

2013

Comparing Consequentialist Solutions to the Nonidentity Problem

Emily K. Ott

Claremont McKenna College

Recommended Citation

Ott, Emily K., "Comparing Consequentialist Solutions to the Nonidentity Problem" (2013). *CMC Senior Theses*. Paper 635.
http://scholarship.claremont.edu/cmc_theses/635

This Open Access Senior Thesis is brought to you by Scholarship@Claremont. It has been accepted for inclusion in this collection by an authorized administrator. For more information, please contact scholarship@cuc.claremont.edu.

CLAREMONT McKENNA COLLEGE
COMPARING CONSEQUENTIALIST SOLUTIONS TO
THE NONIDENTITY PROBLEM

SUBMITTED TO
PROFESSOR S. ANDREW SCHROEDER, Ph. D.

AND
DEAN GREGORY HESS, Ph. D.

BY
EMILY K. OTT

FOR
SENIOR THESIS
SPRING SEMESTER
APRIL 29, 2013

TABLE OF CONTENTS

Acknowledgments.....	1
Introduction.....	2
Chapter I: The Nonidentity Problem.....	5
Chapter II: Nonidentity Case Studies.....	12
<i>Same-Number Cases</i>	12
<i>Different-Number Cases</i>	14
<i>Overview of the Nonidentity Problem</i>	16
Chapter III: Proposed Consequentialist Solutions.....	19
<i>What Consequentialism Entails</i>	19
<i>The Totalist Solution</i>	21
<i>The Repugnant Conclusion</i>	23
<i>The Averagist Solution</i>	28
<i>The Mere Addition Paradox</i>	29
Chapter IV: Comparing Consequentialist Solutions.....	34
<i>Problems for Average Consequentialism</i>	34
<i>Modifying Totalism – Critical Level Principles</i>	38
Conclusion.....	41
Works Cited.....	43

ACKNOWLEDGEMENTS

First and foremost I would like to thank Professor Schroeder for his thoughtful, patient guidance and careful editing throughout the process of writing this thesis. I am also grateful for the knowledge and experience I have gained from the entire CMC Philosophy Department throughout the past four years. I would also like to thank Clare Riva and Sean Smith for their contributions to the final content and coherence of this final product.

Finally, I would like to thank my parents for their unwavering love and support, not only throughout the last four years, but for the entirety of my educational development. I truly appreciate it.

INTRODUCTION

The philosophical puzzle known as the nonidentity problem is a complex and profound issue which is interesting not only because of the attention it has received from philosophers and academics, but also because of its power to alter every living person's perspective towards his or her own existence. An overarching aim of any normative ethical theory is to explain why acts are right or wrong.¹ The most evident answer to the question of what makes an act wrong is that it causes harm. What the nonidentity problem asks us to do is to articulate just what makes an act wrong, by considering who it may harm. In many cases, we can definitively identify the object of our harmful acts – by throwing a rock at Joe, I harm Joe, and so my action is wrong. But, in other cases, the people whom our actions may harm have not been born yet. This fact calls into question the very structure of moral law – are immoral acts wrong because they harm someone like Joe, or can actions be wrong for reasons beyond what they can be expected to do to any existing person?² As this paper will show, an adequate moral theory must justify the fact that an action can be wrong even if there is no specific person who it would negatively affect – it must solve the nonidentity problem. Consider the following:

Two potential parents are considering conceiving a child. Prior to conception, either one (or both) of these two people have the option of taking a pill which will slightly increase the pleasure they experience during intercourse. However, taking the pill would also guarantee that the child they conceive would be born with mild dyslexia, a

¹ Normative ethics, as opposed to comparative ethics, involves defining how people should ideally act versus how they actually do. It seeks to find an ideal moral theory.

²Roberts, Melinda. "The Nonidentity Problem." Web.

treatable but incurable handicap. Aside from this handicap, however, their child's life would be good. Consider three scenarios in which the parents vary the timing and manner of conception; one in which the couple pauses and does not take the pill, one in which the couple pauses and takes the pill, and one in which the couple does not pause. If neither of them takes this "pleasure pill," we can assume that their child will be born without any other defects. If the couple does not pause to take the pill, a perfectly healthy individual will be born from a single distinct sperm, out of billions of possible ones. If the couple does pause for several seconds, they would virtually guarantee that a different sperm would reach the egg, because new sperm are constantly dying and being regenerated at a rate of about 1,500 per second.³ Therefore, any pause would eliminate the possibility that *the same* perfectly healthy individual would be born as would have if they had not hesitated. If the couple does pause for several seconds to take the pill, yet another different sperm would reach the egg. In each of these scenarios, a different sperm would reach the egg, and none of these three scenarios can occur simultaneously. And, in each of these scenarios, only one distinct child can possibly be conceived.⁴ Since only one child is conceived in each case, only one child's existence can be influenced by the decision of his or her parents to either take the pill or refrain from doing so, so *no one* would be rendered worse-off if one or both potential parents took the pill.

This conclusion hinges on our inclination to say that if a handicapped child is conceived, his existence with a mild handicap would still be better, all things considered, than his never having existed at all. Consider how we would feel if, after taking the pill,

³ Dell'Amore, Christine. "How a Man Produces 1,500 Sperm a Second." *National Geographic*. Web.

⁴ Even in incidences of genetically identical siblings, only one egg is fertilized, which splits into separate embryos 1-9 days after conception. In incidences of fraternal siblings, two genetically distinct eggs are fertilized by two genetically distinct sperm.

the couple had decided to abort their child due to his handicap rather than give birth to him. Even if you deliberately cause your child to be mildly dyslexic, it is better for him to be handicapped and alive than to never have existed at all, which would have been the case had no one taken the pleasure pill. However, we would still consider it wrong of these parents to deliberately handicap *this* child.⁵ This example poses the nonidentity problem. In this thesis, I will argue that the nonidentity problem poses a substantial obstacle to moral theories which attempt to explain why we have a responsibility to prevent causing harms to people who do not yet exist, but may exist in the future. I will then discuss two versions of consequentialism, totalism and averagism, which have attempted to circumvent the nonidentity problem. After proving that these two theories each create an additional problem, I will argue for a modification of one, which has the potential to subvert the most substantial obstacle faced by its simpler form.

