and person-affecting goods jointly constitute the good because each contributes non-instrumental value to outcomes, then it is possible that there is *one and the same* kind of reason (a moral reason) to promote each *different* kind of good. We can compare these goods, not by appealing to a third party *impersonal* good, but by appealing to the moral reason to promote each good. Supposing that person-affecting and non-person-affecting goods generate the same kind of moral reason, we can compare their relative value by comparing the *strength* of the moral reason to promote each good. This solves the comparison problem without positing the existence of the impersonal good.

I have so far attempted to debunk three justifications for positing the existence of the impersonal good, construed as something distinct from the comparative or non-comparative

person-affecting goods to which it corresponds. I now turn to providing a positive reason to deny the existence of the impersonal good. One worry for Maximizing Consequentialism or any theory that posits a pro tanto moral reason to promote the good is that it misidentifies the “source of moral value.”¹³¹ The idea is that a theory that demands that agents promote the good fails to capture the compelling thought that a person’s welfare, as Kagan describes it, “matters in its own right,” and therefore “directly generates a reason, as opposed to generating such a reason indirectly, through the mere fact that aiding [someone] happens to promote the good.” This worry only seems to get traction, however, if we accept McMahan’s characterization of the relationship between person-affecting goods and the impersonal good. On his explanation, person-affecting goods (the things that make peoples’ lives go well) seem to be mere means to increasing the impersonal good, and this seems to misidentify the source of more value. The reason that I ought to benefit you does not derive from the fact that I would thereby promote some impersonal good; it derives from the fact that benefitting you would be good or better *for you*. Kagan, like me, seems to think we can do away with the impersonal good. A person-affecting good for someone *just is* a contribution to the good, because individual welfare *is the stuff* that the good is made of. Thus, the fact that doing something would bring about a person-affecting good for someone generates a direct reason to bring it about. As Kagan puts it, “There is no need to choose between ascribing to ordinary morality the acceptance of a pro tanto reason to promote the good, and, say, a pro tanto reason to promote individual well-being. Given a plausible account of the good, the former simply encompasses the latter, under a different level of description.”¹³² The person-affecting goods (and possibly non-person-affecting goods) matter

¹³¹ Ibid, pp. 57-58

¹³² Ibid, p. 59.

in their own right, directly generating reasons to promote them. Together this collection of goods makes up what I call ‘the good’. The welfare of each person, as a constituent of the overall good, can serve as the source of moral value without appealing to the impersonal good. The good is not some abstract entity beyond the things (welfare, for instance) that constitute it. In Kagan’s words, “...well-being is no mere instrument to the realization of the good—it is the very *stuff* of which the good consists...the overall good simply *is* the overall well-being of individual persons.” On this view, B is better than A because B contains more person-affecting good for population x, and that is the whole story. To say that B contains more good for population x (because of its additional person-affecting value) just is to say that B contains more of the good. Of course I have not yet offered a defense of Pro tanto or Welfare, but by eliminating impersonal value from our ontology of goods pertaining to individual welfare I hope to have headed off one of the more significant criticisms of theories that posit a moral reason to promote the good: that they misidentify the source of moral value.

IV.1.4 Welfare, Pro tanto, and the Good

Consider Welfare and Pro tanto again:

Welfare: The welfare of individuals is one of the things that constitute the good.

Pro tanto: There is a moral reason to promote the good, and this reason generates a

moral requirement in the absence of sufficient countervailing reasons.

It should be clear by now that the kind of good I have in mind as it appears in Welfare and Pro tanto is not abstract or mysterious. It is nothing more than the collection of the person-affecting goods (experiences of pleasure, preference satisfactions, etc.) that obtain in a world. If one is inclined to talk about the good in terms of states of affairs, then I am content to say that the good is just the collection of the non-instrumentally valuable states of affairs (e.g. a subject *s*' enjoying a pleasurable experience, a subject *s*' having her preference satisfied, or an object's exemplifying certain aesthetic properties, etc.). The Welfare principle is thus a fairly weak principle. It simply claims that the welfare of individuals, whatever individual welfare turns out to consist in, is one of the things that constitutes the good, where the good includes (at least) all of the person-affecting goods that obtain in the world. Since a unit of welfare must either be a comparative or non-comparative good for the individual who enjoys it, Welfare should not turn out to be controversial. It is true that B is better than A in figure 1, because more person-affecting good obtains at B than A. That is all there is to say.

Pro tanto claims that there is a moral reason to promote the good. This is tantamount to claiming that there is a moral reason to promote the person-affecting goods (leaving open the possibility that there is a moral reason to promote non-person-affecting goods). In fact we have already seen one reason for accepting Pro tanto. Positing a moral reason to promote the good is the best way to avoid positing the existence of the impersonal good while rendering person-affecting and non-person-affecting goods comparable for the purposes of assessing the overall moral value of outcomes. There is no further, abstract, mysterious, impersonal good that is the

real goal in promoting the person-affecting goods. In IV.2 I argue that Pro tanto, construed this way, is the best explanation of a subset of our central moral convictions.

Some have objected that the term ‘the good’, described as something that is non-instrumentally valuable and not relativized to some particular perspective or domain, is non-referring. Philippa Foot raises this worry most forcefully for Utilitarianism.¹³³ On one construal of her criticism, it seems to have some plausibility. If Foot merely means to claim that there is no further impersonal good, apart from the collection of person-affecting goods (and non-person-affecting goods), then I agree with her criticism. If however, Foot means to claim that person-affecting goods cannot make some outcomes better than others in a perspective-neutral sense, then I think she is mistaken. Perhaps it is easiest to individuate the good, described as I have done, by contrasting it with several other distinct kinds of goodness. There is goodness *at*. An assassin, for example, might be good *at* killing diplomats. There is goodness *for*. A rainy summer, for example, might be good *for* a forest. But there also seems to be another, distinct kind of goodness. This distinct kind of goodness is non-instrumentally valuable, and it determines the ordering of outcomes in terms of moral value. It is good regardless of its further effects and its presence makes outcomes better. The welfare of individuals is one kind of this good, and non-person-affecting goods like beauty might be another kind. This is the good that I have in mind as it appears in Welfare. One way to get a grip on this kind of goodness is to recall figure 1. B is clearly better than A. Foot would contend that we can truly say B is better than A *for* x, and that is all there is to say. But surely it is reasonable to ask whether it is *good* that B is

¹³³ See Philippa Foot, ‘Utilitarianism and the Virtues,’ *Mind* 94 (1985): 196-209 for some examples designed to show that goodness simpliciter is not a concept of which we have a firm grasp. Julia Annas has also raised a worry like this in personal correspondence.

better for x than A . It may be that B is better than A *because* it contains more person-affecting good *for* population x , but this is not to concede that all goodness is relativized to particular aims or perspectives. In fact, the claim that B is better than A because B is better for x than A is precisely the claim that I defended in section IV.1.2 and IV.1.3. This is consistent with claiming that B just is better *simpliciter*, in a perspective neutral way, than A , and that we have a moral reason to bring about B rather than A if we can.¹³⁴ I see nothing mysterious about this kind of talk. In fact, I think relativizing all goodness to particular perspectives is a strained and under motivated effort to undermine consequentialist reasoning.

IV.1.5 Preliminaries about the Modality of Persons

I have distinguished so far between several kinds of goods, and I have shown what I take ‘the good’ to refer to as the term appears in Pro tanto and Welfare. ‘The good’ refers to at least the collection of person-affecting goods that obtain in a world. One world is better than another in at least one respect if more person-affecting goods obtain in that world than the other. Several more distinctions will be important for the discussion in the rest of this chapter and the next. I will frequently talk about *actual* people, *possible* people, and *merely possible* people. An *actual* person is someone who exists, has existed, or will at some point exist.

¹³⁴ It is worth noting that accepting a view about the good that holds that something is good in an unqualified sense if and only if it is good for someone might commit one to rejecting theories of the good according to which the values of desert, justice, or art are constitutive of the good, but I will refrain from engaging in this debate.

Actual person: A person who exists at some point given that you chose Act 1 (Act 2).

After the point in time at which it is determined that a person will, at some point, come to exist, it is true at all later times that the person is an actual person. Most people will agree that the welfare of *actual* people matters morally. If we can avoid doing something that would negatively affect the person-affecting goods of actual people, most will agree that this gives us a moral reason not to do it. The moral reasons we have to respect the welfare of actual people can generally be accounted for by the comparative account of harm. A *possible* person is someone whose existence is not yet settled. A possible person might come to exist or she might not come to exist depending on the choices that people make.

Possible Person: A person who will exist if Act 1 is chosen but who will not exist if you Act 2 is chosen.

A person only counts as a possible person until the point when it is determined whether she will or will not come to exist. Once that point is reached she is either an actual person or a *merely possible person*. A merely possible person is someone who might have existed but who, due to the way events have unfolded, never will.

Merely Possible Person: A person who will never exist given that you chose Act 1 (Act 2) but who would have existed at some point if you had chosen Act 2 (Act 1).

Nixon's sister, who would have been Secretary of State, had she come to exist, is a merely possible person in this sense. At all times it is true of every possible person that she is either an actual person or a merely possible person.

Let's consider these distinctions in the context of the case of Wilma. Before Wilma makes the decision to conceive Pebbles, Wilma is an actual person, but her two prospective children Pebbles and Rocks are each possible people. Each of them either will or will not exist, depending on what Wilma decides to do. Once Wilma has decided to conceive Pebbles, Pebbles and Wilma are the actual people, and Rocks is a *merely* possible person. Rocks could have existed but will never exist given Wilma's decision to conceive Pebbles.

It is worth heading off a worry now that, left unaddressed, might distract the reader. One might wonder whether we can successfully refer to Rocks, or to possible or merely possible people more generally, given that they do not and will never exist. One might worry that by talking this way I am implicitly assuming that possible people are some sort of special being that exists in outcomes in which it is stipulated *not* to exist. This talk might seem problematic, because it might seem that I am taking a stand on the side of possibilism against actualists who hold that '[i]f an individual exemplifies a property or stands in a relation in a world, it must exist (i.e. be actual) in this world'.¹³⁵ I do not think that I need to take a stand in this debate, however. Surely actualists and possibilists will both accept that we can say true things about what might

¹³⁵ Bykvist (2007: 339). For a survey of the debate between actualists and possibilists see C. Menzel, 'Actualism' in E. N. Zalta (ed.) *The Stanford Encyclopedia of Philosophy* (Winter 2008 Edition) URL = <<http://plato.stanford.edu/entries/actualism/#ActualistResponses>>.

have happened had events in history unfolded differently than they actually did. In other words the following statements seem unproblematic:

- (i) Had the Confederate South never seceded from the Union, then migration patterns in the US would have been very different.
- (ii) Had migration patterns in the US been very different, then different couples would have met and formed relationships than the couples that actually met.
- (iii) These couples would have conceived different children than they actually conceived.
- (iv) The people who would have been conceived by these couples would have done things and had experiences.

Surely (i)-(iv) could turn out to be true despite the fact that they involve making claims about individuals who are in fact merely possible people. (i)-(iv) refer to events, states of affairs, and people who might have, but never actually existed, and most will agree that this is unproblematic. As long as we think this is unproblematic, then my use of the terms ‘possible person’ and ‘merely possible person’ are similarly unproblematic. The metaphysical challenge is figuring out *how* we can talk about possible and merely possible people, not *whether* we can. Meeting this metaphysical challenge is beyond the scope of my project.¹³⁶

IV.1.6 Preliminaries Concerning the Moral Relevance of Possible People

It is not obvious that we should consider the welfare of *possible* people in our moral

¹³⁶ Lewis (1986) provides a framework for explaining how claims (i)-(iv) are true as well as a framework for providing referents for the terms ‘possible person’ and ‘merely possible person.’ I am not endorsing Lewis’ view. I am merely pointing out that most philosophers are concerned with figuring out *how* counterfactuals like (i-iv) turn out to be true, not *whether* they turn out to be true.

decision-making, and it is perhaps even less obvious that we should consider welfare of *merely possible* people, since the latter will certainly never exist. There are at least two formulations of Welfare and two formulations of Pro tanto. What we say about the moral status of a person's modal status will depend on which of the formulations we accept. First consider two competing ways of formulating Welfare:

Restricted Welfare (RW): The welfare of actual people constitutes the good (bad), but the welfare of possible people, would not constitute the good (bad) were it to be actualized.

Unrestricted Welfare (UW): The welfare of actual people constitutes the good (bad), and the welfare of possible people, would constitute the good (bad) were it to be actualized.

RW is restricted because it claims that only the welfare of people who currently exist, will exist, or have existed can constitute the good. It is one way to arrive at the conclusion that we do not need to include the welfare that possible people would enjoy if they were to come to exist among our moral considerations. According to RW, this is because the welfare of people who might or might not exist *would not constitute the good* even if it actually obtained. RW should be rejected, however, because it asserts a contradiction. All possible people will turn out to be either merely possible people or actual people. If the welfare of possible people is 'actualized', then it must be the case that they are actual people, so RW amounts to the claim that the welfare of actual people constitutes the good and that the welfare of actual people does not

constitute the good. RW therefore asserts a contradiction. UW is the only acceptable option.

Thus, if we accept Welfare then we should accept UW. I will make a few remarks in defense of Welfare, thereby defending UW, in section IV.3.

A restricted version of Pro tanto is more plausible. It concedes that welfare of possible people would constitute the good were it to be actualized, but maintains that only the welfare of actual people matters (in a very restricted sense) when determining what agents have moral reasons to do. This idea is captured by a formulation of Pro tanto that I will call ‘Restricted Pro tanto’ (RP). According to RP, the person-affecting goods that possible people would enjoy, or the person-affecting goods that merely possible people would have enjoyed, if they were to come to exist doesn’t matter morally. Therefore, these goods do not give us moral reasons to bring them about. Consider it below:

Restricted Pro tanto (RP): There is a moral reason (not) to perform an action A if the set of people given A will be better (worse) off given A than the set of people given $\sim A$, but this moral reason generates a moral requirement in the absence of sufficient countervailing reasons **only if some person would be worse off** given $\sim A$ than given A.¹³⁷

RP tanto limits moral concern to actual people without entailing a contradiction of the

¹³⁷ This principle does not rule out the possibility that there could be a moral reason to promote the good in cases where failing to do so would leave no actual person worse off. It simply rules out the possibility that this moral reason could generate a moral *requirement*.

sort implied by Restricted Welfare. RP entails that we have moral reasons to promote and to refrain from reducing the welfare of individuals who actually exist now or who will come to exist, and it entails that *these reasons* can generate moral requirements. RP also concedes that we can have moral reasons to actualize the welfare of a possible person, but it denies that we ever do something morally *wrong* by failing to actualize the welfare of a merely possible person unless, by doing so, we make some actual person worse off. How does this work? The idea seems to be the following: if your act is such that no one who ever actually exists is made worse off by it, then morality *allows* you to do the act – there’s no requirement that’s violated by doing it; but there might well still be a good *moral reason* not to do it and that moral reason could well be generated by the fact that more overall good would result by not doing it. On this view, the fact that one act will result in less overall good than another does not generate moral reasons not to do the act *of the sort* that are capable of generating moral requirements, *unless* some actual person will be worse off if the act is performed.

Let me note right away that it seems strange to call something a ‘moral reason’ when it makes *no difference* to the moral permissibility of an action. Doing so leaves open the possibility that someone might have overwhelming moral reasons to do A but not be morally required to do A. This peculiarity of RP will be revisited later. For now it is enough to point out that RP determines whether an act (Act 1) is wrong by assessing the welfare of all of the people who exist if we perform Act 1 and then determining whether, given that we choose Act 1, a different act (Act 2) would have been better for *those* people. Given that we choose Act 1, if the interests of the people who exist would not have been better served had we chosen Act 2, then RP implies

that choosing Act 1 rather than Act 2 was not wrong.¹³⁸

In Chapter V I reject RP, defending an ‘unrestricted’ version Pro tanto:

Unrestricted Pro tanto (UP): There is a moral reason (not) to perform an action *A* if the set of individuals given *A* would be better (worse) off than the set of individuals would be given $\sim A$, and **this moral reason generates a moral requirement** in the absence of sufficient countervailing reasons.


RP claims that agents, when it comes to promoting the good, can only be morally required to make people better off or to refrain from making people worse off than *those* people would have been. UP, on the other hand, claims that we *also* can be morally required to ensure that the people who come to exist are (not) better (worse) off than some distinct set of possible people would have been had they been the actual people. In this sense UP entails that the welfare of merely possible people is morally relevant. Note that UP does *not* imply that we could have moral reasons to cause merely possible people to come to exist. Such an implication would seem to violate the plausible notion that one can have a moral reason to ϕ only if one is able to ϕ . Since merely possible people are people who will never exist, we cannot have a moral reason to cause a merely possible person to come to exist. UP claims that there is a moral reason (that can generate a moral requirement) to ensure that, among the set of *possible* people, the group of

¹³⁸ Narveson (1967) seems to defend a principle that is very similar if not equivalent to Restricted Pro tanto. Weinberg (2008) defends a view that seems to be supported by Restricted Pro tanto. David Boonin explicitly defends RP in a forthcoming manuscript.

possible people who ‘become’ actual people are better off than those who might have but will never in fact come to exist.

Though neither RP nor UP mentions promotion of the good, I take it that the moral reason posited by each principle is fundamentally related to the promotion of the good. I will assume that *if* there is a moral reason (not) to perform an action A if the set of people given A will be better (worse) off given A than the set of people given $\sim A$, this reason is (at least partly) derived from the fact that performing A promotes the good to a greater degree than $\sim A$. Of course, I have not yet defended the existence of a moral reason to promote the good. That is the subject of section IV.4. First, let us consider figure 1 again to see how UP and RP diverge with respect to practical implications.

Figure 1



Suppose we wonder whether there is a moral objection to choosing B. If we choose B then B will be the actual outcome, and population x will be the set of actual people and y will be merely possible. x will enjoy welfare level 100. RP tells us to ask if x, the actual people given that we chose B, would have fared better in one of the alternative outcomes. If the answer is no,

then RP will not generate a moral requirement to choose any outcome other than B. The answer is no. Outcome B is not worse for population x than A or C, so RP implies that there is no moral reason of the requirement-generating sort to choose any outcome other than B.

UP, on the other hand, asks us, when choosing between A, B, and C, to compare how the actual people given B fare relative to the people who would have been the actual people had an alternative outcome been chosen, even if the actual people in the alternative outcome would have been *non-identical* with the actual people given B. If the actual people in one of the alternative outcomes would have been better off than the actual people in B, then UP generates a moral reason of the requirement-generating sort to choose the alternative outcome rather than B. When we compare how the actual people would fare if we chose C and compare them to the actual people if we chose B, it turns out that the set of actual people in C would be better off than the set of actual people in B.¹³⁹ UP thus says that there is moral reason to choose C rather than B, *and* that this reason is capable of generating a moral requirement. This key difference between RP and UP means that UP but not RP provides a framework for solving the Non-identity Problem.

IV.2 Unrestricted Pro tanto, Restricted Pro tanto, and Solving the Non-identity Problem

Supposing that Unrestricted Welfare is true, UP provides the tools for solving the Non-identity Problem, whereas RP does not. To see that this is the case, let's turn to the immediate

¹³⁹ C is better than B by both an average utility principle and a total utility principle. The average utility for the actual people given C is higher than the average utility for the actual people given B. The total utility for the actual people given C is also higher than the total utility for the actual people given B.

version of the Non-identity Problem presented several chapters back:

Wilma: Wilma's doctor tells her that she has a condition that will cause any child she conceives to suffer from a serious disability if she conceives now. The child will be incurably blind. The doctor also tells her that this result is not unavoidable. If she takes a pill every day for 2 months prior to conceiving a child, then the child she conceives will be fully sighted. She decides not to take the pill for 2 months prior to conceiving, and, as a result, her child is born incurably blind. She names it Pebbles. Note that, had she opted to take the pill for 2 months, the child she would have conceived would have been different from the one she actually conceives, since the sperm and egg that would have merged to create the child would almost certainly have been distinct from the ones that merged to create her actual child.¹⁴⁰

Let's suppose that Pebbles, if she comes to exist, will live for 70 years and enjoy 50 total units of welfare. Let's suppose that the sighted child that Wilma would have had, let's call this child "Rocks", would have lived for 70 years and enjoyed 100 total units of welfare had he come to exist. Lastly, let's suppose that Wilma, because of the minor inconvenience she would endure if she had Rocks instead of Pebbles, enjoys a welfare level 100 if she has Pebbles and a welfare level 99 if she has Rocks (Wilma does not have any preference either for having a sighted child or for having a blind child, but she slightly prefers not to take the pills for 2 months). Let's call

¹⁴⁰ This case is modified from Hanser (2009, pp. 179-180).

the outcome in which Pebbles exists “O1*” and the outcome in which Rocks exists “O2”. The “*” next to O1 signifies that it is the actual outcome.

O1*: Pebbles (50) Wilma (100)

O2: Rocks (100) Wilma (99)

Unrestricted Pro tanto entails that, when determining some act’s moral permissibility, we must first aggregate the welfare enjoyed by the actual people in the actual outcome given some action and then compare that total to the total welfare of all of the actual people given the alternative outcome(s). This means that when comparing O1* and O2 we must aggregate the welfare of Pebbles and Wilma in O1* and compare it with what the welfare of Rocks and Wilma would have been had Wilma conceived Rocks instead of Pebbles. When we aggregate welfare this way the total welfare of O1* comes to 150, and the total welfare of O2 comes to 199. The actual people given O2 are better off than the actual people given O1*.

Let’s suppose for the sake of simplicity that Wilma’s welfare level in some outcome is directly proportionate to the strength of the moral reasons she has to promote it.¹⁴¹ Let us suppose the same about Pebbles and Rocks. A welfare level of 1 corresponds to a moral

¹⁴¹ For now I am setting aside the complication that many might think we can never have a moral requirement to promote our own welfare. I can never be morally required to choose the best thing for myself if such a choice would have no further effects on other people.

“reasons-strength” of 1. If we suppose this relationship between moral reasons and welfare level, then it straightforwardly follows that Wilma has stronger moral reasons to choose O2 than to choose O1*. The total welfare of Wilma and Pebbles is 150 and it corresponds to a reasons-strength of 150. The total welfare of Wilma and Rocks is 199 and it corresponds to a reasons-strength of 199. Thus, according to Unrestricted Pro tanto, again supposing the truth of Unrestricted Welfare, the welfare enjoyed by Rocks and Wilma in O2 generates a moral requirement for Wilma to bring it about in the absence of sufficient countervailing reasons. Note that even if we discount the moral importance of the welfare of Pebbles and Rocks, perhaps because (unlike Wilma) they exist in only one of the two alternative outcomes, we can still arrive at the result that the balance of moral reasons that derive from person-affecting goods counts in favor of choosing O2 rather than O1*.