⁵ Kavka, Gregory S. "The Paradox of Future Individuals." 98.

CHAPTER I: THE NONIDENTITY PROBLEM

As the pleasure pill example above alludes to, is important to understand a few empirical facts about how new people, with new identities, come to exist. The first of these is that every person begins developing from a particular, distinct pair of cells – exactly one egg and exactly one sperm. For our purposes, we can assume that the single female egg in question remains the same distinct cell (it contains the same set of genes) for a time interval of roughly twenty-eight days. However, every time a child is conceived, there is the possibility that any one (and only one) of millions of distinct male sperm cells will reach the egg and ultimately fertilize it. If one of these sperm, carrying a *slightly* different set of genes from the one next to it, reaches the egg first, the child conceived will have a different set of genes than he would have had if he had begun life as a different pair of gametes. The chance that a genetically distinct sperm cell will reach a given ovum changes literally second-by-second. So, if two people wait even seconds longer to conceive a child than they had intended to, or conversely, if they conceive earlier, a distinct person, with a unique genetic makeup, will come into existence. The probability that a distinct individual will be conceived increases proportionately to the amount of time in consideration; a couple who waits one year after first considering conceiving a child will have passed up the opportunity of having conceived billions of distinct individuals.⁶

Barring some apocalyptic turn of events, it is a second fact we can take for granted that there will be people living in the future who do not yet exist today. For a

⁶ Derek Parfit, *Reasons and Persons*. 351-52.

multitude of reasons, it is within the power of current generations to significantly affect the lives of these people who do not yet exist – these *future people*. This may be done by altering either (a) the number of future people, or by altering (b) the identities of future people.⁷ Examples of how this may be done are easily enough imagined. Simple family planning methods (or lack thereof) may either restrict or expand the number of distinct individuals created. Or, as opposed to having a different number of children altogether, parents could vary the times at which they conceive, which would, as detailed above, lead to the creation of distinct individuals. Or, parents may choose to alter the manner in which their child is conceived, perhaps by using in vitro fertilization.⁸ On any given day, a potential mother returning home from work to a potential father one hour later than usual will result in the conception of a completely different child. The identities of future people are under the control of one or both of their parents, and almost always others, too (the boss who keeps a mother late at work, legislators who enact policies which affect the course of parents' daily lives, or inventors who revolutionize transportation methods, for example).

Taken together, these two sets of facts contribute to a phenomenon which Gregory Kavka has termed “the precariousness of human origin,” meaning that the particular people who will exist in subsequent generations depend significantly on the conditions under which their parents procreate. Even the slightest alteration of these conditions ensures the creation of different people coming into existence.⁹ These details will be important to keep in mind when their moral implications are considered later in

⁷ Ibid, 355-356.

⁸ Roberts, Melinda. “The Nonidentity Problem.” Web.

⁹ Kavka, Gregory S. “The Paradox of Future Individuals.” 93.

this paper. From now on, the entities of the sperm and egg cells which have not yet joined together, as well as all future distinct sperm and egg cells which *will* ever join together, and the individuals who are created in the process, will be referred to as *future persons*.

It is important to note at this point that our ability to refer to different future persons hinges on a specific definition of personhood – namely, one in which genetic structure plays a significant role. Some theories of personal identity focus on criteria which are not strictly biological, but may include factors such as formative life experiences and memories, or the continuity of goals and psychological states. For the purpose of this paper, a definition of personal identity which includes and necessitates “the sameness of genetic structure,” or the possession of a distinct set of genes which originated from a particular sperm and egg, will be used.¹⁰ This is because even psychologically-based theories of personal identity are, at the very least, initially informed by our biology – consider how one’s perceptions and formative life experiences would differ if one was born with a handicap, or born incredibly attractive. For this reason, psychologically-based theories will inevitably lead to the same problems as genetically-based theories.

Identity is contingent upon sameness of genetic structure. Likewise, differences in genetic structure necessitate differences in identity. Derek Parfit, one of the principal commentators on the nonidentity problem, highlights this notion when observing in his work, *Reasons and Persons*, that the woman who wonders who we would have been if our parents had married other people “ignores the answer: ‘No one.’”¹¹

¹⁰ Ibid, 93-94.

¹¹ Derek Parfit, *Reasons and Persons*. 351.

There are seemingly infinite numbers of ways in which this sameness of genetic structure is affected by the events surrounding conception, each of which is influenced by one, or both, of the providers of the sperm and egg. As we saw already in the pleasure pill case, the identity of the child is in many ways determined by actions which occurred before he or she begins the first stage of development. This suggests that parents, or anyone involved in the conception of a child (at any point prior), will have a degree of responsibility for how that child turns out. However, this responsibility is not always of a moral nature. Only in cases where people have the option of having different children, particularly different inasmuch as they are worse-off than they could have been, such as in the pleasure pill case, is moral responsibility a consideration.

This point translates to larger-scale cases in which our actions can affect future generations. Moral controversy arises when we consider how the actions of people alive in the present will either positively or negatively affect future people. Nearly every single one of our choices will have some ripple effect on at least one future person, whether this person is born in nine months or in 900 years. However, not every one of these choices is morally relevant. Only acts *which create predictable harms or benefits*, and are chosen *when better alternatives are available*, as they were in the pleasure pill case, are considered for the purpose of the nonidentity problem. This is because, in these instances, agents are fully aware that they will predictably harm future people, and yet they chose to do so regardless of the fact that they could have forgone doing so without harming themselves. Only actions which predictably lead to one individual being worse-off than another, either to a different degree or in a different way, are morally relevant.¹² This is

¹² Ibid. 356-357.

because in these cases, an agent's informed decisions lead directly to a negative consequence for some future person which would not have affected any future people if a different decision, with better consequences, had been made. The nonidentity problem considers only the moral status of acts which have the predictable potential to harm any possible future people.¹³

Against the backdrop of this information, it is now possible to understand why the nonidentity problem is termed as such. The nonidentity problem arises when we consider two commonly-held, but conflicting, intuitions about our actions and their effects on future persons:

1. Some acts whose effects are restricted to persons who do not yet but will exist are morally impermissible (In the pleasure pill example, parents knowingly handicapping their children).
- and*
2. If such an act is wrong, it is wrong because it will affect *some* person in an adverse way; it is wrong because it will (predictably) *harm* one or more future persons, or make things *worse for* one or more future persons than they might have been (No perfectly healthy child is ever harmed because he or she is never brought into existence, in the pleasure pill scenario).¹⁴

The nonidentity problem holds that both of these intuitions cannot be true at the same time. The first intuition means, in simplest terms, that it would be morally wrong to cause harms to people who do not yet exist. Causing harms to future people would cause them to live in a state in which they would be *worse-off* than they would have been if we had not acted so as to cause harm. The term *worse-off* may be used to describe any one of two people who has a lower quality of life, in terms of overall happiness and standard of living, than another. Exactly what constitutes a harm is a more subtle point, and one which will feature more extensively later in this paper. However, for now, it suffices to

¹³Roberts, Melinda. "The Nonidentity Problem." Web.

¹⁴Derek Parfit, *Reasons and Persons*, 361-363.

rely on Parfit's simplification of the first intuition, which claims that if "in either of two possible outcomes the same number of people would ever live, it will be worse if those who live are worse-off, or have a lower quality of life, than those who would have lived."¹⁵

The second intuition, referred to alternately as the Person-Affecting Intuition, View, or Principle, claims that bad acts – that is, morally wrong acts, or acts which cause harm – must be bad *for* someone; they must cause harm *to* a specific person.¹⁶ The Person-Affecting Principle claims that as long as an act does not negatively affect an existing person's wellbeing, this act cannot be considered wrong. This is true even if this act was one of two alternatives, and the chosen alternative has many more negative consequences for future people than the other. Consider a small nation, which, faced with a growing population, chooses to implement policies which will deplete its natural resources, as opposed to choosing policies which favor conservation. The choice of depletion benefits people alive in the present, by allowing them to forgo the strains of rationing food and material resources. It will also lead to a complete exhaustion the country's environmental resources within the span of 200 years. While the depletion option harms no specific future people, it seems as though it is worse, even reprehensible, to choose an act which will predictably create negative effects, when the country's policymakers had the alternative of choosing a different act with positive, or at least better consequences. However, it is evident that the agents' choice has not *harmed* any

¹⁵ Ibid. 369.

¹⁶ Ibid. 363.

specific person at all, because the people who exist in the future owe their very existence to the implementation of the depletion policy.¹⁷

This leaves us with the result that at least some *bad* acts are *bad for* no one at all. In other words, the two intuitions comprising the nonidentity problem are mutually exclusive. If we accept the first, we seem forced, in other words, to reject the second.¹⁸ It seems as though we cannot logically hold both the first intuition *and* the second, the Person-Affecting Principle, to be true at the same time. To better understand the nonidentity problem, it is useful to consult a few examples which highlight how our clashing intuitions about harm are, in fact, problematic.

¹⁷ Roberts, Melinda. "The Nonidentity Problem." Web.

¹⁸ Ibid. Web.

CHAPTER II: NONIDENTITY CASE STUDIES

The nonidentity problem applies to two different types of cases in which our choices affect which future people exist. One type concerns hypothetical scenarios involving the same number of people, but who have different identities than they might otherwise have had, and the second type concerns different numbers of people. The following discussion of these examples clarifies how the nonidentity problem applies to each one.

SAME-NUMBER CASES

The first case which helps to illuminate the conflicting intuitions which comprise the nonidentity problem involves a fourteen-year-old-girl who finds herself pregnant. While she may have compelling moral, emotional, or religious reasons to keep the child, she also has the alternative of not having the baby. Because of her age, the fourteen-year-old is missing several qualities which would enable her to be an ideal parent to her child – for example, she lacks a complete formal education, as well as the maturity to be hired to do a job which would enable her to provide for herself and another person. She would not be able to drive herself to a grocery store, or to a doctor's office in an emergency. Even if she waits to move out of her parents' home to live independently (when she is eighteen and the child is four, at the earliest), her financial state will likely restrict her to living in a less-than-ideal school district. Fully aware of these factors, the girl decides to have her child (Child A) anyway. Had she not done so, *this* child never would have

existed – he owes his existence, however disadvantaged, to his mother’s decision.¹⁹ However, we may still be tempted to say that the fourteen-year-old harmed *this* child by giving him a bad start in life.

Had the fourteen-year-old girl chosen to terminate her current pregnancy and postpone motherhood until she was better-prepared to support a child, there is a chance the child’s father might have been the same as the father in the first pregnancy. The girl would still be having *just one* child. However, for the biological reasons noted in Chapter I, this child (Child B) would not have the same *identity* as the child she could have had when she was fourteen. But, this child would have been *better-off* than the child she would have otherwise brought into existence – “in different outcomes, different people would be born.”²⁰

In this type of same-number case, where the same number of people will come into existence, but some will clearly be better-off than others, Parfit suggests that we appeal to the *Same-Number Quality Claim*, which states that if the same number of people will live in either of two potential scenarios, it would be worse if the people who live in one are worse-off, or have a lower quality of life, than those who live in another.²¹

This claim might seem to suggest that it would be better if only Child B had been brought into existence, and worse if Child A had; or, conversely, that it would have been better if Child A had not existed. However, the fact that it would have been better if Child B had been born does not imply the harsh conclusion that Child A should never have been born *at all*. As long as Child A still has *a life worth living*, that fact that he was born

¹⁹ Derek Parfit, *Reasons and Persons*. 358-359.