Unrestricted Pro tanto does not, by itself entail that Wilma does something *wrong* by choosing O1*, because it says nothing about what kinds of reasons count as sufficient countervailing reasons. This is an issue to be addressed in a later section. UP nonetheless lays the foundation of a solution to the Non-identity Problem, because it explains what is morally objectionable about Wilma’s choice to conceive Pebbles. She fails to bring about the outcome in which the set of individuals who exist are better off than the set of individuals who would have otherwise existed.

Restricted Pro tanto, on the other hand, tells us that there is a moral reason of a requirement-generating sort to promote the good only if some actual individual will otherwise be worse off. This means that even though O2 is better than O1* (because I have argued that even a defender of RP should concede that O2 would contain more good than O1* were it to obtain),

this fact does not matter morally when it comes to evaluating the moral *permissibility* of the choice between O1* and O2. Suppose again that O1* is the actual outcome. Wilma and Pebbles are the actual people given O1*, so, according to Restricted Pro tanto, only their interests matter for the purposes of morally evaluating whether or not Wilma acts wrongly by conceiving Pebbles. Rocks is a merely possible person at the moment of evaluating Wilma's action, RP tells us to ignore how he would fare in O2. RP thus tells us first to calculate the *total welfare of the actual people* given O1* and then ask would any actual people be *worse off* given O2. If no actual people are worse off in O1* than O2, then RP generates no moral objection of the moral-requirement-generating sort to choosing O1*. Adding together the welfare of Wilma (an actual person with welfare level 100) and Pebbles (an actual person with welfare level 50), we get a total of 150 units of welfare. We must then ask whether, given these numbers, O1* is worse for any actual person when compared with O2. Wilma is the only actual person who exists in both O1* and O2, and her welfare in O2 is 99. Pebbles does not exist in O2, so she is not an actual person (even though she might be worse off) in O2. Thus there is no actual person for whom O1* is worse when we compare it with O2. Of course, Restricted Pro tanto concedes that O1* is worse than O2 because the set of actual people is better off in O2 than in O1*. Nonetheless, RP entails that Wilma's choosing O1* over O2 is morally permissible, because O1* is not worse for any of the actual people. Restricted Pro tanto thus leaves us without any helpful tools to explain why Wilma acts wrongly by conceiving Pebbles rather than Rocks. It cannot be used to solve the Non-identity Problem.

In Chapter V I argue that we ought to accept UP rather than RP. The discussion of these issues in later sections will only be sufficiently motivated, however, if it can be shown that we should accept the 'modally neutral' versions of Welfare and Pro tanto, mentioned at the

beginning of the chapter, in the first place. I now turn to defending these principles.


IV.3 The Good of Welfare

We have seen that there are at least two ways of formulating Welfare (UW and RW) and Pro tanto (RP and UP). I have already shown that *if* we accept Welfare, then we should accept UW because RW entails a contradiction. I defend UP in Chapter V. A defense of these more specific principles will only succeed, however, if we accept their more general formulations in the first place; if we think that person-affecting goods do not constitute the good and we think that there is *never* a moral reason to promote the good, where the good refers at least to the welfare of individuals, then the debate between RP and UP will be unmotivated in the first place. Though neither principle explicitly mentions the good, each implicitly depends on the idea that there is a moral reason to promote the good where the good includes (at least) the collection of person-affecting goods. Thus the first major step in defending my solution to the Non-identity Problem is to show that Welfare and Pro tanto (in their ‘modally neutral’ formulations) are central components of ordinary morality. Before my lengthier defense of Pro tanto, let us briefly reconsider some reasons to accept Welfare.

Welfare: The welfare of individuals is one of the things that constitutes the good.

One might wonder why we should think that Welfare is a plausible principle of ordinary morality, but there are good reasons to accept it. Imagine two ways a world might be:

Figure 2



Let's suppose that population y consists of one thousand individuals. Population y consists of the same members in outcomes A and B. In other words, all and only the individuals who exist in A exist in B. In A, every individual in population y enjoys a welfare level of 20. In B, every individual in population y enjoys a welfare level of 100. B is clearly a better outcome than A, and I have already argued in IV.1 that, *if* B is better than A, this fact is best explained by the further fact that it contains more person-affecting good for population y . Therefore, the judgment that B is better than A is strong support for Welfare.

IV.4 Pro tanto

Individual welfare or person-affecting goods seem to be one of the things that constitute the good. Even so, one might wonder if there is any moral reason to *promote* the good so construed. I think there is such a reason so I think that Pro tanto is true.¹⁴² I will follow Kagan in this section by arguing that accepting Pro tanto is important for preserving many of the ordinary moral judgments that most will be loath to give up. Let us begin with his case of the drowning child.¹⁴³

You come upon a child floundering in a shallow pond. If you continue on your way the child will surely drown. If you wade into the shallow pond and pull him out, thereby soaking your shoes and pants, he will survive. Most will react that you are morally required to wade in and save the child despite the minor cost to you of doing so. If you are morally required to save the child, then this requirement is perhaps best accounted for by the fact that there is a moral reason to promote the good. After all, by saving the child you are promoting his welfare, and, I have argued, the good *just is* the welfare of individuals.

It might be objected that in the case of the drowning child your moral requirement to save the child is due to a general moral reason to prevent suffering, but that the prevention of suffering is importantly distinct from the promotion of the good. The case of the drowning child establishes that there is a moral reason to prevent suffering, but it establishes nothing like a moral reason to promote the good, say, by giving your friend

¹⁴² See Shelly Kagan (1989: 16-19) for a more detailed discussion of the *pro tanto* reason.

¹⁴³ I borrow heavily from pp. 47-70 of Kagan's (1989). Peter Singer (1972) originally discusses the example of the drowning child.

\$100 to purchase a new television.

One trouble for this objection is that, although it seems plausible that one of your reasons to save the drowning child is that you would be preventing his suffering, this does not appear to be the whole story. Another reason that you are morally required to save the drowning child is that you would be bringing about more of what is *good for him* by preventing his death, thereby extending his life. The fact that you can, at little cost to yourself, cause it to be the case that he enjoys more welfare over the course of his life than he otherwise would have had he drowned generates a separate moral reason for you to ensure that he does continue to enjoy those things. To see that this is part of the explanation of your having a moral reason to save the child, imagine that the drowning child is unconscious in the pond, because, say, a rock had fallen on his head rendering him unconscious and knocking him into the pond. By saving him you are not preventing any suffering that he might endure (he is not conscious to suffer any pain by drowning), you are simply ensuring that he enjoys more goods by extending his life. The case of the drowning child seems to support Pro tanto even if we are careful to distinguish between the moral reason to prevent suffering and the moral reason to promote the good.

A separate moral distinction might seem to undermine the support for Pro tanto provided by the case of the drowning child. The case of the drowning child might fail to support Pro tanto if it turned out that while there is a moral reason to *prevent the loss* of some good, there is no moral reason to *cause the addition* of some good. The fact that you are morally required to save the drowning child might be explained by the fact that you have a moral reason to prevent the loss of the goods that he would enjoy were he not to

drown. You prevent the loss of some good (by preventing the loss of a longer life that involves more pleasurable experiences) by saving the child from drowning. Thus the case fails to establish that there is a further moral reason to cause the addition of some good. In the case of your friend you merely *cause the addition* of some good by giving him \$100 to purchase a television, and, the opponent of Pro tanto will contend, there is no moral reason to bring about the mere addition of some good.¹⁴⁴

There seem to be two major problems for the attempt to ground the moral requirement to save the drowning child in a moral reason to prevent the loss of a good while simultaneously denying the existence of a moral reason to cause the addition of a good. The first major problem for maintaining the moral distinction between preventing the loss of a good and causing the addition of a good is that it is not clear that there is any genuine metaphysical distinction to ground the moral distinction in the first place. If there turns out not to be any metaphysical difference between preventing the loss of a good and causing the addition of a good (at least no difference to be found in the case of the drowning child), then there cannot be a moral difference that is grounded in the metaphysical distinction. Let us ask, then, in virtue of what is saving the drowning child properly described as *preventing the loss* of some good but not properly described as *causing the addition* of some good.

It seems that losing a good is a temporal notion. It involves comparing the amount of a good that someone enjoys in one state and then comparing this quantity with the quantity they *later* enjoy. So, one plausible way to flesh out the distinction between

¹⁴⁴ Kagan (1989, pp. 54-55) describes this moral distinction. He finds it insufficiently motivated.

preventing the loss of a good and causing the addition of a good is by taking into consideration someone's present state of well-being and asking what would have happened to the person's state of well-being had an agent not intervened.

Prevention: An agent A prevents a loss of a good if and only if a person P is currently enjoying some non-instrumental benefit B, and P would have ceased to enjoy B if A had not intervened.

Addition: An agent A causes the addition of a good if and only if a person P is not currently enjoying some non-instrumental benefit B, and P would commence to enjoy B only if A were to intervene.

The conjunction of Prevention and Addition seems like an intuitive way to flesh out the distinction between preventing of a loss and causing the addition of a good, since it seems to correspond nicely with our ordinary judgments about certain cases.

Adam 1: Adam is eating a tasty apple when Eve attempts to tear the apple from his hands. She would have succeeded had I not stopped her from doing so. Adam continues to enjoy his apple.

Adam 2: Adam is hungry for something refreshing. I pick an apple from a nearby tree and deliver it to him. He eagerly takes a bite.

In *Adam 1* it intuitively seems that I prevent Adam from losing a good. In *Adam 2* it seems that I cause Adam to enjoy an additional good. The analyses just offered can account for these judgments. In the former Adam is currently enjoying some benefit (the taste of a juicy apple) and my intervention prevents him from losing that benefit. In the latter Adam is not currently enjoying some benefit and it is only through my intervention that he comes to enjoy a benefit. This plausible way of distinguishing between preventing the loss and causing the addition of a good does not, however, clearly vindicate the judgment that I prevent the loss of a good in the case of the drowning child. When I come upon the drowning child, he is not currently enjoying some benefit (since he is unconscious when I arrive upon the scene). On a couple of plausible theories of welfare, his welfare level is most accurately described as being at the zero level. We can suppose that he is not enjoying any pleasure, that he is not suffering any pain, and that none of his desires are being satisfied or frustrated. Moreover, he will only come to enjoy future goods (e.g., future pleasurable experiences that accompany a longer life) if I intervene, saving him from drowning. If this description is correct then it would seem that, if we accept Addition and Prevention, then we must accept that I cause the addition of a good but do not prevent the loss of a good by saving the drowning child. My moral reason to save the child derives from a moral reason to cause the addition of a good after all.

My opponent might contend that I am ignoring a benefit that the child is currently

enjoying as he lies unconscious at the bottom of the pond; a benefit that he will continue to enjoy only if I intervene. That is, I am ignoring the fact that the child is enjoying the benefit of *being alive* at the moment I save him, and I prevent the loss of this benefit by intervening. Here it is helpful to distinguish between non-instrumental and instrumental goods. An instrumental good is a good that is good for you only insofar as it is appropriately related to certain non-instrumental goods. Money is good for someone only insofar as it allows him to purchase the things that are required for a happy life. It is not good in and of itself or regardless of its further consequences. Imagine that I were stranded on top of Mount Everest with a suitcase filled with a million dollars in US currency. This money would in no way make my life better, because it would be *of no use to me* given the circumstances. Furthermore, in these bizarre circumstances, it might seem that you would have *no* moral reason not to take it from me. Non-instrumental goods, I have already explained, are goods that are good for someone regardless of their further consequences for that person. Pleasure or the satisfaction of one's desires might be examples of non-instrumental goods. They are good for a person regardless of what else accompanies them. It seems that life is merely an instrumental good, because it is worth having insofar as it is appropriately related to the things that are non-instrumentally good. To see that this is so, consider the case of someone in a vegetative state who has no chance of conscious revival. This person has no chance of having experiences, learning new information, forming relationships. Is it in any way better for her to be alive than to be dead, given that her life provides no chance of securing the benefits listed above? I suspect that most will think that it is not better to be alive in these circumstances than dead. This is because being alive is *useless to her* in these circumstances. Life, then, seems to be a merely instrumental good.

If I am correct that life is instrumentally but not non-instrumentally valuable, then there is a simple response to my opponent who contends that the drowning child is currently enjoying the benefit of being alive at the moment I save him and that I prevent him from losing this benefit. The proper reply is that the ‘benefit’ of being alive, if it is valuable, is of merely instrumental value. Since it is merely instrumentally valuable (valuable only insofar as it is related to something non-instrumentally valuable) we can ask if it is, at the moment I save the drowning child, appropriately related to anything of non-instrumental value. Since the child is unconscious (so he is not enjoying anything of non-instrumental value), and since the child’s life will not yield any further satisfaction of desires or pleasurable experiences in the absence of my intervention, it seems safe to say that the child’s life is devoid of any instrumental value at the moment I save him. His life is useless to him and it thus does not constitute a *benefit* that he is enjoying. Thus, saving the child from drowning does not prevent the loss of a good if the good in question is the good of being alive. Appealing to the good of *being alive* will not suffice to preserve my opponent’s judgment that my obligation to save the drowning child stems from a moral reason to prevent the loss of a good but not to cause the addition of a good.¹⁴⁵

¹⁴⁵ Interestingly, the claim that I cause the addition of a good by saving the drowning child is consistent with maintaining, as seems intuitively correct, that the child will be harmed should he drown in the pond. According to the version of the comparative account of harm that I defend in Chapter II, a person is harmed by an action or event just in case the person is caused to be worse off as a result of that action or event than he otherwise would have been. A person is worse off as a result of some action or event just in case the total value of s’s life is lower as a result of that action or event than the total value of s’s life would have been in the absence of that action or event. If the event of the child falling into the pond results in his drowning, then his life will contain less value than his life will if he is saved. This is because it will be significantly shorter and so will include fewer of the non-instrumental goods that make lives go well. By preventing his drowning I prevent him from being harmed in this way, but not because I prevent the potentially harmful event from occurring (his falling into the pond). I prevent his falling into the pond from harming the child by pulling him from the water, while *simultaneously* causing the addition of many goods that he could only enjoy through my intervention.

Let's suppose that my remarks in the previous section are mistaken. Suppose that there does turn out to be a genuine metaphysical distinction between preventing the loss of a good and causing the addition of a good, and suppose that the fact that you prevent a loss of a good by saving the drowning child is what generates a moral reason for you to save him. Even supposing all of this, there is a further problem for the view that, while there is a moral reason to prevent the loss of a good, there is *no* moral reason to cause the addition of a good. The problem is that maintaining such a moral distinction makes it impossible to account for another ordinary judgment that we are loath to abandon.

Suppose that I have a friend Tom who was born without the use of his legs. The disability is a significant one, but it is painless. It has been bad for him only insofar as it has caused him to enjoy significantly fewer pleasures (or desire-satisfactions or achievements) than he would have enjoyed with the use of his legs. Due to a medical miracle, doctors have discovered that if Tom goes to the hospital in the next 30 minutes they can perform a procedure that will result in the full functioning of his legs. I am the only one who can get him to the hospital in time. By taking him to the hospital I merely cause the addition of some good for Tom, because I merely enable him to enjoy the pleasures that he can only enjoy with the use of his legs.¹⁴⁶ Now imagine that I have a friend, Jerry, whom I find unconscious in his vehicle in the middle of an intersection after colliding with a drunk driver. He is badly injured, but not suffering any pain from the accident (since he is unconscious). If I extract him from the vehicle and take him to the hospital in the next 30

¹⁴⁶ We have already considered (Chapter II) an account of harm according to which being in a state of disability constitutes a harmful state, and we have found the account to be unacceptable. Tom's state is not plausibly described as harmful, so you cannot, by benefitting him, properly be described as preventing or alleviating harm.

minutes I can prevent him from losing the use of his legs, thereby preventing him from losing out on the non-instrumental goods he would only enjoy with the use of his legs. I am the only one who can help him. By taking him to the hospital in the next 30 minutes I *prevent* the *loss* of the goods he would enjoy through the continued use of his legs. If one only has a moral reason to prevent harm by preventing the loss of goods, but no moral reason to promote the good, where promoting the good consists in *causing the addition* of some good, then, though I have strong moral reasons to take Jerry to the hospital, I have no moral reason at all to take Tom to the hospital. Since most will be loath to accept this implication, most ought to reject the view that the moral reason to promote the good includes a moral reason to prevent the loss of a good but *no* moral reason at all to cause the addition of a good. Even if the moral requirement to save the drowning child can be accounted for solely by appealing to a moral reason to prevent the loss of a good, we can account for the moral requirement to take Tom to the hospital only if we accept that there is a moral reason to promote the good where promoting the good is construed as causing the addition of a good.¹⁴⁷ We get this result even if, as some will find plausible, the reason to promote the good by causing an addition of a good is weaker than the reason to prevent the loss of a good.

What, then, should we conclude about the nature of my reason to give my friend \$100 so that he can purchase a new television? My discussion indicates that we should

¹⁴⁷ According to Elizabeth Harman's non-comparative account of harm (discussed in Chapter II), Tom's being disabled constitutes a state that it is harmful for him to be in. Her view is thus consistent with the possibility that your reason to take Tom to the hospital is generated by the fact that doing so would improve his condition from a harmed state to a non-harmed state rather than the fact that doing so would promote the good. I have already rejected Harman's account of harm, however, so it cannot be used to account for your reason to benefit Tom.

conclude that, though my moral reason to cause the addition of some good by giving my friend \$100 to buy a new television might be so weak that it fails to generate a moral *requirement* in most cases (perhaps because the benefit to him would be quite small and the cost to me quite great), I nonetheless have *some* moral reason to give him \$100. This is because the moral difference between my friend who wants a new television and Jerry who wants to enjoy the use of his legs is a difference in degree, not a difference in kind. In both cases I can cause the addition of some good, and in both cases this generates some moral reason for me to do so.¹⁴⁸

IV.5 The Moral Relevance of Possible People? Choosing between Unrestricted Welfare and Unrestricted Pro tanto

In section IV.3 I argued that Welfare and Pro tanto are plausible principles of ordinary morality. If those principles alone could solve the Non-identity Problem, then my work would be almost complete. Welfare and Pro tanto do not obviously solve the Non-identity Problem, however, because the content of the principles is ambiguous. While I think it is plausible that if we accept Welfare than we should accept UW's claim that the welfare of possible people would constitute the good were it to be actualized, I have not yet argued that *if* we accept Pro tanto *then* we should accept UP. Pro tanto, on its own leaves it unclear *which* individuals' welfare generates

¹⁴⁸ One might take the case of Tom to establish only that there is a moral reason to cause the addition of a good if that good is sufficiently important. On this view there are certain thresholds that a benefit must meet before it generates any moral reason for an agent to bring it about. This would render the judgment that I have a moral obligation to cause the addition of a good by enabling Tom to use his legs consistent with the judgment that I have no moral reason whatsoever to benefit my friend by giving him \$100 to buy a new television. I find this view curious as it seems to require that we reject the transitivity of benefits, and this is an implication that most should be hesitant to accept. See Norcross (1997) for a discussion of the transitivity of harms and benefits.

a moral reason to promote it. RP cannot be used to solve the Non-identity Problem but UP can. Whether Welfare and Pro tanto lead us to a solution to the Non-identity Problem will depend on which formulation we accept.

I have already argued that if we accept Welfare, then we should accept UW, because the alternative principle, RW, entails a contradiction. I have not defended UP, however, and the cases to which I appealed in support of Pro tanto do not exclusively support UP. Of course, UP supports the ordinary judgments about the drowning child and Jerry. It entails that you have a strong moral reason to save the child and to restore the use of Jerry's legs because, in each case, you thereby ensure that the actual people are as well off as possible. But RP also supports our judgments about the drowning child and Jerry. It might be that the best explanation of your moral requirement to save the drowning child or of restoring the use of Jerry's legs is that you are preventing some actual person from being *worse off* than *that person* otherwise would have been. It is because Jerry exists and because, in outcomes in which you fail to restore the use of his legs, he also exists and is worse off, that his welfare has some claim on you. If that's right, then your moral reason to aid Jerry does not derive from the fact that aiding him would promote person-affecting goods generally but from the fact that doing so would promote the person-affecting goods of individuals who will exist regardless of what you choose to do. According to RP, it is the fact that Jerry, an *actual person*, would otherwise be a *worse off* actual person that makes it the case that you have a moral reason to promote his welfare. This is the rationale behind RP. Our judgments about the drowning child and Jerry are thus equally supportive of UP and RP, and cannot help us to choose between them.

My aim has been to defend Welfare and Pro tanto thereby providing sufficient motivation

for the importance of the debate between UP and RP. Supposing that I have succeeded in accomplishing this aim, in the following chapter I turn to the debate between UP and RP, arguing that RP has unacceptable implications while defending UP against several powerful objections. I then provide an argument that uses UP to solve the Non-identity Problem.

V

In Defense of Unrestricted Pro tanto: A Solution to the Non-identity Problem

In the previous chapter I defended the Welfare principle, arguing that welfare, in the form of comparative and non-comparative person affecting goods, constitutes the good. I also defended the Pro tanto principle, arguing that there is (at least sometimes) a moral reason to promote the good where the good includes at least individual welfare in the form of person-affecting goods. I then distinguished between two ways of limiting the scope of Pro tanto. According to Restricted Pro tanto, there can be a moral requirement to perform some action only if some actual person would be worse off if the action were not performed. According to Unrestricted Pro tanto, there are moral reasons of the moral-requirement generating sort to perform actions that would make it the case that the set of actual individuals are better off than the set of actual individuals otherwise would have been. It is worth considering the principles again to start this chapter:

Restricted Pro tanto (RP): There is a moral reason (not) to perform an action a if the set of people given A will be better (worse) off given A than the set of people given $\sim A$, but this moral reason generates a moral requirement in the absence of sufficient countervailing reasons **only if some actual person would be worse off** given $\sim A$ than

given A.¹⁴⁹

Unrestricted Pro tanto (UP): There is a moral reason (not) to perform an action *a* if the set of individuals given *A* would be better (worse) off than the set of individuals would be given $\sim A$, and **this moral reason generates a moral requirement** in the absence of sufficient countervailing reasons.

Each of these principles is well supported by my arguments in the previous chapter, but they are mutually exclusive. One of them must be rejected. In the previous chapter I showed that only UP is capable of solving the Non-identity Problem. Because my aim in this chapter is to solve the Non-identity Problem, in section V.1 I begin my defense of UP, raising several objections to RP. In V.2 I flesh out the heretofore under described parts of UP by providing an analysis of sufficient countervailing reasons. Providing this analysis is important for the purposes of demonstrating that UP avoids the pitfalls of Maximizing Consequentialism. In V.3 I offer the ‘Welfare Argument’ as a proposal for a solution to the Non-identity Problem. This argument combines UP with some plausible supplementary premises to yield the conclusions that Wilma acts wrongly by conceiving Pebbles rather than Rocks and that the wealthy society acts wrongly by choosing depletion rather than conservation. I conclude the chapter in section V.4 by considering several objections to my solution to the Non-identity Problem. I argue that none of

¹⁴⁹ This principle does not rule out the possibility that there could be a moral reason to promote the good in cases where failing to do so would leave no actual person worse off. It simply rules out the possibility that this moral reason could generate a moral *requirement*.

these objections prove successful.

V.1 Problems for RP

At this point one might wonder what reasons there are, if any, for preferring UP to RP. In this section I argue that there are several powerful reasons to accept UP rather than RP. I offer two kinds of cases involving the creation of additional people. Our judgments about these cases can be easily accommodated if we accept UP but not if we accept RP. I then consider a couple of examples offered by Derek Parfit that are designed to support what he calls the ‘No Difference View’.¹⁵⁰ UP provides a tidy, plausible explanation for ordinary judgments about Parfit’s cases, whereas RP’s explanation yields deeply problematic implications. For these reasons, I conclude that UP is a more plausible moral principle than RP.

V.1.1 Creation Case 1: Happy Lives

The first kind of problematic case for Restricted Pro tanto involves the creation of happy lives. If we accept Restricted Pro tanto, then we must accept that there is no number of additional happy people, the existence of which depends on an agent performing some action rather than another, which could make it the case that the agent is morally required to perform that action. Consider:

¹⁵⁰ E.g., Parfit (1984: 367; 2011: 219).

Alexandre: Alexandre is thirsty. An omniscient, omnibenevolent, but non-omnipotent being appears before him. The being tells Alexandre that he can perform one of two actions to quench his thirst, each of which will bring about a different state of affairs. He can (A) pick up the glass of water sitting next to him and drink from that glass, or he can (B) walk across the room with an empty glass of water, fill that water glass, and then drink from it. (B) will constitute a minor inconvenience for Alexandre. If Alexandre chooses (A) then he will cause it to be the case that the current generation of human beings is the last ever to exist. The trillions of individuals who would have come to exist in the future to enjoy lives worth living will not come to exist. Let's call this population of happy individuals whose existence depends on Alexandre's choice 'x.' If he chooses (B), however, then x will come to exist and will live their happy lives. He also knows, because he accepts UW, which states that the welfare actual individuals contributes value to states of affairs, that (B) is better than (A). We can represent the case diagrammatically as follows:

(A): Alexandre (100)

(B): Alexandre (99) x (n people at welfare level 100)

Even if we stipulate that no actual person is caused to be worse off by Alexandre choosing (A) (we can suppose that none of the people who currently exist will be worse off if they are prevented from conceiving children), it is hard to believe that Alexandre does not do

something morally wrong if he chooses (A), thereby preventing the coming into existence of *x*. UP implies, plausibly, that there is a serious moral objection to Alexandre's choice of (A) over (B), because the collection of actual individuals in (B) is significantly better off, on a whole, than the collection of actual individuals in (A). But the verdicts of Restricted Pro tanto conflict with this ordinary intuitive judgment. Because there is no actual person for whom (A) is worse, Restricted Pro tanto entails that the welfare of trillions of people who will never exist because he chose (A) *does not matter* for the purposes of evaluating the permissibility of his choice. Because the trillions of people will never exist given that he chooses (A) for the sake of avoiding a minor inconvenience, the (now merely possible) welfare they would have enjoyed does not generate moral reasons of the requirement-generating sort for him to choose (B) rather than (A). How Alexandre (setting aside the other people who currently exist) fares in (B) relative to (A) is the only morally relevant consideration for the purposes of evaluating his choice, according to RP, and Alexandre fares better in (A) than in (B). Therefore, accepting RP generates no objection to Alexandre's choosing (A) rather than (B). But this is hard to accept.

It is worth emphasizing an even more peculiar implication of RP regarding Alexandre's choice. RP implies that (i) Alexandre's choice of (A) is morally permissible despite the fact that (ii) the defender of RP must agree that (B) is the overwhelmingly better choice from the standpoint of the good *and* thus (iii) that the overwhelming majority of Alexandre's moral reasons favor choosing (B). This first creation case thus brings out a counterintuitive implication of RP, but it also demonstrates a peculiar tension embodied by the principle: that one can fail to

be morally required to do what one has overwhelming moral reasons to do.¹⁵¹

*V.1.2 Creation Case 2: The Problem of Suffering*¹⁵²

The case of Alexandre seems rather problematic for Restricted Pro tanto, but the proponent of the principle might be willing to bite the bullet and accept the implication that Alexandre does not act wrongly by choosing A rather than B. Let us see, then, what the principle implies about a second kind of creation case. Consider

James: James is thirsty. An omniscient, omnibenevolent, though non-omnipotent being appears before him and tells him that he can do one of two things, each of which will bring about a different outcome. He can either (C) pick up the glass of water sitting next to him and drink from that glass, or (D) he can walk across the room with an empty glass of water, fill that water glass, and then drink from it (a minor inconvenience). Let's now suppose that some lives can be 'miserable' in the sense that intrinsic bads far outweigh the goods. The being tells James that if he chooses (C) then he will cause it to be the case that many future generations of human beings will come to exist who will have miserable lives in which the suffering heavily outweighs the enjoyment. Let's call the set of these miserable people *y*. If he chooses (D) then *y* will never come to exist. We can picture the

¹⁵¹ Someone who denies that moral reasons generate moral requirements might be happy to accept this kind of implication. See Norcross (2006) for a defense of a 'scalar consequentialist' view of this sort. The defender of RP, however, believes that some moral reasons generate moral requirements. He simply denies that certain kinds of moral reasons generate moral requirements.

¹⁵² As noted in the last chapter, I borrow this example and its name from Holtug (1998).

case diagrammatically as follows:

(C): James (100) y (n people with welfare level -100)

(D): James (99)

Surely most will agree that James would act wrongly by choosing (C) rather than (D). The outcome corresponding to (C) is clearly worse than the outcome corresponding to (D). This much is entailed by Unrestricted Welfare, a principle that I've already shown we should accept. Unrestricted Pro tanto can easily take this fact into account in order to accommodate our judgment that James does something morally objectionable by choosing (C) rather than (D). The actual people if James chooses (C) will suffer an exceedingly negative welfare level whereas the actual people if he chooses (D), though they will be non-identical with the people in (C), will enjoy an on balance positive level of welfare. Unrestricted Pro tanto tells us to compare the combined welfare of the actual individuals at each of the two outcomes. If the set of actual people in (C) are worse off than the set of actual people in (D) then UP entails that there is a moral reason of the requirement-generating sort to choose (D) rather than (C). The fact that y is a merely possible population given that James chooses (D) does not preclude the welfare of its members from entering into the moral evaluation of his choice. Because the set of actual people fare better, given (D), than the set of actual people fare, given (C), UP entails that James has most moral reason to choose (D) and that James is morally required to choose (D) in the absence

of sufficient countervailing reasons.¹⁵³ This is the intuitively correct verdict about the case.

RP does not clearly entail that James is required to choose (D), but, with a small supplement, I think the view also generates this intuitively correct verdict. Restricted Welfare asks whether any actual person is worse off in (D) than in (C). Well, James is slightly worse off in (D), so RP generates some small objection to his choosing (D). But, RP also asks whether any actual person would be worse off if James chooses (C). If James chooses (C) then trillions of actual people would live miserable lives. Let us suppose for the moment that this means that their existence would be *worse* for them than non-existence. Provided this supposition, which I will return to shortly, it seems plausible that (C) is much worse than (D) for some of the actual people. In fact, it is much worse than (D) for a lot of them. Both (D) and (C) turn out to be worse for some of the actual people, so RP implies that there is something objectionable about each choice. But this does not, on its own, tell us which if either of the options is morally permissible or required. RP needs to be supplemented to morally adjudicate between choices, each of which will be worse of some of the actual people. We might add the following supplement:

RPT: In cases where each outcome will be worse for some of the actual people, RP requires us, in the absence of sufficient countervailing reasons, to choose the outcome that will be *least* worse for the actual people.

¹⁵³ I will later question the assumption that an agent's own welfare can generate moral reasons of the requirement-generating sort to act one way rather or the other. For now, I set this complication aside.

No matter what James chooses, his choice will be worse for some actual person than the alternative would have been. The outcome that obtains if James chooses (D) is slightly worse for James than (C). The outcome that would obtain were James to choose (C) would be worse for the trillions of suffering people who *would be actual people* were he to choose (C). It seems then that no matter what James chooses what he does will be worse from the standpoint of the interests of some actual person. Since Restricted Pro tanto tells us that we have a moral reason of the requirement-generating sort to refrain from performing actions that would be worse for some actual person, and since the outcome that results from either (C) or (D) would be worse for some set of actual people, RPT says that James should make the choice that is the *least worse* from the standpoint of the interests of actual people. (D) is the least worse for the actual people, because James is only slightly worse off in (D) than in (C), and population *y* is much worse off in (C) than in (D). RP combined with RPT thus entails that James has most moral reason to choose (D) and that this fact generates a moral requirement for James to choose (D) over (C) in the absence of sufficient countervailing reasons. Despite initial appearances, the case of James gives us no reason to prefer Unrestricted Pro tanto over Restricted Pro tanto, because both principles entail that James ought to choose (D).

A defender of RP who offers this reply faces a couple of serious challenges. The first arises from the observation that the sense in which (D) is worse than (C) for James from the standpoint of his interests is quite different from the sense in which (C) is worse than (D) from the standpoint of the interests of the trillions of sufferers. James *exists* in both outcomes, and one of the outcomes involves James possessing a lower welfare level than the other, whereas *y*, because it only exists in one of the outcomes, only possesses a *welfare level* in one of the outcomes. As presently formulated, however, in order to ensure that the principle gets the right

verdict in the case of James the proponent of RP must maintain that the outcome in which *y* exists is *worse for y* than the outcome in which *y* does not exist. But there are good reasons to doubt that such comparisons are coherent.

V.1.2.A The Comparison Problem

We ordinarily talk as if existence can be worse than non-existence. For instance, we routinely say things like ‘he would be better off dead,’ but there is a serious worry for the view that we can compare how someone fares when she exists with how ‘she’ would have fared were ‘she’ *never* to have existed. The worry is that someone who never exists at a world does not have any welfare-determining properties at that world. Establishing that non-existence is better from the standpoint of the interests of a person who actually never exists seems to require that we compare ‘her’ welfare level in the actual world with what her welfare level would have been in the world at which she exists. But if someone does not have properties at the actual world because ‘she’ never exists at the actual world, therefore failing to possess any welfare-determining properties, then ‘she’ does not have a welfare level at that world with which to compare her welfare level at a world where she does exist. In John Broome’s words:

[I]t cannot ever be *true* that it is better for a person that she lives than that she should never have lived at all. If it were better for a person that she lives than that she should never have lived at all, then if she had never lived at all, that would have been worse for her than if she had lived. But if

she had never lived at all, there would have been no her for it to be worse for, so it could not have been worse for her.¹⁵⁴

It might seem that the proponent of RP can avoid this problem of comparisons by appealing to a ‘whole lives’ version of the counterfactual comparative account of harm that I discussed in II.3. In that section I invoked the notion of a person’s ‘life as a whole’ going better for him than it otherwise would have gone in order to explain why death is harmful to the one who dies. Death, I claimed, makes a person’s life go worse as a whole, because it cuts it short, thereby preventing that life from containing more value. But this ‘whole lives’ explanation of the harmfulness of death will not help the defender of RP explain why coming to exist with a miserable life is worse than non-existence, for, at worlds where ‘a person’ *never* comes to exist there is *no life* that ‘the person’ has that is better than the life that the person would have had were ‘she’ to have come to exist. Appealing to the contribution of non-existence on the quality of a person’s life as a whole is thus a non-starter as a reply to the coherent comparison problem. (Note that UP does not run into the incoherent comparisons problem. When determining whether there is moral requirement to choose (D) over (C), UP does not ask whether some actual person is worse off than *she* otherwise would have been. UP asks whether the people who exist in (C) are worse off than the people who exist in (D), and this is very different matter. To answer this question, we need only determine the total welfare that the people who would *exist* in (D) would enjoy and then compare it to the total welfare enjoyed by the people who would *exist* in (C). UP

¹⁵⁴ Broome (1999, p. 168). See also Narveson (1967, p. 67), Parfit, (1984, pp. 395, 489), Heyd (1988), Bykvist (2007), and recently Herstein (2013).

does not require comparisons between how someone fares when she exists and how she fares in world where she never exists.)

One way that a proponent of RP might reply to the incoherent comparison problem is to say that non-existence is better than existence from the standpoint of a person's interests just in case it *would be rational to prefer* non-existence to existence. How do we determine whether it would be rational to prefer not to exist than to exist? It will not do to appeal to the idea that it is rational to prefer non-existence to existence just in case one's life is worse than non-existence. Alternatively, we might hold that it is rational to prefer non-existence to existence just in case, say, the intrinsic bads in one's life outweigh the intrinsic goods. The defender of RP might then revise her principle as follows: there is a moral reason of a requirement-generating sort (not) to perform some action just in case it would be *rational* for some actual person to prefer that you (not) perform the action, even if the resulting outcome would not be (worse) better, for some actual person. But this does not entirely solve the problem for the defender of RP. It seems difficult to deny that one's rational preferences track *betterness for* someone, for it can be rational to choose outcomes that are intrinsically bad for one as long as they are *better* than the alternatives. For instance, it is rational for me to prefer a 1,000 volt electric shock to a 10,000 volt shock even though both events result in extreme pain. This means that, given the rationality reply, the defender of RP cannot avoid a commitment to the claim that non-existence is *better* than existing with a life that involves more pain than pleasure. But, as we saw, it seems that someone who does not exist does not have the properties necessary to support the judgment that non-existence is better for her than existing with a miserable life. Some will be unmoved by the problem of coherent comparisons. They will say that some lives are clearly worse than never coming to exist, and our conviction about this judgment should not waiver in the face of its

apparent metaphysical problems. I find this reply unsatisfying, but it poses a dialectical impasse, so I will move on.

V.1.2.B The Asymmetry

Setting aside the issue of incoherent comparisons, the problem of suffering highlights an interesting asymmetry entailed by RP. RP treats the moral significance of individuals with happy lives very differently from individuals with miserable lives. This feature of RP is explicit in the principle, but it is tempting to overlook it. Let us assume that comparisons between existence and non-existence are coherent, contrary to my remarks above. In the case of James, (C) would be worse than (D) for the trillions of people who would exist with miserable lives. If this claim is unproblematic, then so is the claim that, in the case of Alexandre, (B) would be better than (A) for the trillions of happy individuals who would exist in (B). So the following claim is true: (B) is better than (A) for some of the actual people. But RP says that, while we can be morally required to refrain from making a choice if the choice would be *worse* for some of the actual people, we cannot be morally required to make a choice if doing so would be better for some of the actual people. This asymmetry is what is needed to generate the implications that it is morally permissible for Alexandre to choose (A) rather than (B) and that it is morally wrong for James to choose (C) rather than (D). In the first case (B) is *better* than (A) for some of the actual people but (A) is not *worse* for any actual people (in fact it is better for Alexandre). In the second case, (C) is *worse* than (D) for some of the actual people even though (D) is not *better* for any actual people (in fact it is worse for James).

A second asymmetry is embodied by RP's tie-breaking principle RPT. RPT directs us, when all of our choices are worse for some actual person, to make the choice that will be least worse for some of the actual people. The view does not have a corresponding tie-breaking principle pertaining to situations when, no matter what we do, our action will be *better* for some of the actual people. We could easily modify RP and RPT to eliminate this awkward asymmetry:

RP2: There is a moral reason (not) to perform an action *a* if the set of people given *A* will be better (worse) off given *A* than the set of people given $\sim A$, but this moral reason generates a moral requirement in the absence of sufficient countervailing reasons **only if either some actual person would be worse off** given $\sim A$ than given *A* **or some actual person would be better off** given *A* than $\sim A$.¹⁵⁵

RPT2: In cases where each outcome will be worse (better) for some of the actual people, RP tells us to choose the outcome that will be the least worse (most best) for the actual people.

Can the defender of RP and RPT justifiably rule out RP2 and RPT2? Are there good reasons to accept the asymmetries embodied by RP and RPT? One might argue that the asymmetry embodied in RP and RPT is justified on the grounds that the moral imperative not to harm is more stringent than the moral imperative to benefit. Alexandre benefits actual people by

¹⁵⁵ This principle does not rule out the possibility that there could be a moral reason to promote the good in cases where failing to do so would leave no actual person worse off. It simply rules out the possibility that this moral reason could generate a moral *requirement*.

choosing (B), but he does not harm any actual people by choosing (A). James, on the other hand, harms actual people by choosing (C) rather than (D). RP's asymmetry is justified by the fact that our moral aim should be to avoid as much comparative harm to actual people as possible, not to achieve as many comparative benefits for actual people as possible. A second but related justification might be that there is no one who exists and *has a legitimate complaint* if Alexandre chooses (A) (the trillions of people who would have been happy might have a complaint, but they don't exist), but there is someone who exists and has a legitimate complaint if James chooses (C) (the trillions of miserable people who are worse off than they otherwise would have been). According to these replies, the justification of the asymmetry embodied in RP and RPT appears to be the idea that we should avoid making it the case that *some actual person has a legitimate complaint against us*. On one way of analyzing the notion of having a legitimate complaint, someone has a legitimate complaint just in case *she is caused to be worse off than she otherwise could have been*. No actual person has a legitimate complaint against Alexandre if he chooses (A), so RP generates no objection to his choosing (A). Many actual people have a legitimate complaint against James if he chooses (C), so RP generates a strong objection to his choosing (C). Perhaps an appeal to legitimate complaints can serve as the underlying justification of the asymmetry embodied by RP and RPT. But RP yields strange results in certain choice situations if the notion of *having a legitimate complaint* is analyzed this way. Consider a choice between the following options:

(E): x (100) y (50)

(F): x (50) y (100)

(G): z (1)

(E) and (F) are equally preferable from the standpoint of the interests of the *group* including x and y, but (F) is worse for x than (E), and (E) is worse for y than (F). If someone has a legitimate complaint just in case she is caused to be worse off than she otherwise would have been through no fault of her own, then x has a legitimate complaint in (F) and y has a legitimate complaint in (E); each is worse off than she otherwise would have been. It seems that (G) is the only choice that yields no legitimate complaints on the part of an actual person, so it seems that RP would direct us to choose (G). But surely (G) is the *least* morally desirable outcome of the three options. Given the choice between (E), (F), and (G), one ought to choose either (E) or (F).

The defender of RP and RPT might revise the analysis of the notion of having a legitimate complaint presented above. He might say that someone has a legitimate complaint about your choice just in case (i) she is worse off than she otherwise would have been as a result of your choice, *and* (ii) no other actual person would have been worse off to the same degree had you chosen differently. This would yield the result that neither x nor y has a legitimate complaint about either (E) or (F). I am not sure that this analysis will suffice. A slight modification of the original case involving (E), (F), and (G) seems to pose just as deep a problem for this analysis of having a legitimate complain. If one simply changes (E) so that y has 51, y won't have a legitimate complaint in (E), but x will in (F), so (G) would still turn out to be preferable to (F) in this case. Though (E) might be slightly preferable to (F), given this modification, surely this should not yield the result that both (E) *and* (G) are preferable to (F).

Let us now suppose that I am mistaken and that the defender of RP and RPT can avoid the move to RP2 and RPT2 by emphasizing the moral significance of legitimate complaints. This allows him to retain the judgments that James acts wrongly by choosing (C) rather than (D) but

that Alexandre does not act wrongly by choosing (A) rather than (B). The most troubling problem for the defender of RP's reply to the problem of suffering is that the reply seems also to entail that RP morally requires agents to bring about the worst outcome in certain cases. The response by the defender of RP to the problem of miserable lives claims that James, given that either (C) or (D) will be worse for *some of* the actual people, ought to bring about the outcome that is *least worse* for the *total set* of actual people, and this fact explains why he does something morally objectionable by bringing about miserable population *y* in order to avoid a minor inconvenience to himself. But now suppose that Walter has the following options available to him:

- | | | | |
|-----|----------------|----------------|--------------|
| (H) | <i>x</i> (50) | <i>y</i> (100) | <i>z</i> (1) |
| (I) | <i>x</i> (100) | <i>y</i> (50) | <i>z</i> (1) |
| (J) | | | <i>z</i> (2) |

If Walter chooses (H) or (I) then populations *x*, *y*, and *z* will come to exist. If he chooses (J) then only population *z* will come to exist. According to RP's explanation of the original case involving James and the miserable population, James' moral reason not to create miserable population *y* stems from the facts that either (C) or (D) will be worse for *some of* the actual people and that (C) would be more worse for the set of actual people (James will be slightly better off but population *y* will be significantly worse off) than (D). But the kind of reason appealed to here has highly counterintuitive implications for what Walter should do in the case

above. Walter can either (H) bring x into existence with welfare 100 and y into existence with welfare 50, and z into existence with welfare 1, (I) bring x into existence with welfare 50 and y into existence with welfare 100, and z into existence with welfare 1, or (J) bring z into existence with welfare 2. (H) is worse for the actual population y than (I), since y 's welfare is 50 units lower in (H) than in (I). (I) is worse for the actual population x than (H), because x 's welfare is 50 units lower in (I) than in (H). Both (H) and (I) are worse for z than (J). (J) is the only option that is not worse for any of the actual people, because only z exists in (J), and (J)'s welfare is higher in (J) than in either (H) or (I). Most importantly, (J) is the only option in which no actual person has a legitimate complaint; it is the only option that avoids making some actual person (z) worse off without making some other (set of) actual person(s) equally worse off. RP thus implies that Walter has most moral reason of a requirement-generating sort, other things being equal, to bring about (J). It also implies that Walter is morally required to bring about (J) in the absence of sufficient countervailing reasons. But I suspect that most will find this result highly counterintuitive. (J) is clearly the morally least preferable outcome. It seems then, that RP only avoids the counterintuitive implication in cases involving miserable lives at the cost of being committed to equally counterintuitive implications for this last case involving the choice between (H), (I), and (J).

V.1.3 The No Difference View

I have argued that RP has problematic implications for two kinds of creations cases. I now argue that RP has trouble accommodating two kinds of cases that Derek Parfit has appealed to in support of what he calls 'the No Difference View'. According to the No

Difference View, the fact that an act harms someone (by making her worse off) makes no difference to its moral permissibility. In other words, the view holds that harm is a morally irrelevant consideration. Parfit contrasts the No Difference View with what he calls the ‘Two Tier View’. On the Two Tier View, the fact that an act would harm someone does make a difference to its moral permissibility. Parfit offers the following description of the Two Tier View: ‘though we always have reasons not to cause future lives to be less worth living, these reasons would be weaker if, because these lives would be lived by different people, these acts would not be worse for any of these people.’¹⁵⁶ To see how the Two Tier View and the No Difference View diverge with respect to their practical implications compare the case of Wilma with the similar but importantly different case of Betty:

Betty is pregnant with Bam Bam. She goes to the doctor for a routine checkup and the doctor tells her that, as things stand, the child she has conceived will be born blind. This condition is not inevitable, however. If she takes a pill every day for the next two months then Bam Bam will be perfectly sighted. Like Wilma, Betty cannot be bothered to take the pills, so Bam Bam (the child she is pregnant with now) is born blind rather than sighted. He is thereby caused to be worse off than he otherwise would have been.