²⁰ *Ibid.* 358-359.

²¹ *Ibid.* 360.

is not morally objectionable. However, the fact that Child A's existence is not morally objectionably does not entail that it was *right* of the fourteen-year-old girl to give birth to *him*.²² The Same-Number Quality Claim does not solve all instances of the nonidentity problem, but does help to deal with same-number cases.²³ It does not, however, address a second scenario: one in which a different number of people, with necessarily different identities, come into existence.

DIFFERENT-NUMBER CASES

A variety of social or economic policies, especially those which involve the depletion or conservation of resources, population control, or environmental practices, have effects on future generations which help to illustrate the nonidentity problem on a grander scale.²⁴

One notable example of this involves a country's choice between two energy programs, either of which would satisfy its energy needs for the next several centuries. The first of these alternatives would involve a sizable investment in solar energy systems; the second would involve a reliance on cheaper nuclear fission power plants. The latter option would be cheaper by ten percent than the first, in terms of cost per unit of electrical output. The solar option has no detrimental effects on the environment, now or in the foreseeable future. The second option would produce radioactive nuclear waste, which may be disposed of in containers which would prevent it from escaping for at least

²² For the purposes of this paper, "a life worth living" will describe any life in which the person living it does not feel that it would be better for them to have never existed. It is possible to imagine cases in which a life may objectively be considered not worth living, such as one described by Parfit, of a child who is "so multiply diseased that his life will be worse than nothing. He will never develop, will live for only a few years, and will suffer pain that cannot be wholly relieved."

²³ Derek Parfit, *Reasons and Persons*. 361.

²⁴ *Ibid.* 361-362.

several decades. However, it is guaranteed that these containers will inevitably break down, causing thousands of future people to die prematurely as a result of radiation poisoning. But, these future people will not be so negatively affected by radiation poisoning during the time which they are alive that their lives will have been considered overall not worth living.

Because building the nuclear plants would alter the manner of conception of thousands of people, the thousands of individuals who would be harmed by radioactive waste would never have existed had the solar energy option been implemented instead. And, because we assume that the radiation victims live lives worth living, we are still justified in selecting the nuclear option by appealing to the Person-Affecting Principle discussed earlier; no persons exist to whom harm can be assigned.²⁵ This is because, as we saw in the depletion example, the people who live to develop radiation poisoning owe their very existence to the fact that the nuclear option was implemented. However, before considering the nonidentity problem, many of us were still nonetheless concerned about the welfare of future generations. Because the cost of the solar program is not prohibitively expensive, we intuitively believe that we should choose this one, in order to avoid responsibility for the deaths of thousands of future people. When we realize that the individuals comprising these future generations will not be the same individuals who would have existed under programs such as the nuclear energy option, this moral concern for their wellbeing does not diminish. The idea that we do not morally differentiate

²⁵ Ibid. 371-72.

between the identities of the people who will exist in the future, whether they come into existence under the nuclear plan or the solar one, is called the *No-Difference View*.²⁶

This example again shows that, in assigning blame for harms to future generations, we must consider the *predictable* effects of our actions. We realize that, if we choose the policy which guarantees radiation poisoning, our choice will not be worse *for* those future people whom it later harms. We realize that, if we choose the policy which guarantees radiation poisoning, this will cause harm to many people in the distant future. But we also know that, if we had chosen the safer policy, the people who are harmed would never have been born. (The higher taxes their parents would have paid to fund the solar program would have prevented them from being able to support children, for example.) And, since these people were in fact born, and live lives worth living, we know that our choice of the more dangerous policy would not be worse for them. However, the nonidentity problem calls on us to explain why we have a moral reason not to chose the more harmful policy.²⁷

OVERVIEW OF THE NONIDENTITY PROBLEM

If the couple in the pleasure pill example chooses to handicap their child, if the fourteen-year old girl gives her child a bad start in life, if the small nation chooses policies favoring depletion, or some country implements nuclear energy, the quality of life for future people who exist as a result of these choices will undeniably be lowered. However, the Person-Affecting Principle does not imply that these choices will have bad

²⁶ Ibid. 367.

²⁷ Ibid. 378.

effects; because of the facts about identity, these given choices will be bad for no one. Would the effects of these decisions be worse if they *were* worse for particular people?

If we appeal to the Person-Affecting Principle rather than the Same-Number Quality Claim, our answer would be yes. The Same-Number Quality Claim describes the effects which we believe to be bad. But, since we also hold the No-Difference view, we answer no to this question. Furthermore, we believe that it makes no difference whether these effects are also bad according to the Person-Affecting Principle. The Person-Affecting Principle draws a moral distinction between persons where we intuitively believe no distinctions should be drawn. According to the Same-Number Quality Claim, our choice has a bad effect. That result leaves us with this problem: where an exclusively future-directed act is wrong, what makes it wrong, if not that it is *bad for* someone or another? What theory explains why it is wrong? We need a new moral theory to explain this to us. The nonidentity problem shows that this theory will not be able to consider the Person-Affecting Principle, since the people to whom we are attempting to prevent harm will only exist if we implement harmful policies.²⁸ So, we need a new principle which will mandate that we do not chose harmful policies, while not relying on the wellbeing of future people as a justification.

The form this new theory will take is not obvious, since the nonidentity problem is not the only problem that the formation of this theory will encounter. Many theories which avoid the Person-Affecting Principle attempt to explain how an act which is *bad for* no one can be wrong, and several philosophers have proposed these types of theories. The remainder of this paper will deal with two such proposed solutions, namely, total and

²⁸ Ibid. 378.

average consequentialism. However, these theories run up against at least as many problems as they solve; respectively, they face the Repugnant Conclusion, or the Mere Addition Paradox.²⁹

To summarize what I have argued, I will briefly reiterate what I have claimed about the nonidentity problem thus far. Choosing to have a child as a teenager or choosing a policy of depletion policy does not harm any person, since the people who exist as a result of these choices in fact owe their very existence to them. Since these acts do not harm anybody, they seem like they cannot be morally objectionable. If we initially believed that these choices were objectionable this was only because we mistakenly failed to take the nonidentity problem into account. Nevertheless, if giving birth to Child A or causing future generations to live in a barely inhabitable wasteland still seems intuitively wrong, as it should on a plausible moral theory, a solution to the non-identity problem must be found.