Presumably most will react that Betty, like Wilma, acts wrongly by failing to take the pills. If we hold the No Difference View then the fact that Bam Bam will be caused to be worse off than he otherwise would have been does not matter for the moral assessment of Betty’s action. Wilma and Betty have all and only the same moral reasons to take the pills. If we hold the Two Tier View, however, we must say that there is something

¹⁵⁶ Parfit (2011: vol. 2, 219).

especially morally objectionable about Betty failing to take the pills that is missing in the case of Wilma. Betty has some extra reason to take the pills because her failure to do so harms someone by causing Bam Bam to be worse off than he otherwise would have been.

If it turned out that the No Difference View was true, this would seem to lend some strong support to UP. The truth of the No Difference View would entail that Wilma and Betty have the same moral reasons to take the pills. The discussion of harm-based solutions in Chapter II concluded that Wilma's action, if it is wrong, is *not* wrong because of its harmful effects. The most plausible explanation for the parity of reasons between Betty and Wilma, then, is likely to be a principle like UP. RP, for instance, cannot countenance the fact, if it is a fact, that Wilma and Betty have all and only the same moral reasons to take the pills. This is because RP generates a moral objection to an action only if that action would bring about an outcome that is worse for some actual people than some available alternative. Because Wilma's failure to take the pills is not worse for any actual people (Pebbles would not have existed had Wilma taken the pills), RP generates no moral objection to her action. But for the same reasons, RP *does* generate a moral objection to Betty's failure; Betty's failure will bring about an outcome that is worse for some actual person (Bam Bam) than taking the pills would have been. UP, on the other hand, is compatible with the possibility that Wilma and Betty have all and only the same moral reasons to take the pills. UP compares the welfare of all of the actual people in *each* outcome when assessing the moral reasons that an agent has to perform some action. UP takes no stance on how to *weigh* the welfare of people who will exist in each available outcome against the welfare of people who will only exist in some of the available outcomes, but it is consistent with the possibility that they are of equal moral significance

(an implication of the No Difference View). UP is thus consistent with the possibility, entailed by the No Difference View, that Wilma and Betty have the same moral reasons to take the pills whereas RP is not. Because RP is not consistent with this possibility, the truth of the No Difference View would spell yet more trouble for RP. Parfit offers two cases that he thinks show that we should reject the Two Tier View in favor of the No Difference View.

The first case involves choosing between two medical programs. One of the programs (P1) will affect a group of mothers who are already pregnant. The other program (P2) will affect women who will soon conceive a child, but who have not yet done so. If we choose P1 then the group of pregnant mothers will be tested for a condition that shortens the life of their unborn children. The women who test positive will be cured of this condition and their unborn children's lives will be extended. If we choose P2 then the group of women who will soon conceive a child will be tested for a condition that will shorten the life of whatever child they conceive. Those women who test positive for the condition will be cured of the condition, but the process will require them to delay conception by two months. Thus the women who are not yet pregnant will conceive different children than they would have conceived if they are cured of the life-shortening condition. If neither program is chosen (let's call the outcome in which neither program is chosen 'P3') then none of the mothers—neither the pregnant women nor the women who will soon conceive a child—will be cured of the conditions that will shorten the lives of their children. Consider the following representation of the effects of the two programs on two mothers and their children.

Program 1: (A) 70 (B) 50

Program 2: (A) 50 (C) 70

Program 3: (A) 50 (B) 50

Amy has the life-shortening condition, and she is pregnant with child A. If we choose Program 1 then Amy will be diagnosed and cured of the condition, and child A will live an extended life. Bonny, on the other hand, who is not yet pregnant, will remain undiagnosed and will conceive child B who will live a shortened life. If we choose Program 2 then Amy will remain undiagnosed and uncured, and the child that she is now pregnant with will live a shorter life than it otherwise would have lived. If we choose 2 then Bonny, who will soon become pregnant, will be diagnosed and cured, and she will conceive child C with an extended life. If we choose program 3 neither woman will be diagnosed or cured and both Amy's unborn child A and Bonny's not-yet-conceived child B will have the life-shortening condition. Parfit claims that intuitively there is no moral difference between 1 and 2, but the Two Tier View entails that 1 is morally preferable to 2 because it prevents A from being caused to be worse off than A otherwise would have been. Because this is implausible, he claims, we ought to reject the Two Tier View in favor of the No Difference View. If this is the correct conclusion to draw from this case, then this provides significant support for UP over RP, because only UP is compatible with the No Difference View.

Difference View and thus, once again, to UP.

Note, however, that Parfit only considers two methods of evaluating medical programs 1, 2, and 3 and options 4, 5, and 6 in case 6.¹⁵⁸ These are the No Difference View and the Two Tier View. The defender of RP might contend that this is a false dichotomy, since his view also has implications for which medical programs we should choose and for which options in case 6 we should choose. If RP generates the intuitively correct verdict about medical programs 1, 2, and 3 and options 4, 5, and 6 for case 6, then the two cases will fail to support the No Difference View over RP and thus will fail to support UP over RP. We cannot rule out, *a priori*, that RP will generate the wrong verdict about the cases. So what does RP entail about the cases?

On one description of RP, it seems to get the wrong verdict about the cases, because the principle seems to entail that medical program 2 is morally preferable to program 1. Child B is, after all, worse off in 2 than in 1, but there is no child who is worse off in 1 than in 2. As stated, RP is the view that there is a moral objection to choosing some action only if there is some actual person who is worse off in that outcome. As stated, then, RP entails that program 2 is morally worse than 1, because it is worse for child A, who exists whether program 1 or 2 is chosen. This would leave the defender of RP unable to account for the apparent fact that there is no moral reason to prefer program 1 over program 2. There is, however, a second way of describing RP according to which it is concerned, not with whether *some actual individual* is worse off in one outcome relative to another, but rather

¹⁵⁸ He has by this point ruled out what he calls the 'strong person-affecting principle' which states that one outcome can be worse than another if and only if it is worse for someone.

whether the *group* of actual individuals is worse off in one outcome relative to another. If, we interpret RP as telling us to focus our attention on *the interests of all and only those people who actually exist* given that we chose Program 2, then RP seems to get the right verdict about the case involving the medical programs. Consider a reformulation of RP that captures this interpretation of the principle:

Restricted Pro tanto* (RP*): There is a moral reason (not) to perform an action *a* if the set of actual people given *A* will be better (worse) off given *A* than the set of actual people given $\sim A$, but this moral reason generates a moral requirement (not) to choose *A* in the absence of sufficient countervailing reasons **only if the group of actual people given $\sim A$ (A) would be worse (better) off given $\sim A$ (A) than given *A* ($\sim A$)**.¹⁵⁹

RP* tells us to take the individuals who actually exist given each medical program and ask if these individuals *as a group* would have fared better had we chosen the other alternative. If the actual individuals as a group, given the choice of some program, would have fared better given the alternative program, then RP* generates a moral objection to the chosen program. If they would not have fared better in the other alternative, then RP* generates no moral objection to the chosen program. Suppose we choose program 1. A will live for 70 years, given program 1, and B will live for 50 years. The total duration of their

¹⁵⁹ This principle does not rule out the possibility that there could be a moral reason to promote the good in cases where failing to do so would leave no actual person worse off. It simply rules out the possibility that this moral reason could generate a moral *requirement*.

lives will be 120.¹⁶⁰ To find out if RP* generates any objection to program 1 we must ask if A and B as a group would have fared better in program 2. The answer is that they would not have. In program 2, A will live for 50 years, and B will not come to exist. Thus the total duration of the combined lives of A and B in program B is 50 years, so RP* generates no objection to choosing 1 over 2. Now suppose we choose program 2. If we choose 2 then A will live for 50 years and C will live for 70 years. The total duration of A and C's combined lives will be 120 years. In program 1, A will live for 70 years, and (C) will not come to exist. Thus the total duration of A and C's combined lives is only 70 years given program 1. 120 years is more than 70 years, so A and C do not fare worse in 2 than in 1. Therefore, RP* generates no objection to choosing 2 rather than 1. RP* seems to generate the correct verdict that there is no moral difference between programs 1 and 2. It is permissible to choose either program. The case thus seems to fail to support the No Difference View, and thus fails to support UP over RP*.


RP* also seems to fare well when it comes to Parfit's case 6. For each of outcomes A-C, the principle asks us to consider how the actual people fare given that outcome and then to compare it to how those actual people would have fared had the alternative outcome been chosen. For each of choices 4, 5, and 6 it is clear that the actual people would not have fared better, as a whole, had one of the alternatives been chosen. Suppose for instance that we choose 4. Adam lives for 70 years and Bernard lives for 40 years for a total of 110 years given 4. Had 5 been chosen, though Bernard would have lived for 90 years, Adam would not have existed. Thus their total years lived in 5 would only be 90 compared with

¹⁶⁰ Let us stipulate that the quality of each individual's life will be the same regardless of which program we choose.

110. Had 6 been chosen neither Adam nor Bernard would have come to exist, so the total number of years lived by them would amount to nothing. 110 is not worse than 90, nor is it worse than nothing, so RP* generates no objection to 4. For the same reason, RP* also generates no objection to 5 or 6 either. Case 6, like the case of the medical programs, seems not to give any reason to accept the No Difference View over RP, and it thus fails to support UP over RP.

So far, Parfit's cases seem not to favor the No Difference View over RP (formulated as RP*), so they seem not to support UP over RP. Looking more closely at the case of the medical programs, however, RP* turns out to have implications that are deeply problematic. Consider the options in the form of the following figure:

Figure 3



RP* asks us to compare how the actual people fare in an outcome with how those very people would have fared had an alternative outcome been selected. If the actual people would not have fared better in any of the alternative outcomes then RP* generates no moral objection to choosing that outcome. This procedure seems to arrive at the correct verdicts

about the permissibility of programs 1, 2, and 3. 1 and 2 are both permissible, but 3 is impermissible in the absence of sufficient countervailing reasons, because it is worse for group of children that includes A and B than program 1.

It seems, however, that there is another question we might ask about programs 1, 2, and 3, and the answer RP* gives to this question yields highly problematic results. The question is whether two outcomes are *morally equivalent* to each other. We can ask, for instance, whether 1 is morally equivalent to 2. To determine whether two outcomes, say 1 and 2, are morally equivalent we ask whether there is anything morally objectionable about one of them relative to the other and vice versa. Moral equivalence appears to be a transitive relation. If 1 is morally equivalent to 2, and 2 is morally equivalent to 3, then 1 is morally equivalent to 3. Let us see what RP* entails about the moral equivalence of 1, 2, and 3. Given that we choose 1, 1 is not worse for group of the actual people than 2 would have been. Given that we choose 2, 2 is not worse for the group of the actual people than 1 would have been. 1 and 2, then, seem to be *morally equivalent* by the lights of RP*. Given that we choose 2, 2 is not worse for the group of the actual people than 3 would have been. A's life would be 50 years in 2 and B's life would be 70 years in 2, for a total of 120 years. A's life would be 50 years in 3 and (B) would not exist in 3. 120 is greater than 50, so 2 is not morally objectionable relative to 3 by the lights of RP*. Given that we choose 3, 3 is not worse for the set of the actual people than 2 would have been. B's life is 50 years and C's life is 50 years, for a total of 100 years. B's life would be 50 years in 2 and C would not exist. 100 years is more than 50, so RP* says that 3 is not morally objectionable relative to 2. Because 2 is not objectionable relative to 3 and 3 is not objectionable relative to 2, 2 and 3 are morally equivalent. But 1 and 2 are morally equivalent. By transitivity, 1 and 3

are morally equivalent. But we have seen that RP* also entails that 3 is morally objectionable relative to 1, and two outcomes cannot be morally equivalent if one is morally objectionable relative to the other. It turns out, then, that RP* and the transitivity of the *moral equivalence relation* together entail both that 3 is and is not morally equivalent to 1. This is a contradiction, so we must reject one of the suppositions that got us there. We can reject the supposition that the moral equivalence relation obtains between two choices or outcomes if and only if neither of two alternatives is morally objectionable (preferable) relative to the other. This supposition seems entirely uncontroversial, however, so we should prefer to retain it if possible. Alternatively, we might reject the supposition that the moral equivalence relation is transitive. Some have proposed that certain kinds of relations between outcomes or choices do not obey transitivity, but examples that purport to support the intransitivity of these relations generally involve comparisons between values that are not commensurable (e.g. knowledge and beauty).¹⁶¹ In comparing programs 1, 2, and 3 there does not appear to be two incommensurable kinds of value being compared, so there is no reason to think that the moral equivalence relation that obtains between them is intransitive. I suggest that we reject RP*, thereby avoiding the implication that 3 is morally equivalent to 2, rather than reject a natural and plausible characterization of the moral equivalence relation. But of this move is tantamount to accepting UP, to whose implications for the medical program case I now turn.

UP does not entail the problematic conclusions that program 3 simultaneously is and is not morally equivalent to program 1. UP like RP* entails that 1 is better than 3. Also,

¹⁶¹ Chang (1997: 2).

UP, like RP*, entails that 1 is morally equivalent to 2. Unlike RP*, however, UP is consistent with the possibility that choosing program 3 is morally objectionable relative to program 2. This means that, even if we suppose the transitivity of the moral equivalence relation, UP avoids the implication that 3 simultaneously is and is not morally equivalent to 1. By the lights of UP, 3 is worse than 1 according to all measures. The ability of UP to generate the correct verdicts about the three medical programs and case 6, verdicts that are in concert with the No Difference View, provides yet one more reason to prefer UP over RP.

V.2 Countervailing Reasons

In the previous section I argued that RP cannot satisfactorily accommodate our judgments about two kinds of creation cases. It also entails a contradiction when applied to cases designed to motivate the No Difference View. These problems give us strong *prima facie* reasons to prefer UP to RP*. We now have the foundation of a solution to the Non-identity Problem. UP entails that there is something morally objectionable about Wilma's choice to conceive Pebbles rather than the sighted child she could have conceived. It also entails that there is something morally objectionable about the wealthy society's choice of depletion rather than conservation. In each case the agent(s) in question fail to act in accordance with their moral reasons to promote the good. But UP does not claim that the facts about the welfare of actual people are all that matters in determining the moral permissibility of an action. UP states that there is a moral reason (not) to perform an action *a* if the set of individuals given *A* would be better (worse) off than the set of individuals would be given $\sim A$, and this moral reason generates a moral requirement *in the*

absence of sufficient countervailing reasons. I have not yet said much about countervailing reasons other than a few introductory remarks in Chapter IV. An account of sufficient countervailing reasons is important in two respects. First, without an account of sufficient countervailing reasons, the possibility remains that both Wilma and the wealthy society *do* in fact have sufficient countervailing reasons (whatever these turn out to be) to act the ways that they do. If they do have sufficient countervailing reasons then UP does not generate a moral requirement for them to either choose to conceive the sighted child or to choose conservation. The second reason is that I have claimed that my solution to the Non-identity Problem is *moderate* in the sense that it entails a moral theory that is less demanding than Maximizing Consequentialism. Without an account of sufficient countervailing reasons, however, the possibility remains that the moral theory entailed by UP is in fact no less demanding than Maximizing Consequentialism. For these reasons it is important that I explain and defend the notion of ‘sufficient countervailing reasons’ I have in mind as the term appears in UP.

V.2.1 Countervailing Reasons and Maximizing Consequentialism

I have claimed that my view is distinct from a Maximizing Consequentialist view because it makes space for morally relevant considerations other than the promotion of the good. This aspect of my view is supposed to be captured by the ‘sufficient countervailing reasons’ clause at the end of UP. I begin this section with some preliminaries that are meant to show precisely how the mere inclusion of sufficient countervailing reasons blocks the collapse of UP into Maximizing Consequentialism. First, consider a rough definition of ‘countervailing reason’ (CR).

CR: A reason r is a countervailing reason just in case it counts against some action a that is supported by a reason r' ; r must arise from a different kind of moral consideration than r' .

r and r' refer to two *different kinds* of reason. This is meant to preclude the possibility that, say, a reason to promote the good would count as a CR against an agent being morally required to perform an action that itself would promote the good. For instance, if I have \$20 that I can give to either Jon or Jane, the fact that I would benefit Jane by \$20 generates a reason that counts against my giving the money to Jon and vice versa. But this reason is not a CR against giving the money to Jon, on the definition above, because the reason that counts against giving the money to Jon (the reason generated by the fact that I would benefit Jane by \$20 if I gave the money to her) is the *same kind* of reason that counts in favor of giving the money to Jon. They have the same *source*. Both the reason to give the money to Jon and the reason to give the money to Jane are generated by the fact that the actions would promote the good. In a later section I will elaborate a bit more about what kinds of non-good-promoting CRs I have in mind as they pertain to UP.

CRs, on this conception, might come in two varieties, each of which might count against an agent being morally required to perform some action. One variety of CR is a reason that itself generates a moral requirement under certain circumstances. I refer to this kind of reason as a 'countervailing moral reason' (CMR). For instance, the fact that an action a would constitute

breaking a promise to someone might be the kind of facts that generate a reason r against me being morally required to perform a , and these facts themselves might, under the right circumstances generate a moral requirement *not* to perform a . The thought is that, other things being equal, I am morally required not to break promises, and the fact that performing some action a would involve harming or breaking a promise generates a reason that counts against me being morally required to perform a even if I have some moral reason that favors performing a . A second kind of countervailing reason counts against an agent being morally required to perform some action but cannot *itself* generate a moral requirement. I refer to this kind of reason as a ‘countervailing non-moral reason’ (CNMR). Recall the case from the previous chapter in which, at little or no cost to myself, I can take my friend Tom to the hospital so that he will undergo a procedure to render his legs fully functional. Now imagine that in addition to taking him to the hospital, I must donate one of my legs to Tom in order to enable him to enjoy the full use of his legs. I suspect that, while most think that I am morally required to take Tom to the hospital, most will be less confident in this case that I am morally required to donate my *own* leg to Tom’s cause. Most will likely have the further reaction that I am not morally required *not to* donate my leg to Tom’s cause if I wish to. This pair of reactions is best explained by the fact that, though I have strong reasons of self-interest not to donate my leg, and these reasons count against my being morally required to maximize the good by giving Tom the full use of his legs, these reasons cannot generate a moral requirement for me to favor my own interests. Perhaps it is a brute feature of morality that one cannot be morally required to promote her interests. If this is correct then I have a powerful CNMR but not a CMR that counts against me being morally required to donate my leg to Tom’s cause.¹⁶² It would be permissible for me to refrain from

¹⁶² Joshua Gert (2003) distinguishes between two *dimensions of normative strength*. He thinks that moral

donating my leg, but it would not be wrong for me to donate my leg either.

The existence of countervailing reasons of either the CMR or the CNMR variety, given the analysis in the previous paragraph, is incompatible with Maximizing Consequentialism. MC cannot make room for countervailing reasons as I have defined them, because MC *only* acknowledges the moral relevance of reasons to promote the good. Because there is only one kind of morally relevant reason (the moral reason to promote the good), according to MC, there cannot be morally relevant reasons of a different *kind* (aside from reasons having to do with promotion of the good) that count against an agent being morally required to perform the action that would maximize the good. But *countervailing* reasons, according to the definition above, must be different *kinds* of reasons than the reason they countervail. Because Maximizing Consequentialists only admit the existence of one kind of moral reason, they cannot also admit the existence of countervailing reasons.¹⁶³ But what is it for a countervailing reason to be *sufficient* in the sense that it can prevent the moral reason to promote the good from generating a moral requirement? I now turn to describing the notion of a *sufficient* countervailing reason (SCR).

considerations can have both a *justifying* dimension and a *requiring* dimension, and that some considerations can produce strong reasons along one dimension but weak reasons along the other. I take it that his distinction roughly corresponds to what I have been calling *CNMRs* and *CMRs*. *CNMRs* can morally justify a choice of action, but only *CMRs* can morally require a choice of action.

¹⁶³ Again I should note that Pluralistic Consequentialists might object that their view leaves room for countervailing reasons because they recognize different elements of the good. So, that option A promotes happiness is a reason to do A, but that option (B) (inconsistent with A) promotes friendship is a reason to do (B) instead. These are not genuine countervailing reasons, however, for they both derive from a more general concern with promoting a kind of value. Though these reasons may ‘countervail’ one another, they are reasons of the same kind insofar as they derive from a more the more general moral imperative to promote the good.

Sufficient Countervailing Reason: A reason r that counts *decisively* against some agent s being morally required to perform some action a , where s has some moral reason r' to perform a .

An SCR, on this definition, is just a CR (of either the CMR or CNMR variety) that *successfully* or *decisively* counts against some agent being morally required to perform some action. If an agent has an SCR against being morally required to perform an action, then that agent is *not* morally required to perform the action. Not all CRs are SCRs. Sometimes a countervailing reason counts decisively against an agent being morally required to perform some action, but other times it does not. There are cases in which someone is morally required to perform an action despite clearly having morally relevant reasons not to perform the action. Recall the example from section IV.2 involving my friend Tom who has never enjoyed the use of his legs. If I take Tom to the hospital in the next 30 minutes I can ensure that Tom will enjoy the use of his legs for the remainder of his life. It seems that I do something morally wrong if I do not take Tom to the hospital and that this fact is best explained by the existence of a moral reason to promote the good. Now imagine that the only way for me to get Tom to the hospital in the next 30 minutes is to break my promise to Bob that I would help him move into a new apartment. It seems that most will react that I am morally required to take Tom to the hospital in time even if doing so requires breaking a promise to Bob. If I am still morally required to take Tom to the hospital, then one can have a CR not to promote the good that is not an SCR. The fact that taking Tom to the hospital breaks a promise to Bob is a CR not to take Tom to the hospital, but it is not an SCR; it counts against me being morally required to take Tom to the hospital, but it does not

decisively count against being so required. Whether a CR counts as an SCR depends on its strength and on the strength of the reason against which it is being weighed.

Now consider a version of the case in which I can only get Tom to the hospital in time by running over a pedestrian, breaking one of his legs. Most will likely react that I am morally required not to take Tom to the hospital in time for the surgery if doing so would require breaking the pedestrian's leg. This reaction is likely best explained by the fact that I have an SCR (perhaps generated by the fact that I would significantly *harm* or *violate the right* of the pedestrian by breaking his leg) not to take Tom to the hospital in time for surgery. I take this contrastive pair of cases to show that CRs and SCRs are not extensionally equivalent. Sometimes a CR is an SCR, but sometimes it is not.