²⁹Roberts, Melinda. "The Nonidentity Problem." Web.

CHAPTER III: PROPOSED CONSEQUENTIALIST SOLUTIONS

While dozens of solutions to the nonidentity problem have been proposed, none of them have managed to rectify the two conflicting intuitions without significant flaws. Consequentialist theories provide us with a means of explaining how an act can be wrong, while making things worse for no one. Thus, while consequentialism rejects the Person-Affecting Principle, it achieves the same result as the theory which Parfit seeks to discover in *Reasons and Persons*.

WHAT CONSEQUENTIALISM ENTAILS

Consequentialism is, at its core, a school of ethics which claims that the moral rightness or wrongness of an action depends only on the consequences which it can be said to produce. Many plausible ethical theories exist – why examine the nonidentity problem from the view of its consequentialist respondents?

In *Reasons and Persons*, Parfit aims to construct a complete moral theory which would determine which actions are right as a result of their consequences – specifically, the consequences which they bring about for future people. Parfit’s moral theory aims to maximize the good that can be brought about for future people, thus, the framework of the nonidentity problem is already built upon consequentialist foundations. Perhaps the best answer to “why consequentialism?” is articulated by ethicist Melinda Roberts, who argues that “the consequentialist notion that we ought to bring about the most good that we can – that to do this is to do the *best* we can – seems a truly important, even profound,

insight into moral law. I, at least, find approaches that abandon maximization difficult to understand...³⁰

Consequentialist proposals reject the Person-Affecting Principle in seeking to justify how acts can be wrong, because they mandate that we maximize overall wellbeing. We can do this by increasing the wellbeing of an existing person or population, or by increasing the number of well-off people who exist. If it is not possible to increase the average wellbeing of some or all existing people, consequentialism holds that we *must* increase wellbeing by increasing the total number of well-off people who exist.

Two forms of consequentialism, average consequentialism and total consequentialism, have provided us with ways to avoid two problems faced by attempts to solve the nonidentity problem. Average consequentialism claims that, if other things are equal, it is better if people's lives go better on average, and maintains that acts are judged to be good based on the *average* net good which they produce. Total consequentialism claims that we should maximize aggregate wellbeing, and maintains that moral rightness depends only on the *total* net good of an act's consequences (as opposed to the average net good it produces per person).³¹

Neither of these views is sufficient to address each of two significant problems – namely, the Repugnant Conclusion and the Mere Addition Paradox – at the same time. Average consequentialism and total consequentialism are alike in that they both result in solutions to the nonidentity problem in same-number, different identity cases (such as

³⁰ Roberts, Melinda A. *Child versus Childmaker: Future Persons and Present Duties in Ethics and the Law*. 7.

³¹ Aggregate wellbeing entails maximizing the total amount of welfare across a society, as opposed to maximizing welfare for only certain groups or not concentrating on maximization as an end.

that of the fourteen-year-old girl). However, in different-number, different identity cases (such as the depletion case), they yield significantly different results.³² In what follows, I will explain the ways in which both average and total consequentialist theories propose to solve the nonidentity problem, and why they run up against the Repugnant Conclusion and the Mere Addition Paradox.

THE TOTALIST SOLUTION

Total consequentialism is a form of consequentialism which maintains that we should maximize *aggregate* wellbeing.³³ Also known as totalism, it is the form of consequentialism that claims that moral rightness depends only on the *total* net good of an act's consequences (as opposed to the average net good per person). When applied in the context of the nonidentity cases we have considered thus far, totalism yields the conclusion that if other things are equal, it is better if there is a greater total sum of whatever makes life worth living. This view implies that the total quantity of welfare experienced by a given population is all that matters. Total quantity is calculated by summing the number of units of welfare possessed by each person who lives under a given set of conditions.³⁴ Under totalism, the best outcome is one in which the product of

³² Singer, Peter. *Practical Ethics*. 122–25.

³³ Aggregate wellbeing entails maximizing the total amount of welfare across a society, as opposed to maximizing welfare for only certain groups or not concentrating on maximization as an end.

³⁴ In order to judge whether these people's lives would be worth living, or barely worth living, or not worth living at all, we must rely on some type of metric. Measuring the total number of "utils," or units of wellbeing, possessed by a given population is one method commonly discussed by philosophers as a method of determining how a population may measure welfare. Utilitarianism is an approach to normative ethics, of which consequentialism is a variation, which claims that the morally right action is the action that produces the most good. On the utilitarian view, one ought to maximize the overall good, as measured in utils, or units of good, also called "happiness" or "pleasure."

this calculation is as high as possible – one where the total amount of welfare is maximized.

Totalism solves the nonidentity problem by explaining how acts may be deemed morally wrong without requiring an appeal to the Person-Affecting Principle. Total consequentialism's focus on aggregate wellbeing implies that we do not need to focus on the question of *who* our actions may harm in the future, but instead direct our concern towards the question of how to create a future in which people are as well-off as possible, given the options available to us today.

Suppose that a survey was given to the 4,380,415 citizens who were residents of Kentucky in 2012, asking them to rate their level of wellbeing on a scale from 0 to 100.³⁵ The results of this survey indicate that the average Kentuckian has a level of wellbeing of 86. By summing the individual self-reported levels of wellbeing of the citizens of Kentucky, we find out that the Bluegrass State has a total wellbeing of 375,715,690 units.³⁶ Now, suppose that 100,000 Kentuckian babies are born before December 31, 2013, and that no one in Kentucky dies this year.³⁷ Kentucky's total amount of wellbeing has just increased to 384,315,690 units! Totalism claims that Kentucky in 2013 is morally better than Kentucky in 2012, due the increase in overall wellbeing experienced by its population. Adding new members to the population has directly led to an increase in total wellbeing. However, increasing the number of people in a given population is not always a recommendable way to raise its total level of welfare, as the following section suggests.

³⁵“Kentucky QuickFacts.” *United States Census Bureau*. Web.

³⁶ The same result would be yielded if, instead of averaging the amount of wellbeing experienced by each citizen, we simply summed the results.

³⁷ For this calculation I will assume that each of the newborn babies experiences the average reported level of welfare in Kentucky.

THE REPUGNANT CONCLUSION

The Repugnant Conclusion arises when we attempt to figure out how many people should be living in a given place at a given time, and when there would be too many people living. As this is a paper in philosophy, the answer to this question will not take into account hard numbers or economic data regarding the population level at which the earth will be unable to sustain additional life. Rather, it will focus on determining the point at which continuing to increase the earth's population size will detrimentally affect the welfare of future generations. The answer to this question will be determined by whether or not we believe it is in the interest of future people to be caused to exist; in other words, by answering the question of whether existing people have a moral obligation to produce future people, when doing so may seem to cause them harm.

Consider the following example. A couple is in the process of deciding whether or not to have a child. The couple feels that having a child would confer several benefits on them: they would experience the rewards of loving and nurturing their child and watching him or her succeed, they would fulfill a biological desire, and they would have someone to take care of them in their old age. However, the decision to conceive may also come with a set of drawbacks: raising a family would take time away from their careers, they may need to make significant financial adjustments to be able to support a child, or their relationship may not respond well to a major lifestyle change. If they settle in favor of conceiving, they will be able to provide this child with the conditions necessary to live a life worth living. After much reflection, the couple decides that having the child would be neither better nor worse *for them*. But would it be better or worse for their child to be

caused to exist? In other words, is there a moral obligation for these parents to conceive a future person who would have a life worth living?

Intuitions differ on this point. Some think it would be better if more people live – after all, if every couple decided against having children, what would become of the human race? Others think it makes no difference whether they bring additional people into existence – if *I* do not have a child, someone else surely will.³⁸ This example is small in scope, and might seem inconsequential at first glance. However, when we consider the effects of population growth in terms of not one couple, but of millions of couples, the problematic nature of our moral obligation either to procreate or to refrain from doing so becomes more evident.

Consider now a country with a population of five million people. This population has a finite number of natural resources, and while it currently consumes them at a sustainable level, increased rates of population growth will inevitably pose the threat of depletion. This depletion would be so severe that the lack of access to building materials, medical supplies, and even food and water would render the quality of life of the population very poor – a fact of which this country's residents are aware. Of the country's five million citizens, two million are people such as those described in the above scenario – half of them men and half of them women, all considering conceiving one child (per couple). Because of their country's small size and limited resource supply, these people must take into account not only their personal motivations and justifications, but also how their decisions will affect other existing and future people. If these couples decide to have children, the country's workforce will grow, among other benefits, but

³⁸ Derek Parfit, *Reasons and Persons*. 381.

they must also consider the negative effects of depletion. It will be helpful to assign a metric to this generation in order to quantify its initial level of wellbeing – I will henceforth use a measure of 100 units of welfare.

In one scenario, P_1 , all of the couples decide in favor of having children, adding one million new people to the country's existing population. In an alternate case, Q_1 , half of the couples decide to forgo having children, while the other half bring five hundred thousand new people into existence. Observing the attitudes of previous generations towards the positive and negative outcomes associated with having children, subsequent generations follow in their parent's footsteps (assume that the patterns stipulated in P_1 and Q_1 repeat for five generations). While the Q_2 population has experienced a sustainable growth rate and their lives are just as worth living as the Q_1 generation's lives were, the P_2 population now lives lives which are less worth living than the P_1 population's lives were. While Q_2 retains its welfare level of 100, P_2 's population experiences a level of 60.

Under Q_1 , fewer future people exist after these five generations have passed than otherwise would have if P_1 had occurred instead. Under P_1 , many more people exist after five generations than would have existed if Q_1 had been favored. These additional people owe their existence to the decision made by the original P_1 group to conceive children in every case (even though they were fully aware of the eventual consequences of overpopulation). However, these additional people live lives which are less worth living (by slightly more than half as much as they could have been otherwise). Now, we may ask the same question as in the single-couple case above – do we have a moral obligation to bring future people into existence? What about when we include the caveat that future

people will have lives worth living, but to a lesser degree than they could have otherwise been? It still does not seem plausible to *deny* that we should bring these people into existence – after all, their lives will still be worth living on the whole.

Could the fact that people in P_2 are worse-off be morally outweighed by the fact that there are more people living? We have accepted that, all things equal, it would be *worse* if previous generations ensure for future generations a lower quantity of whatever makes life worth living than they could have, given different alternatives. Conversely, we should agree that if other things are equal, the *best* outcome is the one in which there would be the greatest quantity of whatever makes life worth living. Parfit calls this claim the *Impersonal Total Principle*.³⁹

If we believe that people have no moral reason to bring future happy children into existence, then we must also believe the converse of this - that if people are worse off, this cannot be morally outweighed by increasing the number of happy people. These two beliefs may be summarized by the claim that “If other things are equal, the best outcome is the one in which there is the greatest net sum of happiness per life lived.”⁴⁰

If we believe that people have a moral reason to bring happy future people into existence, it is implied we believe that it is intrinsically valuable to have an extra life lived that is worth living. It follows from this that we also believe that a sufficient increase in the number of lives worth living would outweigh a certain loss of quality of life. Both of these claims deal with the balance between the quality and quantity of lives lived. To condense these beliefs, Parfit uses the phrase, “the *amount* of whatever makes

³⁹ Ibid. 387.