V.2.2 Kinds of Countervailing Reasons

With analyses of CR and SCR in hand, it should be clear that CRs (both CMRs and CNMRs) and SCRs are incompatible with Maximizing Consequentialism. UP thus avoids collapse into MC insofar as it makes space for CRs and SCRs. I now turn to providing a provisional, admittedly incomplete, account of the kinds of reasons I have in mind as examples of CRs. Since we are interested in constraining UP in such a way that agents are not always morally required to maximally promote the good, this list includes only reasons that might count *against* an agent being morally required to promote the good. Offering a comprehensive list of CRs is a serious challenge, but the challenge is not peculiar to someone who wishes to defend UP. It is a challenge for ordinary morality more generally. If the difficulty of providing an

exhaustive list of all countervailing reasons fails to make us hesitant to use them in ordinary moral theorizing (which I think it does not), then the same difficulty should not make us hesitant to use them to make moral distinctions in identity-affecting cases. UP can accommodate (at least) the two most commonly cited kinds of countervailing reasons from ordinary morality: considerations of self-interest and deontological constraints against harming and violating rights. My aim in the remainder of this section is to show that CRs are in fact ubiquitous in ordinary morality. I then sketch a provisional list of the kinds of CRs that might count against an agent being morally required to perform an action that would maximally promote the good. In light of this provisional list, I then reevaluate Wilma's CRs for conceiving a blind child and the wealthy society's CRs for choosing depletion.

V.2.3 Countervailing Reasons in Ordinary Morality

CRs and SCRs are indispensable features of ordinary morality. For example, we must appeal to the notion of a sufficient countervailing reason when determining whether a moral agent has permissibly or wrongfully failed to prevent harm. Consider what I take to be a case of *wrongful* failure to prevent harm.

Cruel Calvin: Calvin is walking to class when he sees a young boy struggling to stay above water in a pond. There is no one else around to hear his flailing. Calvin knows that if he does not help the young boy he will drown. Calvin also knows that if he goes into the pond he will ruin his new fifty dollar jeans.

Unwilling to sacrifice his newly purchased pants, Calvin continues on his way, and the struggling child drowns.

Most will react that Calvin has violated a moral requirement by failing to help the young boy. Calvin has CNMRs of self-interest not to save the drowning child, but these reasons are not SCRs. They do not count decisively against Calvin having a moral requirement to save the child. His CNMRs are not SCRs so he has done something wrong. Now consider a case of *permissible* failure to prevent harm.

Conflicted Calvin: Calvin is walking to class when he sees a young boy struggling to stay above water in a pond. There is no one else around to hear his flailing. Calvin knows that if he does not save the young boy he will drown. Calvin, however, has a unique genetic disorder that causes severely restricted circulation to his lower extremities when he comes in contact with cold water. He knows that he can jump in the pond and successfully save the child, but doing so will cause him to permanently lose the use of both of his legs. He will be wheelchair-bound for the rest of his life. Though it pains him to do so, Calvin does not jump into the pond because he is unwilling to sacrifice the use of his legs. Instead he calls 911, and the struggling child drowns.

Most will react that Calvin does not violate a moral requirement by failing to save the drowning child in this case, but the case is quite similar to the case of *Cruel Calvin*. In both cases Calvin's reason to save the child is generated by the fact that he could thereby prevent harm to the child. In both cases Calvin has CNMRs of self-interest that count against him being morally required to save the drowning child.¹⁶⁴ The only difference between *Conflicted Calvin* and *Cruel Calvin* is the strength of Calvin's countervailing reasons. The agent's countervailing reasons in *Conflicted Calvin* [generated by the fact that Calvin's life will go much better if he enjoys the use of his legs] are SCRs, and they count decisively against Calvin being morally required to save the child, whereas the countervailing reasons in *Cruel Calvin* are not SCRs; they are 'insufficient' to count decisively against Calvin being morally required to save the drowning child. Thus, while Calvin acts permissibly in the second case, he acts impermissibly in the first case. I will assume that these are the correct conclusions to draw about these cases from the standpoint of ordinary morality.

Most will not demand a precise account of the threshold at which CRs become SCRs before acknowledging both that Calvin in *Cruel Calvin* has done something wrong and that Calvin in *Conflicted Calvin* has not. It may be practically impossible to offer a clear distinction between sufficient and insufficient countervailing reasons that encompasses all of our intuitions about possible cases. Fortunately, the difficulty of making a clear distinction between CRs and SCRs does not cast serious doubt about the fact that there is a moral distinction between what *Cruel Calvin* does and what *Conflicted Calvin* does. The distinction does not become arbitrary or

¹⁶⁴ If you think that *Conflicted Calvin* has a special moral obligation to his child then he has moral countervailing reasons in addition to his self-interested ones. This does not affect my argument.

ad hoc in the absence of a precise analysis of countervailing reasons. To the contrary, we comfortably invoke the distinction to differentiate between wrongful harming (failing to prevent harm) and permissibly harming (failing to prevent harm) in ordinary morality.

Since most are not troubled by the imprecise distinction between sufficient and insufficient countervailing reasons in ordinary moral cases that do not involve identity-affecting actions, we should not require a clear distinction is before we are willing to employ the notion of countervailing reasons when assessing the rightness and wrongness of identity-affecting actions.

V.2.4 (At Least) Two Kinds of Countervailing Reasons

The analysis of sufficient countervailing reasons required for my solution to the Non-identity Problem is not controversial or out of the ordinary. I take at least the following considerations to be morally relevant: (a) personal sacrifice and self-interest and (b) moral constraints [constraints against harming, promise-breaking, etc].¹⁶⁵ The moral relevance of considerations (a) and (b) is an indispensable feature of ordinary morality, and I assume that both (a) and (b) can generate CRs that count (sometimes decisively) against an agent being morally required to promote the good.

The moral relevance of considerations (a) and (b) is easily motivated though perhaps rather difficult to prove. There are several threads within the philosophical literature that involve

¹⁶⁵ As Kagan (1988) points out, ordinary moral theorists take the moral relevance of considerations (a) and (b) to make space for what he calls ‘constraints’ against promoting the good and ‘options’ not to promote the good.

attempts either to defend or to refute their moral relevance.¹⁶⁶ I do not think that a complete defense of the moral relevance of considerations (a) and (b) is necessary for my solution to the Non-identity Problem to be successful for the following reason. I have already defended the existence of an unrestricted *pro tanto* moral reason to promote the good (in the form of UP). If considerations (a) and (b) turn out to be morally irrelevant, thereby failing to generate reasons of either the CMR or CMNR variety, then the truth of UP will entail a very demanding moral theory. It will be a moral theory that amounts to some version of MC (Maximizing Consequentialism). In that case UP will still solve the Non-identity Problem, because it will entail that Wilma and the wealthy society act wrongly by failing to *maximally* promote the good. They can't have SCRs that count against maximally promoting the good, because, if Maximizing Consequentialism is true, there are no CRs. *A fortiori*, there are no SCRs. But, my solution to the Non-identity Problem is motivated, in part, by the supposition that most will reject MC precisely because most take considerations (a) and (b) to be morally relevant. Therefore, I proceed supposing that considerations (a) and (b) are morally significant. The inclusion of (a) and (b) is what makes my solution an improved, moderate version of the MC solution. UP can now be formulated more exactly as follows:

Unrestricted Pro tanto* (UP*): There is a moral reason (not) to perform an action *a* if the

¹⁶⁶ Samuel Scheffler (1982) and Doug Portmore (2008) defend versions of consequentialism intended to accommodate the view that personal sacrifice or self-interest is a morally relevant consideration apart from its usefulness in promoting the good. This indicates that even consequentialists see the moral significance of self-interest for moral theorizing. Moral appeals to and defenses of rights are ubiquitous in the philosophical literature including Nozick (1974), Regan (1983), Sumner (1987), Thomson (1971) and (1990), and Kamm (1992) and (2007).

set of individuals given A would be better (worse) off than the set of individuals would be given $\sim A$, and this moral reason generates a moral requirement in the absence of sufficient countervailing reasons, where ‘sufficient countervailing reasons’ refers to (at least) (a) reasons generated by the fact that performing A would require a significant personal sacrifice and (b) reasons generated by the fact that performing A would violate a moral constraint (against harming or violating rights).

In the next section I show that UP* does not morally excuse Wilma or the wealthy society for failing to act in accordance with the moral reason to promote the good. They do not have to endure any serious personal sacrifices nor do they have to violate any moral constraints in order to choose the action that better promotes the good. Their countervailing reasons are thus insufficient, even provided an account of sufficient countervailing reasons that is consistent with the central commitments of ordinary morality, to decisively count against them being morally required to conceive the sighted child rather than Pebbles or to choose conservation rather than depletion.

V.3 A Solution to the Non-identity Problem

In the previous sections of this chapter I explained and defended the two central components of my solution to the Non-identity Problem. In V.1 I defended UP against its competitor RP. This principle provides a foundation for a solution to the Non-identity Problem by entailing that there is something morally objectionable about Wilma conceiving Pebbles

rather than a sighted child or the wealthy society choosing depletion rather than conservation. In the above section I conjoined UP with a provisional but initially plausible account of countervailing reasons. The result was UP*. In this section I show that UP* can solve the Non-identity Problem while remaining consistent with the central components of ordinary morality. Let us call the complete conjunction of these claims and the conclusions that follow from them the 'Welfare Argument'.

The Welfare Argument:

P1) UP*: There is a moral reason (not) to perform an action a if the set of individuals given A would be better (worse) off than the set of individuals would be given $\sim A$, and this moral reason generates a moral requirement in the absence of sufficient countervailing reasons, where 'sufficient countervailing reasons' refers to (at least) (a) reasons generated by the fact that performing A would require a significant personal sacrifice and (b) reasons generated by the fact that performing A would violate a moral constraint (against harming or violating rights).

C1) *From 1 and the case of Wilma and the wealthy society:* Because Pebbles has lower lifetime welfare than Wilma's sighted child would have had, the set of people who would exist if Wilma has the blind child will be worse off than the set of people who will exist if she has the sighted child. Because the members of the depleted society will have lower lifetime welfare than the members of the conservation society would have had, the set of people who would exist if

the wealthy society depletes their resources will be worse off than the set of people who would exist if the wealthy society chooses conservation.

P2) Wilma lacks countervailing reasons that count decisively against Wilma's being morally required to conceive the sighted child rather than Pebbles. The wealthy society lacks countervailing reasons that count decisively against its members being morally required to choose conservation rather than depletion.¹⁶⁷

C2) *From 1, C1, and 2:* Wilma violates a moral requirement (acts wrongly) when she conceives Pebbles rather than a sighted child. The wealthy society acts wrongly by choosing depletion rather than conservation.

I have not yet defended (P2) of the Welfare Argument, but it is the least controversial of the premises. In this section I defend (P2). Perhaps the easiest way to justify (P2) is to compare Wilma's case with another one that is similar to it but importantly different. Consider

The Risky Flight: The Jetsons are debating whether to fly to a New York fertility clinic to conceive a blind child or to fly to a London fertility clinic to conceive a sighted child. Jane has a rare condition that will cause her to go blind in one eye if she spends too much time in an airplane. She can endure a

¹⁶⁷ If you think Wilma does have sufficient countervailing reasons, then imagine that the inconvenience is even less significant.

flight to New York, but a flight to London would result in her becoming blind in one eye. This is the Jetsons' last chance to conceive a child. If they do not pick one clinic or the other they will not be able to have a child that is biologically entirely theirs, which is something that they highly value. They decide, because of Jane's health, to fly to New York to conceive an incurably blind child.

In this case, the Jetsons, like Wilma, are choosing the worse of two alternatives from the standpoint of promoting the good. The individuals who exist if the Jetsons fly to New York will be worse off than the individuals who exist if they fly to London. In spite of this fact, most will react that it is *permissible* for the Jetsons to fly to New York. Though we might think that it is regrettable that they have to choose the worse alternative from the standpoint of promoting the good, few will think that they act *wrongly* by doing so. What is the difference between the case of *Wilma* and *Risky Flight* that is responsible for this divergence with respect to our judgments? The most obvious difference (and a difference that UP* can comfortably account for) between *Risky Flight* and *Wilma* is the strength of the countervailing reasons each individual has for choosing the worse of her two alternatives. In the original case, Wilma must endure a minor inconvenience. In *Risky Flight*, Jane must endure a sight-risking voyage in order to conceive the sighted child. UP* at least leaves open the possibility that Jane acts permissibly because it allows considerations such as significant personal sacrifice to play a role in determining the moral status of actions. Wilma's act is wrong in the original case because she fails to choose the better of the available outcomes, *and* she does so without sufficient countervailing reasons. The Jetsons, on

the other hand, have sufficient countervailing reasons to choose the suboptimal alternative. UP* is thereby consistent with the intuitively correct result that the Jetsons act permissibly when they fly to New York.

By contrasting our reactions to Wilma's case with the case of *Risky Flight* we see that our reaction to the case of Wilma depends, in part, on the fact that she seems not to have sufficient countervailing reasons to choose the worse of her alternatives. *Risky Flight* also highlights an appealing feature of UP*: it is consistent with a significantly less demanding moral theory than Maximizing Consequentialism. It does not require that agents always produce the most valuable outcome of those available to them. It simply requires that an agent have *some morally sufficient reason* for failing to produce the best outcome. If one accepts UP* then one can reject (P5) and the conclusion of the Non-identity Problem—which claims that an act is wrong only if it wrongs someone—without fettering morality with Maximizing Consequentialist baggage.

I have now defended each of the premises of my argument for a solution to the Non-identity Problem in some detail. In V.4 I turn to considering several objections to the Welfare Argument. These objections come in the form of counterexamples. The counterexamples purport to show that the Welfare Argument (specifically UP*) turns out to have implications that are more counterintuitive than the conclusion of NIPA, and that it thereby fails the Intuitive Cost Criterion for a successful solution to the Non-identity Problem. If my solution is to succeed, I must show either that my solution does not have these implications or that we should be willing to accept the implications.

V.4 Objections to the Welfare Argument from Problem Cases

I have argued that UP*, together with some plausible suppositions about the cases of Wilma and the wealthy society, yields a solution to the Non-identity Problem that is consistent with the central components of ordinary morality. If any of the claims that I have appealed to in developing my solution to turn out to have implications that are more counterintuitive than the conclusion they are designed to avoid, however, then this would be a good reason to reject my solution. There are two kinds of cases in particular that appear to show that UP* has implications that are more implausible than the conclusion of NIPA. If these cases are successful, then my solution to the Non-identity Problem fails the Intuitive Cost Criterion for a successful solution, and must be rejected.

V.4.1 Rescue Cases

A first kind of case supposedly poses a problem for Pro tanto in both its Unrestricted and Restricted versions (RP and UP*). It involves choosing between saving one person who is perfectly able-bodied and someone who is handicapped in some way.

Peter: Peter walks past a shallow pond where two boys, Bryan and Chris, are struggling to stay afloat. Bryan is blind, but Chris is perfectly sighted. Chris is slightly further away from Peter than Bryan is, so it would be slightly less convenient for Peter to save Chris than Bryan. Peter can only save one of the drowning boys, because, by the time he is able to drag one of them to shore, the other one will have drowned.

Peter seems closely analogous to the case of Wilma with the exception of one detail. The two children in the case of Peter already exist, whereas the two children in the case of Wilma do not yet exist. As in the case of Wilma, if Peter chooses the sighted child over the blind child he will more effectively promote the good. Also, as in the case of Wilma, Peter can do so at little or no cost to himself and without violating anyone's rights. Neither child, we might suppose, has an absolute right to Peter's assistance. If each child did have such a right, then Peter would do something wrong by violating one of the child's rights no matter whom he saved. But this would violate the rule that ought implies can, so it seems that the children must not have such a right. RP and UP* say that an agent is morally required to promote the good except in cases where doing so would (a) require a significant personal sacrifice or (b) violate someone's rights, so RP and UP* would seem to entail that Peter does something morally objectionable by saving Bryan, the blind child, rather than Chris, the sighted child. But this is highly counterintuitive; it seems that Peter acts morally permissibly whether he saves Bryan or Chris. But if that is correct, then both RP and UP* ought to be rejected on the basis of their counterintuitive implications for rescue cases.

V.4.1.A Reply from Inconsistency

Both the RP and UP* principle supposedly entail that Peter acts impermissibly by saving Bryan rather than Chris. This implication might seem sufficiently counterintuitive to reject both versions of the principle, but there are actually good reasons to accept the implication. The judgment that Peter acts permissibly by saving Bryan rather than Chris, is undermined by another judgment that most will have about the case. Most will judge that Peter violates a moral

requirement if he decides not to save either child. If that is the case then we are saddled with a peculiar set of judgments about Peter:

No Right (NR): Neither child has a right to Peter's assistance that is violated by Peter's failure to save him.

Either Or (EO): It is morally permissible for Peter to save either the blind child or the sighted child.

Neither (N): It is morally impermissible for Peter to save neither child.

If EO and N are both true, and I suspect that our confidence in N is just as high as our confidence in EO, then the inclusion of NR poses a problem. How can it be wrong for Peter to fail to save either child if neither child has a right to his assistance? If indeed N is true, given NR, this should undermine our confidence in EO, for, as far as I can tell, the best explanation for N, given NR, is that failing to save at least one child constitutes choosing the action that fails to promote the good in the absence of sufficient countervailing reasons. In other words, the best explanation of N, given NR, is either RP or UP*. But RP and UP* are inconsistent with EO, because RP and UP* supposedly entail that Peter acts wrongly by saving the blind child rather than the sighted child. I cannot think of any reason to have more confidence in EO than N or NR. Note that only if we

reject N can the objection to UP* be sustained. If NR is false, then someone's right (presumably the child that is left to drown) will be violated by Peter's action regardless of which child he saves, so UP* does not clearly entail that Peter violates a moral requirement to promote the good by saving Bryan rather than Chris. Rejecting EO will obviously not do to sustain the objection, since EO is the reaction that is supposed to generate the problem of UP* and RP in the first place. Since we have no reason to think that N is less plausible than EO or NR, in the absence of further arguments, there is only a 1/3 chance that the objection from rescue cases succeeds against UP*. This likelihood is actually smaller than 1/3, since I have provided independent arguments for UP* in previous sections, and UP* is inconsistent with EO.

V.4.1.B Reply from Extreme Differences

In my first reply to the objection to UP* from rescue cases I pointed out a tension between our reactions to the case of Peter. Most of the ways of resolving this tension, I argued, entail the failure of the objection from rescue cases. It is more likely, given our other commitments about the case of Peter, that UP* gives the correct verdict about the case than that it does not. For my second reply I offer a revised version of the case in which the discrepancy between the victims' quality of life is exaggerated. I argue that our reactions to this version of the case are inconsistent with our reactions to the case of Peter, and this gives us some reason to doubt that the objection from rescue cases succeeds. Consider

Lois: Lois has two children, Stewie and Meg, both of whom will soon die if they do not receive immediate medical attention. Meg's life is barely worth living. That is, the intrinsic goods she enjoys barely outweigh the intrinsic bads she suffers. Her life will continue to be barely worth living if she receives the life-saving attention. Stewie's life, on the other hand, is well worth living. The intrinsic goods he enjoys vastly outweigh the intrinsic bads, and his life will continue to be well worth living if he receives the life-saving medical attention. Lois can only provide the life-saving medical attention for one of the children. The other one will die. It is slightly more convenient for Lois to save Meg.

Most will think that Lois is morally required to save Stewie rather than Meg. Lois does something wrong if she chooses to save Meg, the child whose life is barely worth living. Since neither child has a right to Lois' assistance, this fact is likely best explained by the truth of either RP or UP*. But then what do we say about the case of Peter? That case is just like the case of Lois except that the welfare disparity between the children is smaller. But this is a difference in degree rather than a difference in kind. If the fact that one child is better off than the other generates a moral requirement to save the better off child in the case of Lois, then it seems that the same fact would generate a moral requirement to save the better off child in the case of Peter. This should make us much less confident that UP* gives the wrong verdict about the case of Peter, since it clearly gets the right verdict in the case of Lois.

It might appear that this reply depends on a questionable slippery slope form of reasoning. The success of my reply depends on the truth of the claim that the difference between the case of Lois and the case of Peter is a difference in degrees and the claim that a difference in degrees

cannot make a difference in what one is morally required to do. Strictly speaking, the latter claim is false. Differences in degrees do make a difference in what one is morally required to do, and UP* actually acknowledges this fact. For instance, UP* includes significant personal sacrifice as a morally relevant consideration that sometimes, but not always, generates SCRs against promoting the good depending on the *degree* of sacrifice required on the part of the agent. The difference between a significant and a negligible personal sacrifice is thus a difference in degrees. Thus, one might object that my reply from extreme differences supposes that differences in degrees cannot make a moral difference to what one is required to do, and yet I am already committed to the view that differences in degrees of personal sacrifice *can* make a moral difference to what one is morally required to do. This may seem problematically inconsistent. There is no inconsistency to be found here, however.

Earlier I made a distinction between countervailing moral reasons (CMRs) and countervailing non-moral reasons (CMNRs). CMRs are capable of generating moral requirements, whereas CMNRs can *only* prevent moral reasons from generating moral requirements. CMNRs and CMRs are just subsets of the more general categories of non-moral reasons and moral reasons. Personal sacrifice generates non-moral reasons, and these reasons do not generate moral requirements. Promotion of the good generates *moral reasons*, and these reasons generate moral requirements. According to the moral picture that I have developed in this chapter and the previous chapter, determining whether an agent is required or permitted to ϕ involves pooling together all of an agent's moral and non-moral reasons and then *weighing* them against one another. Some of these reasons may count in favor of ϕ ing. Others may count against ϕ ing. When the balance of an agent's reasons favors (counts against) ϕ ing, and at least some of the reasons that 'tip the balance', as it were, are moral reasons, then the agent is morally required (not) to ϕ . When the balance of an

agent's reasons favors (counts against) ϕ ing, and the reasons that tip the balance are non-moral reasons only, then the agent is permitted either to ϕ or not to ϕ . It is in this sense that differences in degrees *make a difference* to what agents are morally required to do. The degree of sacrifice required by an agent if she chooses (not) to ϕ determines the strength of her non-moral reasons (not) to ϕ . If these reasons are strong enough then they may tip the balance of her reasons in favor of (against) ϕ . In that case the agent is permitted either to ϕ or not to ϕ .