⁴⁰ Parfit calls this *The Impersonal Average Principle*, p. 386

life worth living.”⁴¹ In the context of the P_2 and Q_2 generations, we can now claim that compared to the Q_2 people, each of the P_2 people has less of “whatever makes life worth living.” But, we can assume that each life in P_2 is still more than half as worth living as each life in Q_2 . Since there are about twice as many P_2 people, this population collectively has more of what makes life worth living.

If we accept the Total Principle, we are forced to accept the conclusion that lives lived at P_2 are better in a moral sense than lives lived at Q_2 . However, this is clearly untrue – P_2 lives are in at least one way worse than lives in Q_2 , as they are valued at 60 as opposed to 100 – and they are not better in any way.⁴²

Imagine now that lives lived one hundred generations after P_1 (call this P_3) are *barely* worth living at all. Depletion has progressed to the point where the lack of building materials, medical supplies, food, and water has made P_3 's lives all but unbearable – P_3 has a welfare level of 1. However, the number of people living in P_3 is massive enough that these people still have a greater amount of what makes life worth living than people in Q_2 . In this case, the Total Principle implies that P_3 is better than Q_2 . Herein is where the Repugnant Conclusion lies. The Repugnant Conclusion states that:

“For any population of at least ten billion people (Q_2), all with a very high quality of life, there must be some larger imaginable population (P_3) whose existence, if other things are equal, would be better, even though its members have lives that are barely worth living.”⁴³

⁴¹ Derek Parfit, *Reasons and Persons*. 419. Emphasis added.

⁴² Ibid. 419.

⁴³ Ibid.. 388. Parenthetical notes added.

This is the first issue that any cogent response to the nonidentity problem must avoid. Total consequentialism does not do this, because it inherently implies the Repugnant Conclusion in dictating that we maximize the total amount of wellbeing at all costs.

Parfit believes, and I believe most people would agree, that the Repugnant Conclusion is intrinsically repugnant – the addition of extra lives (the extra people in P_3) does not improve a situation.⁴⁴ As Parfit writes, “If these [extra] lives are worth living, they have personal value. But the fact that such lives are lived does not make the outcome better.”⁴⁵ Average consequentialism provides an alternative solution to the nonidentity problem which does not fall prey to the Repugnant Conclusion.

THE AVERAGIST SOLUTION

Perhaps one observation that can be made about the Repugnant Conclusion is that it involves huge numbers of people, when in fact it can be difficult to conceptualize or appropriately express empathy for such a large population so far temporally removed from us. Average consequentialism provides examples which illustrate how thinking in smaller terms may help provide a solution to the nonidentity problem.

Average consequentialism is the form of consequentialism which claims that if other things are equal, it is better if people's lives go, on average, better. As its name implies, averagism maintains that acts are judged to be good based on the *average* net good which they produce. This view implies that only *quality* matters, and that the average wellbeing *per life* in a population should be maximized.

⁴⁴ Temkin, Larry S. “Intransitivity and the Mere Addition Paradox.” 141.

⁴⁵ Derek Parfit, *Reasons and Persons*. 381.

Averagism also solves the nonidentity problem by explaining how acts may be deemed morally wrong without requiring an appeal to the Person-Affecting Principle. Average consequentialism's focus on maximizing the average level of wellbeing implies that regardless of *who* our actions may harm in the future, the act that maximizes average wellbeing is still morally right. Average consequentialism also avoids the trap of the Repugnant Conclusion, since it states that the *average level* of welfare will clearly be higher in a small population where every person has a high level of wellbeing than in a population where each person has a life barely worth living, although it is exponentially larger. Averagism implies that no decrease in a population's average level of wellbeing can be compensated for, in a moral sense, by an increase in its total wellbeing. However, like totalism, averagism faces its own set of problems in endeavoring to solve the nonidentity problem.

THE MERE ADDITION PARADOX

Carefully consider the following diagram depicting three possible worlds, where the length of the rectangles represents the population's size, and the height represents the population's overall utility:

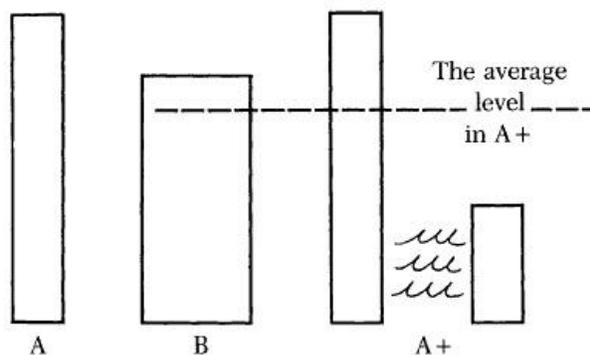


DIAGRAM 2

46

Imagine that Possible World A represents a world in which only the population of Spain in 1491 exists. A+ represents a possible world consisting of two distinct populations with no knowledge of each other's existence. The left rectangle in A+ represents the population of Spain in 1491, while the right rectangle represents only the indigenous population of North America in the same year. Possible World B represents a version of Spain in 1491 which is exactly twice as populous as A, and its residents have four-fifths the aggregate amount of happiness as A's do.

In A+ there is natural inequality – the extra people on the right are less than half as well-off as the other group, through no fault of their own. They are unaware of the fact that they are in a worse-off group than those on the left. There is no social injustice affecting the better-off group which could force them to surrender some of their happiness to the worse-off group. There is also no social injustice affecting the worse-off group; they are not worse-off because the group on the left is repressing or disadvantaging them in any way. If social injustice was present in these diagrams, adjustments would be made in terms of the moral status of Possible World A+.

⁴⁶ Derek Parfit, *Reasons and Persons*. 420.