The upshot of this analysis for the relevance of differences in degrees to what one is morally required to do is the following: in both the case of Lois and the case of Peter it is stipulated that the balance of Lois' and Peter's reasons counts in favor of saving the better off child, *and* the reasons that tip the balance are stipulated to be *moral* reasons. So, according to the plausible analysis just offered, both Lois and Peter are morally required to save the better off child. Once this is settled, it is irrelevant that Lois has *more* moral reasons to save the better off child than Peter has. As long as it has been determined that moral reasons are the reasons responsible for tipping the balance of reasons in favor of saving the better off child, then differences in degrees of moral reasons make no *further* difference in determining what each agent is morally required to do. Differences in degrees of non-moral reasons make a difference to what an agent is morally required to do in exactly the same way. Once it has been determined that the balance of an agent's reasons favors (counts against) ϕ ing and that non-moral reasons only are responsible for tipping the balance, then it is settled that the agent is permitted either to ϕ or not to ϕ . Differences in degrees of one's non-moral reasons play no further role in determining what one is permitted to do. The only difference between moral and non-moral reasons on this account is that the former but not the latter generate moral requirements when they tip the balance of one's reasons in favor of (against) ϕ ing. This means that if Lois is morally required to save the better off child then Peter is

required to save the better off child. Though differences in degrees can make a moral difference in what one is required to do in a restricted sense, given the stipulated facts of the Peter and Lois cases, the difference in degree between Peter's reasons and Lois' reasons cannot make a difference with respect to what each is required to do.

V.4.1.C Reply from Distinguishing Between the 'Quality of Life' and the 'Moral Value of Life'

My third reply concedes that Peter is not morally required to save Chris rather than Bryan despite the fact that saving Chris is the better option from the standpoint of promoting the good. But how can this be? The cost to Peter of saving Chris rather than Bryan is minimal, and neither person has a right to Peter's assistance. UP* thus seems to be left without any sources of countervailing reasons capable of preventing Peter's moral reason to promote the good from generating a moral requirement.

I left open the possibility that my list of countervailing reasons was incomplete, and rescue cases seem to indicate that there may be need to make an addition to the list. Note first that, unlike identity-affecting cases in which an agent chooses *who comes to exist*, in rescue cases like Peter's there are two existing individuals each of whom values his life equally. This indicates that we may want to distinguish between the *quality* of a person's life and the *moral claim* the person has to continue living. This distinction might prove useful in allowing UP* to accommodate rescue cases like *Peter*. Though Bryan has a lower quality of life than Chris, he cares as much about continuing to live as Chris does. It is at least conceptually coherent that someone's moral claim to continue living does not depend on how good her life is, as long as it is worth living. Perhaps, then, when

we are choosing between saving one life rather than another, each person's equal claim to continue living *morally overshadows* the moral significance of the quality of those individuals' lives, rendering the latter consideration negligible with respect to its moral significance.¹⁶⁸ The decisive moral weight of an equal claim to continued life could explain why Peter is morally required to save one of the two children but why it is permissible for him to save either child. Because he cannot save both of them, he can satisfy only one of their claims to continue living. When it comes to saving lives, the quality of the victims' lives is of negligible moral significance next to their equal claims to continued life, and so the fact that Chris' life is better than Bryan's cannot count in favor of saving Chris rather than Bryan.

V.4.2 Creation Cases

The second kind of problem case poses a specific problem for UP* but not RP, and might thus seem to give support for the latter formulation of Pro tanto. This second kind of case involves an agent choosing between creating an additional person and not creating an additional person. Sometimes people are almost indifferent about having children. There would be costs to them of having a child, but those costs would be almost entirely compensated by the benefits of having the child. These people could 'take it or leave it' when it comes to having a child. My view might seem to imply that these people have an obligation to have children. This implication would seem to be a problem. Most do not think that a couple would have a moral obligation to conceive a child in these circumstances, but UP* might seem to entail that they do. In cases of

¹⁶⁸ The conviction that all lives are of equal moral value, regardless of their quality, seems to be steadfast in the disabilities literature. See Nord, Daniels, and Kamlet (2009): S10-S15, at p. S10-S11 for a discussion of this conviction and the problem it poses for cost-effectiveness analysis in healthcare.

indifference a couple is considering two possible outcomes, one of which involves having a child, the other of which does not, and the outcome in which they do not have a child is *very slightly* preferable from the standpoint of the couple's interests. According to UP*, however, the outcome in which the additional happy child exists is morally preferable to the outcome involving no additional child, because the additional happy child's welfare makes a significant contribution to good. To put it in terms of UP*, the additional happy child's welfare makes it the case that the individuals who exist are significantly better off than the individuals who otherwise would have existed would have been. Furthermore, the parents only have to make a minor sacrifice if they conceive the child (because they are almost indifferent). Consider one simple way of representing the two alternatives, one in which the parents conceive the additional child, one in which they do not:

A: Parents (100) Additional Child (100)

B: Parents (101)

It might seem that UP* implies that these indifferent parents have an obligation to have the child. This implication is sufficiently counterintuitive, my opponent might claim, to force us to reject UP* thereby blocking the inference to (C1). I will offer two responses to this objection to UP*. Neither reply denies that indifferent parents have some moral reason to conceive an additional happy child. Rather the replies deny that the existence of such a reason implies that indifferent parents are morally required to conceive an additional happy child. The first reply

proceeds by way of providing a clearer description of the referent of the term ‘significant personal sacrifice’ as it appears in UP*. The term is ambiguous as it stands. Once it is disambiguated it is clear that even indifferent parents must, in virtually every case, endure a significant personal sacrifice in order to conceive a child. My second reply appeals to the special moral significance of the right to reproductive autonomy. Within a certain realm, agents have the right to choose whether or not to reproduce, and they cannot act wrongly by choosing not to reproduce. I argue that it is plausible to think that the scope of this right is wide enough to ensure that even indifferent parents are not morally required to conceive children.

V.4.2.A Getting a Grip on the Notion of Significant Sacrifice

UP* includes a provision about a moral consideration that I have been calling ‘significant personal sacrifice’. UP* says that if bringing about the best outcome requires a significant personal sacrifice on the part of the agent, then the agent is not morally required to do it. Indifferent parents appear to be outside of the bounds of this clause, because these people will be only negligibly worse off if they have a child. Because the indifferent parents appear to be able to conceive a child without incurring a significant personal sacrifice and without violating anyone’s rights, UP* seems to entail that they have an obligation to conceive a child. This is supposedly problematic. I now attempt to show that it is not clear that UP* entails that the indifferent parents are morally required to conceive a child.

To begin it is important to note that ‘significant personal sacrifice’ is actually ambiguous between two interpretations. One of these interpretations entails that the indifferent parents do

not incur a significant personal sacrifice by conceiving a child, whereas a second interpretation entails that they do. On one reading, let's call it the *net sacrifice* (NS) interpretation, someone incurs a significant personal sacrifice by bringing about some outcome rather than another if bringing about that outcome leaves her significantly worse off than she otherwise would have been, all things considered, in the alternative outcomes.

NS: S incurs a significant personal sacrifice to bring about some outcome X if things go worse for S in X, all things considered, than in the alternative outcomes.

According to NS, someone incurs a significant personal sacrifice if the *net cost* to her is sufficiently great. On this reading, determining whether someone incurs a significant personal sacrifice is simply a matter of comparing how someone fares in one outcome with how she fares in another outcome. I think the net sacrifice interpretation correctly provides a *sufficient* condition for when someone makes a significant personal sacrifice, but it fails to capture several characteristic examples of significant personal sacrifice. Take, for instance, the young professional who pinches his pennies and forgoes a social life now so that he might enjoy a great fortune later. He seems to make a significant personal sacrifice now even though he will be ultimately better off than he otherwise would have been. We need a second interpretation of 'significant personal sacrifice' that can accommodate this kind of case. On what I will call the 'gross sacrifice' (GS) interpretation, some action requires a significant personal sacrifice if there are significant costs associated with the action (loss of certain benefits or suffering of some

kind), even if those costs are ultimately compensated for.

GS: S incurs a significant personal sacrifice to bring about some outcome X if S would have to forego (suffer) a significant quantity of goods (bads) that she would not have to forego (suffer) in some available alternative.

GS can accommodate the penny pinching young professional, and a multitude of similar cases. The efforts of the gold medal-winning Olympic gymnast pay off when she finally stands atop the podium. If her earlier efforts did not constitute a significant personal sacrifice then it would not make sense to say that they *pay off* when she wins her event. She forgoes goods earlier so that she will enjoy greater goods later. The fact that this leaves her better off than she otherwise would have been does not entail that forgoing those goods was not a significant personal sacrifice. It seems clear, then, that there are at least two interpretations of ‘significant personal sacrifice’ that correspond to our ordinary concept.

According to the NS interpretation of ‘significant sacrifice’, the indifferent parents do not incur a significant personal sacrifice by conceiving a child. On the NS interpretation, then, UP* entails that the indifferent parents are morally required to conceive the additional child. Because the *net cost* to the indifferent parents is quite small, they do not incur a significant personal sacrifice by conceiving a child. On the other hand, given some plausible assumptions, the GS interpretation of ‘significant sacrifice’ implies that indifferent parents *do* incur a significant personal sacrifice by conceiving a child. To get this result we need only suppose, as seems

plausible, that the indifferent parents would have to forgo the kinds of benefits that most parents must forgo in order to conceive (and raise) an additional child. They give up time, money, energy, and sleep that they will never get back, and they will worry about the well-being of their child for the remainder of their lives. Each of these costs can be expected to detract, albeit temporarily, from the parents' well-being. Furthermore, these costs seem to be unavoidable for people who have children, except in highly unusual circumstances. Even if these forgone goods are later compensated for by the joys associated with parenting, GS entails that indifferent parents incur a significant personal sacrifice. According to the GS interpretation of 'significant personal sacrifice,' UP* entails that indifferent parents, in ordinary cases, are not morally required to conceive an additional happy child. This is the desired result.

Though it is clear that GS is a plausible interpretation of 'significant personal sacrifice', it is perhaps less clear that this kind of personal sacrifice is morally relevant. Can the fact that an agent would have to forgo (suffer) goods (bads)—that will later be compensated for—in order to perform some action that would otherwise be morally required count against an agent being morally required to perform that action? In particular, can the fact that indifferent parents would have to undergo a significant personal sacrifice in the GS sense prevent the moral reason to promote the good from generating a moral requirement to conceive a child?

In fact, there are many cases in which the fact that an agent would forgo (suffer) goods (bads) to perform an action that she would otherwise be morally required to perform seems to override a moral requirement to perform the action *even when* the agent would subsequently receive benefits that fully compensate him for his sacrifice. Consider the following variation of *Conflicted Calvin* (from section 4.2).

Compensated Calvin: Calvin is walking to class when he sees a young boy struggling to stay above water in a partially frozen pond. There is no one else around to hear his flailing. Calvin knows that if he does not save the young boy he will drown. Calvin also knows, however, that he has a unique genetic disorder that causes severely restricted circulation to his lower extremities when he comes in contact with cold water. He knows that he can jump in the pond and successfully save the child, but doing so will cause him to suffer the extremely painful loss of both of his legs. He will be wheelchair-bound for years until his generous leg-insurance policy will pay for bionic replacements legs. These legs will allow him to run faster and jump higher than ever before. The benefits he will receive in the long term, as a result of jumping in the water, will exactly outweigh the cost of the suffering he will endure during the process of and as a result of losing his legs. Calvin does not jump into the pond, because he is unwilling to suffer the pain and temporary loss of the use of his legs despite fact that he will receive an equivalent compensating benefit later. Instead he calls 911, and the struggling child drowns.¹⁶⁹

¹⁶⁹ One might think that *full* compensation for Calvin's sacrifice would simply be to restore the use of his legs. This type of compensation, however, seems to negate his sacrifice in such a way that we no longer have a case of genuine sacrifice. Also, the case is supposed to be analogous to the case of potential parents. The sacrifices that parents make for their children (time, money, emotional and physical stress) are of a different *kind* than the compensating benefits that accompany having children (companionship, satisfaction of one's biological urges, and feelings of accomplishment derived from raising a child to adulthood). A parent rarely regains the sleep, prosperity, or free time that is lost by taking care of a child, and most would not expect to.

Calvin will, eventually, be fully compensated for his sacrifice. In the end, he will not be worse off than he otherwise would have been, so the NS interpretation implies that saving the drowning child does not require Calvin to incur a significant personal sacrifice. This seems wrong. It seems that Calvin would in fact incur a significant personal sacrifice if he were to save the drowning child. Moreover, it seems that Calvin does not act wrongly by failing to save the child precisely *because of* the sacrifice that he would incur. The case thus supports the moral significance of the GS interpretation of ‘significant personal sacrifice’. Personal sacrifices, interpreted as forgoing (suffering) goods (bads) that one would otherwise enjoy—even when those forgone (suffered) goods (bads) will later be compensated for—can override moral requirements.¹⁷⁰

The proposal that significant personal sacrifice, in the GS sense, matters in determining what we are morally required to do might seem to commit us to some highly counterintuitive results. Suppose that if Calvin saves the drowning child he will immediately incur a \$1,000,000 debt. He will subsequently be paid \$1,000,000 once he returns to shore. Suppose that incurring a \$1,000,000 debt is a large enough personal sacrifice to prevent Calvin from being morally required to save the child. On the view of significant personal sacrifice that I just defended, the fact that his debt will be immediately forgiven is irrelevant. Calvin is still not morally required to

¹⁷⁰ It is worth noting that our intuitions about moral requirements to conceive children may be driven by the fact that potential parents are always greatly uncertain about whether the benefits of conceiving a child will really will compensate for the sacrifices. In the real world, we can never be certain of this. When it comes to reproductive cases, we are never sure whether parents are really better off having kids, or even only slightly worse off. We also can never be sure that the addition of the kid to the world is a net gain. Additionally, we hear a claim of obligation in at least a quasi-legal sense, and we are pretty sure that the bad effects of forcing people to have kids will outweigh whatever good effects will come from the kids. Insofar as our intuitions are the result of confusing the legal sense of obligation with the moral sense, we should doubt their usefulness in undermining UP*.

save the child. This seems absurd.

This objection can be overcome if we recognize a distinction between sacrifices involving *instrumental* goods and sacrifices involving *intrinsic* goods. Money is an *instrumental* good; it is valuable only insofar as it is related to an intrinsic good. Pain and frustration associated with the loss of Calvin's limbs in the original *Compensated Calvin* case are *intrinsic* bads; they are bad regardless of their further effects. It is consistent with my analysis of morally relevant significant personal sacrifice offered so far that sacrifices of instrumental goods that have no further effects on the *intrinsic* goods that someone possesses are not morally relevant.

For this reason, a counterexample to my solution to the case of indifferent parents that appeals to instrumental goods (bads) cannot succeed. The sacrifices that indifferent parents make when they conceive a child involve intrinsic goods, and the objection being considered here provides no reason to think that sacrifices of *intrinsic* goods are morally irrelevant. Thus the original *Compensated Calvin* case continues to support the view that the sacrifice required by indifferent parents is morally significant, and that this sacrifice might be great enough to prevent the moral reason to conceive an additional child from generating a moral requirement.

I have argued that even indifferent parents would have to make great personal sacrifices, in the GS sense, in order to raise a child (I imagine any actual parent can attest to this), and these sacrifices might often constitute sufficient countervailing reasons not to conceive a child. Any compensating benefits they will later enjoy do not erase the moral significance of these sacrifices. The important consequence of my remarks in this section is that one can accept UP* and maintain that indifferent parents do not, in ordinary cases, have an obligation to conceive children. But this consequence is consistent with the conclusion that Wilma acts wrongly by

conceiving Pebbles rather than Rocks, for the case of the indifferent parents is importantly different from *Wilma*. Indifferent parents must make significant sacrifices in order to have a child. Unlike indifferent parents, Wilma has already decided to accept the significant sacrifices associated with raising a child. The indifferent parents have not. Wilma must make only a minor sacrifice, on either the NS or the GS interpretation of ‘significant personal sacrifice’, to ensure that her child is not born incurably blind. Wilma’s extra sacrifice to conceive a sighted child is stipulated to be so minor that it is not sufficient to prevent the moral reason to make the set of people better off than the set of people otherwise would have been from generating a moral requirement.

The preceding remarks also provide a way of dealing with a second, supposedly peculiar implication of UP*. Suppose that Wilma has three choices available to her, each of which corresponds with a particular outcome:

- | | | | |
|----|--------------------------|-------------|--------------|
| a) | Conceive a sighted child | Wilma (99) | Rocks (100) |
| b) | Conceive a blind child | Wilma (100) | Pebbles (50) |
| c) | Conceive no child at all | Wilma (50) | |

Wilma will be happiest if she conceives a blind child and least happy if she conceives no child (either a blind or a sighted child would make her life go significantly better). On my view, Wilma has most moral reason to do (a), since the set of people who exist is as well off as possible in (a). But she has *second most* moral reason to do (b), since (b) will be the second best

outcome with regard to the welfare of people who exist. She has *least* moral reason to conceive no child at all. Moreover, conceiving no child at all is the worst option for Wilma. Thus Wilma appears to be permitted to do the thing that she has least moral reason to do (and that is the worst option for her) and yet morally prohibited from doing the thing that she has *more* moral reason to do and that is the best option for her. I have skirted this issue thus far by considering a version of the case in which Wilma has already decided to conceive a child, so her relevant choice is between (a) and (b). We can, however, easily imagine a scenario where Wilma has not yet made her choice to conceive a child, and she of course has the option of changing her decision. In these cases my view seems to face the challenge of explaining why it is permissible for Wilma to do (a) or (c) but not (b). My first response to this challenge is to appeal, once again, to the moral significance of the GS interpretation of ‘significant personal sacrifice.’ Even though Wilma’s life will go significantly better, all things considered, if she conceives a child, there will no doubt be major sacrifices that she must also endure. This would apply to the case of Wilma in the following way: (a) and (b) both require significant sacrifices for Wilma that are associated with conceiving and rearing a child. (c) does not require any of these sacrifices associated with conceiving and then rearing children, and both (a) and (b) result in outcomes that are better for Wilma, all things considered. On the other hand, (a) requires only a negligible sacrifice for Wilma relative to (b). This allows for the possibility that (c) is permissible despite the fact that Wilma has the least moral reason to choose (c). It is permissible because of the significant sacrifices associated with (a) and (b). This also yields the verdict that, given that Wilma has chosen to conceive a child, (a) is morally required relative to (b). This is because the significant personal sacrifice associated with choosing (a) is not greater than that associated with choosing (b). Though it sounds strange to say that (c) is permissible and that (b) is impermissible, provided

that Wilma has stronger moral reasons to choose (b), this makes perfect sense once we understand the specific role of personal sacrifice as a countervailing moral consideration.

V.4.2.B The Right to Reproductive Autonomy

I am inclined to offer a second response to the apparent problem involving Wilma's three options. This response also constitutes a response to the case of indifferent parents. This second response appeals to the significance of reproductive autonomy. *Reproductive autonomy*, sometimes called, *reproductive self-determination*, often refers to a woman's freedom to choose whether or not to birth children at any point before or after a child is conceived. To avoid taking a stance on the complicated issue of abortion, I will use the term more narrowly. By *reproductive autonomy* I have in mind only the freedom to choose whether or not to conceive children in the first place. Even those who hold that abortion ought to be prohibited on the grounds that it is morally impermissible will accept that agents generally have a right to choose whether or not to conceive children in the first place. Reproductive autonomy, understood in this narrow sense, is thought to be of uncontroversial special significance in the legal realm. To see one instance of the presumption in favor of reproductive autonomy, consider the debate (or lack thereof) about parental licensure. Parenting is remarkably similar to other activities that require licensure before someone can engage in them. Like possessing a firearm, practicing medicine, law, or psychiatry, and driving a motor vehicle, each of which requires a license before it is legally performed, parenting is potentially very harmful to others. It requires a certain level of competence, and there are fairly reliable measures for determining someone's competence at performing the activity. Yet the proposal that potential parents should be required to acquire a state-issued

license before they are allowed to conceive children has gained little or no traction in the political realm.¹⁷¹ This indicates that, within the legal realm at least, regulation of reproduction is off limits.¹⁷²

In the legal realm, individuals are given a wide berth to behave as they see fit when it comes to their reproductive choices. I want to take this observation about the legal realm and apply it to the moral realm. How might this go? Well, it seems that within certain limits individuals have reproductive rights (RRs) of the following form

RR: An agent may permissibly choose to conceive or not to conceive a child as long as doing so does not result in _____.

How we fill in the blank will determine the strength and scope of reproductive rights, and it may turn out that reproductive rights are exceptionally strong. The right described by RR is one according to which agents enjoy a particular sphere of autonomy over their behavior that, with the exception of bizarre circumstances, is left untouched by other moral considerations like promotion of the good. Within this sphere one may choose to act as she wishes without influence from other moral considerations. Reproductive rights are certainly thought to have special

¹⁷¹ Hugh LaFollette (1980) proposes that parent licensing is sufficiently similar to other licensed activities and that the state is justified in requiring people to acquire licenses before conceiving children.

¹⁷² Interestingly, most are perfectly willing to accept that adopting parents should be required to go through extensive measures before being legally authorized to adopt a child. See for instance: http://www.adoption-services.org/domestic_adoption_types/adoption_requirements.htm. This is not in obvious tension with a presumption in favor of *reproductive* autonomy, however, since adoption is clearly not reproduction.

significance in the legal realm. Most do not blink at certain legal restrictions on the right to bear arms, the right to freedom of religion, or the right to freedom of expression, yet most will balk at even minimal constraints on freedom of reproduction. This constitutes a *prima facie* reason to think that reproductive rights have special moral significance as well. I do not intend to offer a robust defense of the special moral significance of reproductive rights outlined here. My intent is simply to point out that *in fact*, in ordinary morality, there appears to be widespread agreement that such rights *are* of special moral significance, and that this fact may very well explain the widespread view that people are almost never morally required (not) to conceive children.¹⁷³ The challenge, then, is to incorporate the apparent special moral significance of reproductive rights into UP* to yield the plausible result that neither Wilma nor the indifferent parents is required to conceive a child. It seems to me that this challenge is easily met. Consider my attempt

UP**: There is a moral reason (not) to perform an action *a* if the set of individuals given *A* would be better (worse) off than the set of individuals would be given $\sim A$, and this moral reason generates a moral requirement in the absence of sufficient countervailing reasons, where ‘sufficient countervailing reasons’ may refer to (at least) (a) reasons generated by the fact that performing *A* would require a significant personal sacrifice, (b) reasons generated by the fact that performing *A* would violate a moral constraint (against

¹⁷³ I leave open the possibility that our conviction that no one is ever morally required to conceive a child in virtue of the fact that the child would have a life worth living is a function of confusing legal rights with moral rights. What we mean when we judge that no one can be morally required to have children is that we would not *want a society* in which the state *demand*s that some people conceive children. It also warrants mention that it is not the case in the actual world that anyone would be morally required to have children in virtue of the fact that those children would have lives worth living. In fact people should be donating their money to famine relief.

harming or violating rights), and **(c) reasons generated by the fact that the choice (not) to perform A falls within a special class of choices protected by some sphere of moral autonomy.**

The inclusion of (c) would seem to be *ad hoc* except that above I attempted to show that ordinary morality already includes a provision intended to protect certain choices—reproductive choices in particular—from the influence of other moral and legal considerations. I do not mean to restrict (c) to reproductive choices. Perhaps other kinds of choices fall into the protected category, but, for my purposes it need only be the case that certain reproductive choices fall into that category. In particular, choices about *whether or not* to conceive a child must fall into the protected category for my solution to avoid the implication that Wilma or the indifferent parents are morally required to conceive a child rather than no child at all. Note that the inclusion of (c) allows us to retain the result that Wilma acts wrongly by conceiving Pebbles rather than Rocks, because this choice is not a choice about *whether* to conceive but rather about *when* to conceive. It is plausible, and at least coherent, to maintain that the former but not the latter kind of choice falls under the special class of choices protected by some sphere of moral autonomy.¹⁷⁴

¹⁷⁴ There is a puzzle here about how to explain an apparent moral asymmetry that seems to pose a problem for the idea of a reproductive right. It seems that someone can be morally required not to conceive a child in cases where that person would have a life that is ‘worth not living’—a life in which the intrinsically bad states outweigh the good. At the same time it seems that one is never morally required to conceive a child merely because that person \ have a life worth living.

V.4.2.C A Dilemma: How Demanding is Morality?

My earlier replies in V.4.2.1 and V.4.2.2 were designed to show that it is not clear that UP* implies that indifferent parents are required to conceive a child. In what follows I will attempt to put intuitive pressure back on my opponent by reiterating some implications of rejecting UP*. Let's say that we reject UP* because the principle entails that potential parents sometimes have an obligation to conceive. What do we say, then, about our reactions to cases like *Wilma*? We think she has done something wrong by causing the blind child to exist instead of the sighted child. Yet it is very difficult, if we reject UP*, to explain why this is the case. Our reactions to thought experiments are generally taken to reflect *prima facie* facts about morality, but if we accept the conclusion of the Non-identity Problem then we must also accept that our reaction to *Wilma* is simply mistaken. The Non-identity Problem is, in this respect, a counterexample to any argument that leads to it. UP* provides a theoretical justification of our reactions to non-identity cases but—supposing that the view implies that parents are sometimes morally required to conceive children—at the (small) cost of accepting a slightly more demanding morality.

Here we are faced with a tradeoff of counterintuitive implications. If my UP* is correct, morality may turn out to be slightly more demanding than we thought. On the other hand, if the conclusion of the Non-identity Problem is correct then morality is concerned solely with wrongs to existing individuals and is significantly less demanding than we thought. The prospect of having fewer moral duties may seem appealing. We should not, however, lose sight of the fact that the conclusion of the Non-identity Problem is morally repugnant. Accepting the conclusion of the Non-identity Problem would commit us to seriously reforming our moral beliefs about

cases like *Wilma* and the *wealthy society*. Let us turn one more time to the case involving society's choice between depleting and conserving our natural resources, briefly discussed in the first chapter:

Depletion: Imagine that we have a simple choice to make. We can continue depleting the Earth's resources to the point that the environmental damage will be irreversible and self-perpetuating. As a result, the welfare of some future generation will be so low that their lives will be barely worth living. We will enjoy some nice conveniences (AC, driving to work, leaving the lights on, etc.) by continuing to deplete the Earth's resources. These conveniences, however, will not come close to comparing to the major illnesses, famine and over-population suffered by the future generation. OR, we can begin to conserve our resources and give up certain conveniences. If we choose the second course of action future generations will enjoy welfare levels equal to ones enjoyed by current people. If the identities of the people in the future generation would be the same regardless of our actions then choosing to deplete would be morally monstrous. Can it really make all the difference between moral permissibility and moral monstrosity that the identities of future individuals will be different if we choose to deplete rather than conserve?

If the individuals who will exist if we choose to deplete our resources will have lives just

barely worth living, and we accept the conclusion of the Non-identity Problem, then we should believe that there is *nothing wrong* with us choosing to deplete. But most believe that it is deeply wrong to deplete our resources in this case. Yet, if we accept the conclusion of the Non-identity Problem, we must accept this implication. In the first chapter I said that the fact that a solution to the Non-identity Problem would solve the Non-identity Problem cannot count in favor of preferring that solution over *other* solutions. This was the motivation for including the *Non-ad hocness Criterion* for a successful solution. However, in chapters II and III I eliminated all other popular candidate solutions from contention. Premises (1-4) of NIPA, I have argued, are true. We are now left with accepting P5 of NIPA, thereby accepting the conclusion of NIPA, or rejecting P5 by appealing to the Welfare argument. Given the dialectical situation, it seems fair at this point to count the conclusion of NIPA as a reason to accept the Welfare Argument. Even if the Welfare Argument implies, despite my attempts to avoid the implication, that some potential parents are morally required to have children, the implication is not grounds for rejecting the Welfare Argument, especially if the alternative is accepting the conclusion of the Non-identity Problem.

VI

Practical Applications

In Chapter I, I described the circumstances under which the Non-identity Problem might arise, presented some criteria that a solution needs to satisfy in order to be successful, and I pointed out the significance of the Non-identity Problem for several practical ethical issues. In Chapters II and III I raised some serious difficulties for solutions to the Problem that have been offered in the literature. Chapters IV and V constituted the development and defense of my solution to the Non-identity Problem. In these chapters I argued that a solution to the Non-identity Problem is available if we accept the existence of a pro tanto reason to promote the good and that this reason applies to actions that affect both the quality of life and the identities of the people who will come to exist. One might be left wondering, however, whether my solution is useful for solving any practical moral issues. If I have solved the problem in the abstract, but my offers no guidance for the kinds of real-world issues for which the Non-identity Problem arises, then the discussion thus far is open to the criticism that it is mere philosophical musing. For this reason, in this Chapter I turn to explaining the practical application of my solution for various issues in applied ethics.

VI.1 Environmental Policy

One issue concerns the moral obligation to conserve rather than deplete our natural resources for the sake of benefitting future generations. Our position today may not be altogether

different from the position of the wealthy society that I discussed in describing what I called the ‘eventual’ version of the Non-identity Problem. Most believe that if today’s society, by suffering some minor inconvenience now, can avoid making future generations significantly worse off, then today’s society ought to suffer that minor inconvenience. In Chapter I, I described how the Non-identity Problem arises for this view. The trouble was that the identities of the individuals who will come to exist in the future might very well depend on the kind of environmental policy we adopt now. If (i) the people who will come to exist if we deplete will have different identities than the people who will come to exist if we conserve, and (ii) the people who will exist if we deplete our resources will have lives that are worth living, then the Non-identity Problem rears its head. What, if anything, does my solution to the Non-identity Problem say about this practical application of the Non-identity Problem? To answer this question we must consider, one more time, Unrestricted Pro tanto (UP*), and apply that principle to the current situation with respect to resource conservation.¹⁷⁵ Before doing so, however, it is important to note that the conditions of the real-world, unlike philosophical examples, are empirically messy and, as a result, epistemically opaque. For instance, there is not an entirely settled view about the expected negative effects of various potential environmental policies on future generations. There appears to be a fairly strong consensus that elevated levels of CO₂ emissions, a result of industrialization and increased human consumption, will negatively affect future generations by causing catastrophic climate change, but some (albeit a minority) have denied this.¹⁷⁶ There is also no

¹⁷⁵ Though UP** was the final formulation of Unrestricted Pro tanto in the previous chapter, it will suffice for my purposes here to consider the simpler version of the principle.

¹⁷⁶ The Intergovernmental Panel on Climate Change is widely recognized by many scientific organizations as the most reliable source of information about the state of global climate change. See “The Science of Climate Change,” a statement attributed to sixteen international science organizations, for an endorsement of the IPCC’s methods. For some other sources that support the claim that climate change is occurring as a result of human behavior see J.

settled view about the *costs* of the policies that would avoid those negative effects on people who are alive today.¹⁷⁷ I do not wish to underplay the importance or difficulty of solving these empirical disputes, but solving them is not the job of philosophers. In answering the question, then, I will suppose that some environmental policies will have greater negative effects on future generations than others.¹⁷⁸ I will also consider but not attempt to settle the issue of what sacrifices these policies would require by people alive today. Consider UP* one more time.

Unrestricted Pro tanto* (UP*): There is a moral reason (not) to perform an action *A* if the set of individuals given *A* would be better (worse) off than the set of individuals would be given $\sim A$, and this moral reason generates a moral requirement in the absence of sufficient countervailing reasons, where ‘sufficient countervailing reasons’ refers to (at least) (a) reasons generated by the fact that performing *A* would require a significant personal sacrifice and (b) reasons generated by the fact that performing *A* would violate a moral constraint (against harming or violating rights).

J. McCarthy (2001); National Academy of Sciences Committee on the Science of Climate Change, *Climate Change Science: An Analysis of Some Key Questions* (National Academy Press, Washington, DC, 2001); Arctic Climate Impact Assessment, *Impacts of a Warming Arctic* (2004); National Oceanic and Atmospheric Administration, *State of the Arctic Report* (2006). There are, of course, a handful of dissenting opinions: the American Association of Petroleum Geologists have released studies that seem to cast doubt on the claim that human consumption has played a role in rising global temperatures. See *Geological Perspectives of Global Climate Change* (2001) for one such study. See also Dennis Avery’s “500 Scientists Whose Research Contradicts Man-Made Global Warming Scares,” sponsored by The Heartland Institute (2007).

¹⁷⁷ See *Stern* (2007) for a pessimistic estimate of the economic costs required to curb climate change. According to one study conducted by the Green Growth Action Alliance on behalf of the World Economic Forum (2013), curbing the world’s use of fossil fuels to avoid the negative effects of climate change (e.g., worsening floods, heat waves, rising sea levels) would cost \$700 billion per year.

¹⁷⁸ This much seems plausible. For instance, it would be very bad for future generations if, through the use of nuclear weaponry, people today turned the world into a radioactive wasteland. A policy that had this result would clearly have a *more negative impact* on future generations than one that did not.

Suppose that some resource policy would have the fewest negative effects on future generations relative to the available alternative policies. Let us call this the ‘optimal policy’. Suppose also that the conditions that give rise to the Non-identity Problem obtain. The verdict of UP* regarding contemporary society’s obligation to choose the optimal policy rather than the alternatives depends on the *cost* of that policy for people alive today (and the cost for people who will come to exist regardless of which policy we choose) and on whether that policy would violate any *constraints*. One thing that is clear is that UP* entails that there is a moral *reason* to choose the optimal policy and that this reason generates a moral requirement in the absence of sufficient countervailing reasons. This means that, *if* it is possible to choose the optimal policy, and the costs of that policy for people alive today (and for people who will come to exist regardless of which policy we choose) are relatively minor, and the policy would not violate any moral constraints, then UP* entails that there is a moral obligation to choose the optimal policy.

Consider one policy that might satisfy this description.¹⁷⁹ Suppose that society has the option to impose an emissions-curbing policy that would require its citizens either to travel by carpooling or by riding some form of public transportation during peak commuting hours. Let us suppose that such a policy, implemented on a sufficiently large scale in the most polluting societies, would have a significant impact on CO₂ emissions. Now suppose that decreasing CO₂ emissions would have a positive effect on the welfare of future generations and that the conditions that give rise to the Non-identity Problem obtain. In other words, the future generations who will come to exist if we enact the emissions-curbing policy (let’s call them ‘A’) will be better off than but non-identical with the future generations who will come to exist if we

¹⁷⁹ I do not mean to claim that the policy I describe is *actually* the optimal policy with regard to the welfare of future generations. I simply use it as an example to illustrate the verdicts of UP*

do not enact the policy (let's call them 'B'). This supposition is plausible when we notice that even minor changes in the timing of conceptions can have a significant impact on the identities of who comes to exist. Presumably, if everyone in a society starts carpooling or taking public transit to and from work rather than driving, then this will have at least minor effects on the timing of conceptions. According to UP*, there is a moral reason for today's society to enact the emissions-curbing policy in virtue of the fact that if we choose the emissions-curbing policy then A would come to exist and would be better off than B. But does UP* entail a moral requirement to choose the emissions-curbing policy? As noted above, the answer to this question depends on (i) the cost to people alive today of enacting the policy and on (ii) whether enacting the policy would violate any constraints. With regard to (i), it seems that the costs of taking public transit or carpooling in order to travel during peak commuting hours would constitute a minor inconvenience for some (those with easy and affordable access to public transit) but a significant sacrifice for others (those who live in rural areas). On this picture UP* justifies enacting an emissions-curbing policy that requires only those people with easy and affordable access to public transportation or carpooling to take advantage of that access. Of course the possibility remains that enacting an emissions-curbing policy would result in unforeseen but significant economic costs. I leave it to economists to determine the significance of the economic costs of various environmental policies. In the absence of significant economic costs it seems that the kind of emissions-curbing policy described above could be enacted without significant costs to some living members of society (or to future members of society who will come to exist regardless of which environmental policy we choose). Supposing that the costs to the members of society affected by the emissions curbing policy are not significant, we must now ask (ii) whether enacting the policy would violate any constraints. Different theories will yield different

answers.¹⁸⁰ I will simply point out that relatively uncontroversial policies exist in the United States today that are no more and no less restrictive than a policy requiring people to travel by carpool or by public transit. Companies are required to ensure that their facilities satisfy EPA emissions standards. Vehicles are also required to pass emissions inspections before they can be eligible for registration. Vast expanses of public lands are closed to motor vehicles or logging for the sake of preserving their natural condition. These policies significantly restrict what a person may do with her private property, and some of them, arguably, have significant economic costs. Most will nonetheless take some subset of these policies to be uncontroversial. Indeed, the most plausible justification for these policies is that they will benefit both current *and* future generations. While I do not want to delve into debates on the nature of constraints, there are environmental policies in place today in the United States that are relevantly similar to the emissions-curbing policy described above, and few people think that these policies violate any constraints. So, given certain empirical conditions, UP* entails that there is a moral obligation for society today to enact policies, like the emissions-curbing policy described above, that will have a positive effect on future generations. The question remains whether these empirical conditions obtain.

VI.2 Preimplantation Genetic Diagnosis

My solution has more decisive moral implications for the use of certain genetic technologies. In Chapter I, I discussed Preimplantation Genetic Diagnosis (PGD) and

¹⁸⁰ A strong libertarian, for instance, might object to this kind of policy on the grounds that it violates individuals' rights to their property by preventing them, say, from using their motor vehicles as they see fit.

Preimplantation Genetic Screening (PGS), kinds of technology the application of which can affect both the identity of the people who come to exist as well as the quality of life of the people who come to exist. PGD refers to a process of testing embryos fertilized using in vitro fertilization for genetic abnormalities before pregnancy. This is typically done when at least one of the parents has a known genetic abnormality. Embryos that test positive for the abnormality will be rejected in favor of ones that do not. PGS uses similar technology, but it is used by parents neither of whom is known to have a genetic abnormality. PGD might be useful for the purposes of testing an embryo for Huntington's Disease (HD), an uncommon but devastating genetic abnormality found in some humans. The symptoms associated with HD typically start to appear in a person's 30s or 40s. The first symptoms involve changes in mood and cognition. These are followed by problems with coordination that progress until they are completely debilitating. Later symptoms include difficulty chewing and swallowing, and seizures. Eventually a person with HD suffers from full dementia and requires full-time care until death.¹⁸¹ It is a horrible disease, but it is also peculiar in a way that makes it a candidate for a kind of disease that showcases the Non-identity Problem; a person who suffers from HD can enjoy 30-40 years of a very satisfying life before experiencing any symptoms. Now let us consider the situation of a person who plans to conceive a child but who knows that she carries the genetic abnormality responsible for HD. She is considering whether she ought to use PGD to ensure that the child she has does not have the disease. There are two moral questions pertaining to this person's situation. First, is she morally required to use PGD? Second, if she uses PGD is she morally required to choose to implant an embryo that does not have the genetic abnormality responsible for HD?

¹⁸¹ Montoya, Price, Menear, Lepage (2006)

Before turning to the answers my solution gives to these questions, note that the case we are considering meets the conditions that generate the Non-identity Problem. The first condition is that if one embryo has the genetic abnormality responsible for HD and another does not, then the people who would develop from these embryos are non-identical. This supposition is easy enough to demonstrate, for suppose for a moment that two embryos would develop into one and the same person. Both embryos *could be implanted* in distinct uteri. If they both developed to become children, then one person would have HD and the other person would not. But then one and the same person would both have HD and not have HD. This is a contradiction, so we can safely reject the claim that the person who would develop from the embryo that has the genetic abnormality would be the same person as the person who would develop from the abnormality-free embryo.

The second condition is that the person who would come to exist if the embryo with the genetic abnormality were implanted would have a life worth living. It is perhaps more difficult to show that this condition obtains, since the symptoms of HD clearly diminish the quality of someone's life. Note, however, that most HD victims can enjoy 30-40 years of a normal life before they even begin to suffer the early symptoms of the disease. Suppose that it is only in the final stages of the disease that a person's life becomes worse than no life at all. The life expectancy is 20 years after the onset of the first visible symptoms.¹⁸² This means that a person with the disease enjoys 30-40 years of an unimpaired life that are then followed by 20 years of an increasingly difficult and presumably less happy life. Since the early stages of the disease are not incapacitating, the first several years after the onset of HD symptoms are only slightly worse

¹⁸² Ibid

than the first 30-40 years of a person's life. The final stages of their life would have to be pretty awful to justify the claim that the person's life was on balance not worth living. I also suspect that if you were to ask someone with HD if they regretted coming to exist they would say that they do not. So, supposing (i) that the person who would develop from the embryo with the genetic abnormality responsible for HD is not identical with the person who would develop from the embryo without the abnormality and (ii) that the person who will develop from the embryo with the genetic abnormality responsible for HD will have a life worth living, then the conditions that give rise to the Non-identity Problem obtain. Now we can turn to my solution's answers to the two questions posed above. Does the woman have an obligation to use PGD, and if she does use PGD does she have an obligation to choose the embryo that does not have the genetic abnormality responsible for HD?

UP*, with one minor complication, seems clearly to entail that the woman is morally required both to use PGD and, given that she does use PGD, to choose an embryo that does not have the genetic abnormality responsible for HD. Let us answer the second conjunct first. Let's suppose that the child who would develop from the non-HD embryo would be twice as well off as the child who would develop from the HD embryo. The child with HD would enjoy 50 units of welfare over his whole life whereas the child without HD would enjoy 100 units of welfare. Given that she has chosen to use PGD, she is able to guarantee to a reasonable degree that the child without HD comes to exist. Because the child without HD would be better off than the child with HD, UP* entails that she has a moral reason to bring the child without HD into existence and that this reason generates a moral requirement in the absence of sufficient countervailing reasons. Do sufficient countervailing reasons obtain? If bringing the child without HD into existence would result in significant costs to her that she could avoid by bringing the

child with HD into existence instead, then UP* does not entail that she is morally required to implant the non-HD embryo. There is no reason to think that such conditions obtain in ordinary cases involving PGD, however, so this would be an extremely unlikely scenario. If bringing the child without HD into existence would violate a constraint (say by infringing someone's rights), then, again, UP* does not entail that the woman is morally required to do so. The only candidate for a victim of a rights infringement in this case, however, is the embryo that is not chosen for implantation (the HD embryo in this case). Does the unselected embryo have its right infringed? This is an especially tricky question. Answering it requires delving into the moral status of embryos, and this requires determining an embryo's relationship to the person it will become if it is allowed to develop, etc. Of course, if the HD embryo would have its right infringed by not being implanted, then the non-HD embryo would also have its right infringed if it were not implanted. Since the woman cannot, given that she has chosen PGD, avoid infringing one of the embryos' rights, and since ought implies can, then it seems that the fact that her act will infringe a right should not factor into the moral verdict regarding which embryo she ought to implant. For the purposes of this first question, then, I will bracket off the possibility that the embryo might have its right infringed. It seems, then, that UP* entails that the woman is morally required to implant the embryo *without* the genetic abnormality responsible for HD.

A separate question is whether UP* requires the woman to use PGD in the first place. To answer this question we must employ the notion of expected value. There is a 50% chance that a person conceived by a parent affected by the abnormality responsible for HD will inherit the abnormality. We are still supposing that the child with HD's life would be half as good as the life her child without HD would have had, other things being equal. The non-HD child's life would have 100 units of welfare and the HD child's life would have only 50. This means that, if the

woman chooses not to use PGD in the first place, there is a 50% chance that the child she has will have a life with 50 units of welfare rather than 100. This means that the expected welfare of the person who will come to exist if she conceives without using PGD is 75 units. If, on the other hand, she uses PGD *and* selects the non-HD embryo for implantation, then the expected welfare of the person who comes to exist is 100 units. UP* thus entails that, in the absence of sufficient countervailing reasons, the woman is morally required to choose to use PGD and to then choose to implant the non-HD embryo. Do sufficient countervailing reasons obtain? There is an interesting question that arises here if we think that embryos have rights or can be harmed. By choosing to use PGD, fertilizing two embryos (one with the HD abnormality and one without) the woman guarantees that one of them will not be implanted. But if embryos can be harmed and have rights, then by choosing to use PGD the woman guarantees that she will harm (by preventing the continued existence of) and thereby violate the rights of one of the embryos. If embryos have rights, then it might turn out that UP* does not entail that the woman is morally required to use PGD after all. In fact, it might entail that she has a moral obligation not to use PGD. Exploring the issue of the moral status of embryos is interesting and important, but it is not the subject of the discussion here; I will bracket off this question and restrict my conclusion to the following claim: If the woman does not guarantee that she will violate the rights of one of the embryos by using PGD, then UP* entails that she is morally required to use PGD, assuming the personal sacrifice to her of doing so is minimal.

VI.3 Reproductive Cloning

Reproductive cloning is a method of producing offspring using the exact genetic DNA of someone who already exists. Individuals produced using this method would be genetically identical to their 'parent'. Though the method continues to face scientific obstacles, it is likely to be a viable reproductive option in the near future. In cases where, say, one partner was to die before a couple had any children, the other partner might wish to clone the deceased in order to have a child with her partner's genes. Cloning might also be a desirable option for infertile couples or for same-sex couples who want to avoid having offspring to whom a third party donor has contributed nuclear DNA. A common objection to reproductive cloning is that the cloned offspring would suffer significant psychological harms as a result of the parents 'forcing a shared identity' on the offspring. The concern is that cloned offspring will be psychologically harmed because of the pressure they will feel to emulate the life or personality of the parents from which they were cloned.¹⁸³ The trouble with this harm-based argument against reproductive cloning is that reproductive cloning cases seem to possess the features that give rise to the Non-identity Problem, and we have seen that harm-based solutions to the Non-identity Problem do not succeed. The Non-identity Problem arises for reproductive cloning cases because (i) the offspring who are conceived through reproductive cloning would not otherwise have existed, and (ii) offspring who are the product of cloning, in most cases at least, will have a life that is not worse than non-existence.¹⁸⁴

¹⁸³ The President's Council on Bioethics (2002). See Havstad (2010) for further discussion of the claim that it is impermissible to use cloning technology if it will result in psychological harm to the clone.

¹⁸⁴ Brock (2003), Gillon (2003), and Harris (2005) raise this worry for harm-based arguments against reproductive cloning.

I showed in Chapter II that the Non-identity Problem resists solutions that appeal to harm. This stubborn feature of the Non-identity Problem combined with the fact that instances of reproductive cloning are instances of the Non-identity Problem seems to be a decisive problem for harm-based objections to human reproductive cloning. If there is a moral objection to reproductive cloning it does not stem from the harmful effects of the technology on the resulting offspring.¹⁸⁵ Whether UP* entails that reproductive cloning is morally objectionable will depend on the details of particular cases. Consider once again the two kinds of cases where a pair of potential parents might wish to use reproductive cloning. First, someone may wish to conceive offspring that are genetically related to a deceased partner. Second, an infertile or same-sex couple may wish to conceive a child without using genetic material from a third party donor.

Let us suppose that someday human reproductive cloning will be technologically viable. A woman who wishes to conceive a child that is genetically related to her deceased partner is considering whether or not to try reproductive cloning. She has two courses of reproductive action available to her. She can conceive using the genetic material from a third party donor or she can create a clone of her deceased partner using his genetic material. If she chooses to use reproductive cloning then she will conceive a child that is genetically related to her deceased partner, but due to the psychological suffering associated with being a clone, this child will be worse off than the child she would conceive if she were to use a sperm donor. Let's call the child conceived using reproductive cloning 'Jack' and the child conceived using a donor's sperm 'Jill.'

¹⁸⁵ If the opponent of human reproductive cloning claims that the harmful effects of human reproductive cloning are a result of the parents *treatment* of the cloned offspring, then the Non-identity Problem doesn't arise for *this kind* of harm-based objection. But, this kind of harm-based objection to human reproductive cloning is not actually an objection to human reproductive cloning per se. It is, rather, an objection to treating one's cloned offspring a certain way.

UP* entails that the fact that that Jill would be better off than Jack gives her a moral reason to conceive Jill that generates a moral requirement in the absence of sufficient countervailing reasons. How strong this reason is will depend on how much better off Jill would be than Jack, so let's suppose (as we did when comparing the HD embryo with the non-HD embryo) that Jill would be twice as well off as Jack because of the psychological trauma that Jack would suffer as a result of being a clone. The woman in this case, then, has reasons of the same strength to conceive Jill that the woman had in the Huntington's Disease case to implant the non-HD embryo. Whether UP* generates a moral requirement for the woman to conceive Jill rather than Jack will depend on the countervailing reasons that the woman has to conceive Jack. This, I suspect, is where typical cases of reproductive cloning will diverge morally from typical cases of PGD. In cases involving the use of PGD to avoid, say, implanting an embryo with the genetic abnormality responsible for HD, it is unlikely, barring complications mentioned earlier, that a potential parent would need to undergo a significant personal sacrifice or to violate any constraints in order to implant a non-HD embryo rather than an HD embryo. In fact, it is likely to be in *parent's interests* to implant the non-HD embryo because of the burdens associated with being the parent of someone who suffers from HD.

The case of reproductive cloning in order to preserve genetic ties to a deceased partner is more complicated, especially if we suppose that the woman who is deciding whether to conceive Jack or Jill has a strong interest in conceiving a child who will be genetically related to her late partner. Suppose that it matters *deeply* to her that her child be genetically related to her deceased partner. It seems to be a difficult question, then, whether the woman has sufficient countervailing reasons to conceive Jack rather than Jill given the importance to her of conceiving a child that is genetically related to her deceased partner. It might turn out, then, that conceiving Jill would

constitute a significant sacrifice for the woman. According to UP*, significant personal sacrifice generates a sufficient countervailing reason against being morally required to perform some action. Because it is unclear whether conceiving Jill rather than Jack would constitute a *significant* personal sacrifice for the woman, it is unclear whether UP* would say that the woman is morally required to conceive Jill rather than Jack. I suspect that similar considerations are relevant in the case of same-sex or infertile couples who prefer not to conceive a child using genetic material from a third party donor. In cases like these it seems that the child's genetic relationship to (one of) the parents matters deeply to the couple, otherwise they would not be considering the reproductive cloning option in the first place. If that is true, then again it is not clear that UP* entails that the couple is morally required to conceive the non-cloned child even if that child would be better off than the cloned child the couple prefers to conceive.

UP* does not provide a decisive verdict regarding the moral permissibility of reproductive human cloning, but I do not think this indicates a flaw in the principle. It nonetheless gives important guidance regarding the kinds of considerations that are morally relevant to answering what is a complex moral question. I take it to be a virtue of UP* that it accommodates nebulous empirical facts such as the expected disparity in welfare between a cloned person and a person conceived using conventional fertilization and the expected personal sacrifice associated with conceiving a child that is not fully biologically yours. A moral theory should not be expected to generate a moral verdict until all of the facts are in. Once the empirical facts of reproductive cloning are clear, UP*'s moral verdict about particular cases of reproductive cloning will be clear as well.

VI.4 Conclusion

The aim of this chapter has been to show that my solution to the Non-identity Problem, far from being mere philosophical musing, can guide our decision-making about some very difficult practical moral issues concerning future generations. My solution to the Non-identity Problem provides clear moral guidance, once the empirical facts are settled, for issues including environmental policy, Preimplantation Genetic Diagnosis, and reproductive human cloning. My solution to the Non-identity Problem provides guidance for other issues as well, including the debate about reparations for slavery and various issues in animal ethics.¹⁸⁶ My solution to the Non-identity Problem thus has significant theoretical value as well as practical importance. We must consider how well actual people will fare as well as how well possible people would have fared in determining what agents are morally required to do. The quality of the lives of future people matters morally even if the actions we perform now determine who comes to exist in the future to live those lives.

¹⁸⁶ For instance, many people think that puppy mills are morally objectionable. But if we think that the animals conceived in puppy mills would not have otherwise existed and that they have lives worth living then the Non-identity Problem arises.

References

- Adams, Robert. 1979. "Existence, Self-Interest, and the Problem of Evil." *Nous*, vol. 13, pp. 53-65.
- Anglin, Bill. 1977. "The Repugnant Conclusion." *The Canadian Journal of Philosophy*, vol. 7, no. 4, pp. 745-754.
- Arctic Climate Impact Assessment. 2004. *Impacts of a Warming Arctic*.
- Arrhenius, Gustaf. 2003. "The Person-Affecting Restriction, Comparativism, and the Moral Status of Potential People." *Ethical Perspectives*, vol. 10, no. 3-4, pp. 185-195.
- Arrhenius, Gustaf. 2009. "Can the Person-Affecting Restriction Solve the Problems in Population Ethics?" M.A. Roberts, D.T. Wasserman (eds.), *Harming Future Persons: Ethics, Genetics, and the Non-Identity Problem*, International Library of Ethics, Law, and the New Medicine.
- Benatar, David. 2008. *Better Never to Have Been: The Harm of Coming Into Existence*. Oxford.
- Boonin, David. 2008. "How to Solve the Non-Identity Problem." *Public Affairs Quarterly*, vol. 22, pp. 127-157
- Bradley, Ben. 1998. "Extrinsic Value." *Philosophical Studies*, vol. 91, pp. 109-126.
- 2011. *Well Being and Death*. Oxford.
- Brock, Dan W. 1995. "The Non-Identity Problem and Genetic Harms -- The Case of Wrongful Handicaps." *Bioethics*, vol. 9, nos. 3-4, pp. 272-75.

----2003. "Cloning Human Beings: An Assessment of the Ethical Issues Pro and Con." In *Contemporary Issues in Biomedical Ethics*, TL Beauchamp and L Walters eds.

----. 2005. "Preventing Genetically Transmitted Disabilities while Respecting Persons with Disabilities." In David Wasserman, Jerome Bickenbach and Robert Wachbroit, eds., *Quality of Life and Human Difference: Genetic Testing, Health Care, and Disability*, pp. 67-100.

Broome, John. 1999. *Weighing Goods*. Oxford.

----2004. *Weighing Lives*. Oxford.

Bigelow, John, Campbell, John and Pargetter, Robert. 1990. "Death and Well-Being." *Pacific Philosophical Quarterly*, vol. 71, pp. 119-40

Brock, Dan W. 2005. "Shaping future children: Parental rights and societal interests." *Journal of Political Philosophy*, vol. 13, no. 4, pp. 377-398.

Buchanan, Allen. E, Brock, Dan W., Daniels, Norman, and, Wikler, Daniel. 2000. *From Chance to Choice: Genes and Social Justice*. Cambridge: Cambridge University Press.

Avery, Dennis. 2007. "500 Scientists Whose Research Contradicts Man-Made Global Warming Scares." Sponsored by The Heartland Institute.

Bykvist, Krister. 2007. "The Benefits of Coming into Existence." *Philosophical Studies*, vol. 135, pp. 335-362.

Chang, Ruth (ed.) (1997). *Introduction, Incommensurability, Incomparability, and Practical Reasoning*. Harvard University Press.

Feinberg, Joel. 1986. "Wrongful Life and the Counterfactual Element in Harming." *Social Philosophy and Policy*, vol. 4, pp. 145-178.

Feit, Neil. 2002. "The Time of Death's Misfortune." *Nous*, vol. 36, pp. 359-383.

Feldman, Fred. 1991. "Some Puzzles About the Harm of Death." *The Philosophical Review*, vol. 100, no. 2 April, pp. 205-227.

----1992. *Confrontations with the Reaper*. Oxford University Press: New York.

Forbes, Graeme. 1980. "Origin and Identity." *Philosophical Studies*, vol. 37, pp. 353-362.

Foot, Philippa. 1985. "Utilitarianism and the Virtues." *Mind*, vol. 94, pp. 196-209.

----1994. "Killing and Letting Die." Reprinted in Bonnie Steinbock and Alastair Norcross, *Killing and Letting Die*. New York: Fordham University Press, pp. 280-289.

Gert, Joshua. 2003. "Requiring and Justifying: Two Dimensions of Normative Strength."

Erkenntnis, vol. 59, pp. 5-36.

Gillon, Ranaan. 2003. "Human Reproductive Cloning: A Look at the Arguments against it and a Rejection of most of them." In *Contemporary Issues in Biomedical Ethics*, TL Beauchamp and L Walters eds.

Hanser, Matthew. 1990. "Harming Future People." *Philosophy and Public Affairs*, vol. 19, no. 1, pp. 47-70.

- 2008. "The Metaphysics of Harm." *Philosophy and Phenomenological Research*, vol. 77, no. 2, pp. 421-450.
- 2009. "Harming and Procreating." M.A. Roberts, D.T. Wasserman (eds.), *Harming Future Persons: Ethics, Genetics, and the Non-identity Problem*, International Library of Ethics, Law, and the New Medicine, pp. 179-199.
- Hare, Caspar. 2007. "Voices from Another World: Must We Respect the Interests of People Who Do Not, and Will Never, Exist?" *Ethics*, vol. 117 April, pp. 498-523.
- Harris, J. 2005. "The Poverty of Objections to Human Reproductive Cloning." In *Contemporary Debates in applied Ethics*, AI Cohen and CH Wellman eds. Blackwell.
- Harman, Elizabeth. 2004. "Can We Harm and Benefit in Creating?" *Philosophical Perspectives*, vol. 18, pp. 89-113.
- Harman, Elizabeth. 2009. "Harming as Causing Harm." M.A. Roberts, D.T. Wasserman (eds.), *Harming Future Persons: Ethics, Genetics, and the Non-Identity Problem*, International Library of Ethics, Law, and the New Medicine.
- Havstad, Joyce. 2010. "Human Reproductive Cloning: A Conflict of Liberties." *Bioethics*, vol. 24, pp. 71-77.
- Hershenov, David. 2007. "A More Palatable Epicureanism." *American Philosophical Quarterly*, vol. 44, pp. 170-80.
- Herstein, Ori. 2013. "Why 'Non-existent People' Do Not Have Zero Well-being but No Well-being at All." *Journal of Applied Philosophy*, vol. 30, no. 2, pp. 136-145.

Heyd, David. 1988. "Procreation and Value: Can Ethics Deal with Futurity Problems?"

Philosophia. vol.18, pp. 151-170.

----2009. "The Intractability of the Non-Identity Problem." M.A. Roberts, D.T.

Wasserman (eds.), *Harming Future Persons: Ethics, Genetics, and the Non-Identity Problem*, International Library of Ethics, Law, and the New Medicine.

Holtug, Nils. 1998. "Egalitarianism and the Levelling Down Objection." *Analysis*, vol. 58, pp. 166-174.

----2001. "The Value of Coming Into Existence." *The Journal of Ethics*, vol. 5, pp. 361-384.

----2009. "Who Cares about Identity?" M.A. Roberts, D.T. Wasserman (eds.), *Harming Future Persons: Ethics, Genetics, and the Non-Identity Problem*, International Library of Ethics, Law, and the New Medicine.

Huemer, Michael. 2008. "In Defence of Repugnance." *Mind*, vol. 117, pp. 899-933.

Kagan, Shelly. 1989. *The Limits of Morality*, Oxford University Press.

Kamm, Frances. 1992. "Non-Consequentialism, the Person as an End-in-Itself, and the Significance of Status." *Philosophy and Public Affairs*, vol. 21, pp. 354-389.

----2007. *Intricate Ethics*, Oxford University Press.

Kass et al. 2002. "Human Cloning and Human Dignity: An Ethical Inquiry." *The President's Council on Bioethics*.

Kavka, Gregory. 1982. "The Paradox of Future Individuals." *Philosophical Perspectives*, vol. 11, no. 2, pp. 93-112.

Klocksien, Justin. 2012. "A Defense of the Counterfactual Comparative Account of Harm." *American Philosophical Quarterly*, vol. 48, no. 4, pp. 285-300.

Kripke, Saul. 1980. *Naming and Necessity*, Blackwell.

LaFollette, Hugh. 1980. "Licensing Parents." *Philosophy and Public Affairs*, vol. 9, no. 2, pp. 182-97.

Lamont, Julian. 1998. "A Solution to the Puzzle of When Death Harms Its Victims." *Australasian Journal of Philosophy*, vol. 76, pp. 198-212.

Levy, Neil. 2002. "The Apology Paradox and the Non-Identity Problem." *The Philosophical Quarterly*, vol. 52, no. 208 July, pp. 358-368.

Lewis, David. 1986. *On the Plurality of Worlds*. Blackwell.

Luper, Steven. 2007. "Mortal Harm." *Philosophical Quarterly*, vol. 57, pp. 239-251.

Markie, Peter. 2005. "Nonidentity, Wrongful Conception, and Harmless Wrongs." *Ratio*, vol. 18, pp. 290-305.

McCarthy, J. J. Ed. 2001. *Climate Change 2001: Impacts, Adaptation, and Vulnerability*. Cambridge University Press.

- McMahan, Jeff. 2009. "Asymmetries in the Morality of Causing People to Exist." M.A. Roberts, D.T. Wasserman (eds.), *Harming Future Persons: Ethics, Genetics, and the Non-Identity Problem*, International Library of Ethics, Law, and the New Medicine.
- 2006. "Paradoxes of Abortion and Prenatal Injury." *Ethics*, vol. 116, pp. 625–655.
- Menzel, Christopher. 2008. "Actualism." In E. N. Zalta (ed.) *The Stanford Encyclopedia of Philosophy*, URL = <<http://plato.stanford.edu/entries/actualism/#ActualistResponses>>.
- Montoya et al. 2006. "Brain Imaging and Cognitive Dysfunctions in Huntington's Disease." *Journal of Psychiatry Neuroscience*, vol. 31, no. 1, pp. 21-29.
- Narveson, Jan. 1967. "Utilitarianism and New Generations." *Mind*, vol. 76, no. 301 January, pp. 62-72.
- National Academy of Sciences Committee on the Science of Climate Change. 2001. *Climate Change Science: An Analysis of Some Key Questions*, National Academy Press, Washington, DC.
- Norcross, Alastair. 1997. "Comparing Harms: Headaches and Human Lives." *Philosophy and Public Affairs*, vol. 26, no. 2 spring, pp. 135-167.
- 2005. "Harming in Context." *Philosophical Studies*, vol. 123, pp. 149–173.
- 2006. "The Scalar Approach to Utilitarianism." In Henry West (ed.), *The Blackwell Guide to Mill's Utilitarianism*.
- Nord, Erik, Daniels, Norman, Kamlet, Mark. 2009. "QALYs: Some Challenges." *Value in Health*, vol. 12, no. 1, pp. S10-S15.

Nozick, Robert. 1969. "Coercion." *Philosophy, Science, and Method: Essays in Honor of Ernest Nagel*. Eds. Morganbesser, Suppes, White. St. Martin's Press: New York.

----Nozick, Robert. 1974. *Anarchy, State, and Utopia*, New York: Basic Books.

Parfit, Derek. 1976. "On Doing the Best for Our Children." In Michael Bayles, ed., *Ethics and Population*, Cambridge, MA: Schenkman, pp. 100-02.

----1984. *Reasons and Persons*. Oxford University Press.

----1997. "Equality and Priority." *Ratio*, vol. 10, no. 3, pp. 202-21.

----2000. "Equality or Priority?" In *The Ideal of Equality*, ed. M. Clayton and A. Williams, New York: St. Martin's Press, pp. 81-125.

----2011. *On What Matters*. Oxford University Press.

Parker, Michael. 2007. "The Best Possible Child." *Journal of Medical Ethics*, vol. 33, pp. 279-283.

Parsons, Josh. 2003. "Why the Handicapped Child Case is Hard," *Philosophical Studies*, vol. 112, no. 2 January, pp. 147-162.

Persson, Ingmar. 2003. "Two Claims About Potential Human Beings." *Bioethics*, vol. 17 no. 5-6, pp. 503-516.

Peters, Philip G. 2009. "Implications of the Nonidentity Problem for State Regulation of Reproductive Liberty." In Roberts and Wasserman, eds., *Harming Future People*, pp. 17-31.

Portmore, Doug. 2008. "Dual-Ranking Act-Consequentialism." *Philosophical Studies*, vol. 138, pp. 409-427.

Rawls, John. 1971. *A Theory of Justice*. Harvard University Press.

Regan, Tom. 1983. *The Case for Animal Rights*. University of California Press.

Reiman, Jeffrey. 2007. "Being Fair to Future People: The Non-Identity Problem in the Original Position." *Philosophy and Public Affairs*, vol. 35, no. 1, pp. 69-92.

Roberts, M.A. 2002. "A New Way of Doing the Best That We Can: Person-Based Consequentialism and The Equality Problem." *Ethics*, vol. 117, pp. 315-350.

----2003. "Can it Ever Be Better Never to Have Existed At All? Person-Based Consequentialism and a New Repugnant Conclusion." *Journal of Applied Philosophy*, vol. 20, no. 2, pp. 159-185.

----2007. "The Non-Identity Fallacy: Harm, Probability, and Another Look at Parfit's Depletion Example." *Utilitas*, vol. 19, no. 3, pp. 267-311.

Savulescu, Julian. 2001. "Procreative Beneficence: Why We Should Select the Best Children." *Bioethics*, vol. 15, no. 5-6, pp. 413-426.

Savulescu, Julian. 2007. "In Defence of Procreative Beneficence." *Journal of Medical Ethics*, vol. 33, pp. 284-288.

Scheffler, Samuel. 1982. *The Rejection of Consequentialism*. Oxford.

- Shiffrin, Seana Valentine. 1999. "Wrongful Life, Procreative Responsibility, and the Significance of Harm." *Legal Theory*, vol. 5, pp. 117-148.
- Schwartz, Thomas. 1978. "Obligations To Posterity." in Richard Sikora and Brian Barry, eds., *Obligations to Future Generations*, Philadelphia: Temple University Press, pp. 3-13.
- Sikora, S.I. 1975. "Utilitarianism: The Classical Principle and the Average Principle." *The Canadian Journal of Philosophy*, vol. 5, no. 3, pp. 409-419.
- Silverstein, Harry. 2000. "The Evil of Death Revisited." In *Midwest Studies in Philosophy*, vol. 24, ed. French and Wettstein, pp. 116-34.
- Singer, Peter. 1979. *Practical Ethics*, Cambridge University Press.
- 1972. "Famine, Affluence, and Morality." *Philosophy and Public Affairs*, vol. 1, no. 1, pp. 229-243.
- 1978. "Anglin on the Obligation to Create Extra People." *The Canadian Journal of Philosophy*, vol. 8, no. 3, pp. 583-585.
- Smolkin, Doran. 1994. "The Non-Identity Problem and the Appeal to Future People's Rights." *Southern Journal of Philosophy*, vol. 32, no. 3, pp. 315-329.
- 1999. "Towards a Rights-Based Solution to the Non-Identity Problem." *Journal of Social Philosophy*, vol. 30, no. 1, pp. 194-208.
- Steinbock, Bonnie. 2009. "Wrongful Life and Procreative Decisions." In Roberts and Wasserman, eds., *Harming Future Persons*. pp. 155-78.

Stern, Nicholas. 2007. *The Economics of Climate Change: The Stern Review*, Cambridge University Press.

Sumner, L.W. 1987. *The Moral Foundations of Rights*, Oxford University Press.

The American Association of Petroleum Geologists. 2001. *Geological Perspectives of Global Climate Change*.

Temkin, L. 2000. "Equality, Priority, and the Levelling Down Objection." In *The Ideal of Equality*, ed. M. Clayton and A. Williams, New York: St. Martin's Press, pp. 126-61.

Thomson, J. 1971. "A Defense of Abortion." *Philosophy & Public Affairs*, Vol. 1, no. 1 Fall.

----1990. *The Realm of Rights*, Cambridge: Harvard University Press.

----2011. "More On the Metaphysics of Harm." *Philosophy and Phenomenological Research*, vol. 82, no. 2, pp. 436-458.

Tooley, Michael. 1985. *Abortion and Infanticide*. Oxford University Press.

Velleman, David. 2008. "Persons in Prospect." *Philosophy and Public Affairs*, vol. 36, no. 3, pp. 221-244.

Weinberg, Rivka. 2008. "Identifying and Dissolving the Non-Identity Problem." *Philosophical Studies*, vol. 137, pp. 3-18.

Woodward, James. 1986. "The Non-Identity Problem." *Ethics*, vol. 96, no. 4, pp. 804-831.

Wrigley, Anthony. 2011. "The Problem of Counterfactuals in Substituted Judgment Decision-Making." *Journal of Applied Philosophy*, vol. 28, no. 2 May, pp. 169-187.