The only difference between A and A+ is that in A+ a group has merely been added. Every person in this group has a life worth living and affects no one else. There is nothing about A+ which is worse than A. In Possible World B, more people exist, and have a higher average level of utility than all of the people in A+. So, we can assume that there is both on average and in total a higher level of utility, implying that B is better than A+. It seems, *prima facie*, that A+ is not a worse possible world than A. This is because the Spanish population in A is not worse-off in A+, while the indigenous people who exist in A+ are better off in A+ compared to A. In other words, while no one who exists in A *and* in A+ is worse-off in the latter possible world than in the former, there are *more* people living in A+, thus bringing up the total amount of wellbeing.⁴⁷ Additionally, it seems that B is better than A+, because in B, there are not only more people living, but each of them has a higher level of wellbeing than do the inhabitants of A+. These three comparisons together lead to the conclusion that Possible World B is better than Possible World A. Because the following three claims cannot plausibly be consistent, we are faced with the Mere Addition Paradox: A+ is no worse than A; B is better than A+; and B is worse than A.⁴⁸

Let me now show how this is a problem for averagist responses to the nonidentity problem. Mere Addition occurs when parents create one additional person who has a life worth living and who does not detract from the wellbeing of others. This person would not have been created had I not merely added her to the population. To illustrate this point, consider a possible world in which every one of a small number of families lives in

⁴⁷ This assumes that having a life worth living is preferable to nonexistence. However, this does not assume that *extra* lives have intrinsic moral value.

⁴⁸ Derek Parfit, *Reasons and Persons*. 419-21.

a constant state of euphoric bliss, and everyone enjoys the same level of happiness as the next – say they live at an average welfare level of 100 units. This state of affairs continues for generations, until a contaminant leaks into one family’s water supply. Drinking the contaminated water will not harm any existing people, but it will cause an irreversible genetic defect in all future people whom this family brings into existence - the parents of this family know that this contaminant will cause every one of their children to be born with mild dyslexia. Aside from this one disability, all future people born into this family will continue to live under the same condition of sustained euphoria as their ancestors - this ailment will cause all future people from this family to live lives valued at a welfare level of 99 units.

The question arises of whether the small existing population can be morally justified in continuing to procreate, when their doing so would certainly lower the population’s average welfare over time. If the population does decide to continue procreating in the wake of this news, they will *predictably* bring down the world’s average utility. And, because the Average Principle dictates that we maximize the average level of wellbeing, any act which decreases the average level, even slightly, would certainly be wrong.

But surely, the possible world in which this family ceases to procreate is *not* better than one in which they do. Although they would predictably harm all future people born with mild dyslexia, their alternative is to significantly decrease the total number of people added to the population. With such a small population, quantity – the *total* amount of welfare – must carry some weight. Ceasing to procreate cannot possibly be justified on the basis that doing so would lower the *average* amount of welfare.

Averagism claims that it would be better to have one person living who possesses 100 units of wellbeing than to have one million people living, each of whom possesses 99 units. This implies that conversely, it would be better to have one person living who possesses only 2 units of wellbeing than to have one million people living who possess merely 1 unit each. Consider a world in which every one of one million inhabitants lives in constant torture. In this population, one couple knows that if they give birth to a child, this child will also live a life of constant torture. However, *this child* will not be tortured on his birthday every year. Averagism implies that giving birth to this child, who will experience 1.0001 units of welfare as opposed to his parents' 1 unit, is morally permissible. However, it is clear that these parents should refrain from giving birth to a child who will live a life of constant torture.

Averagism provides one solution to the nonidentity problem by mandating that we maximize the average wellbeing of future generations. It also avoids the Repugnant Conclusion. However, it also generates the Mere Addition Paradox, which implies that it would be morally impermissible to have both very happy children if they would bring down the average level of wellbeing, and that it would be permissible to have very miserable children if they would bring up the average level.

CHAPTER IV: COMPARING CONSEQUENTIALIST SOLUTIONS

Upon reflection, I believe it is clear that the worst worlds generated by the Repugnant Conclusion are clearly worse than the worst worlds generated by the Mere Addition Paradox, and thus that the problems faced by total consequentialist theories which seek to overcome the nonidentity problem are much more severe, *prima facie*, than those faced by averagism. However, when the practical implications of each of these problems are considered, it becomes less clear which of the two problems poses a greater challenge. This section will examine the drawbacks of averagism in seeking to defend totalism as the better of the two. Finally, I will defend a modified version of totalism which avoids the Repugnant Conclusion.

PROBLEMS FOR AVERAGE CONSEQUENTIALISM

The first and most significant flaw faced by averagism is that its implications are much more serious than the implications of the Repugnant Conclusion when viewed from a practical standpoint. This is to say that, in the real world, it is impossible to imagine a future in which every living person experiences a total welfare level of 1 unit throughout life. There will, hopefully, never be a point in the future when all of humanity lives lives similar to those described in the worst depletion example, or lives in constant torture. While theoretically problematic, the Repugnant Conclusion will not play into our personal concerns when making the decision whether or not to have children. The Mere Addition Paradox, on the other hand, very likely could. For this reason, the Mere Addition Paradox

is more problematic than the Repugnant Conclusion, and therefore averagism faces a more serious problem than totalism.

As discussed in the previous section, averagism leads to the result that it would be impermissible for parents to have even *extremely* happy children if these children would bring down the average level of wellbeing. While the hypothetical examples previously used to illustrate the Mere Addition Paradox were extreme, the idea which they highlight is in fact one which probably plays into the considerations of many prospective parents. Imagine that you are a member of the one couple who knows that their child will be born with mild dyslexia. You know that, even though this child will be very happy, he will have to grow up surrounded by children who have no disability, and are therefore just slightly better-off than he is. I would imagine that in reality, many parents do their best to avoid giving birth to children who they feel would have a *great* disadvantage compared to their peers. But, if these parents know that their child would be well-off on the whole, they would likely not avoid having this child altogether. This reality clashes with Mere Addition's requirement that we not lower the average level of wellbeing of future generations - it seems highly counterintuitive to say that the world may actually be made morally worse by bringing happy people into existence.

Additionally, most people would hesitate to say that people living in horrible circumstances, such as inescapable abject poverty, slavery, or other extreme oppression, should absolutely have children who will experience the same circumstances. While these children may have lives which are equally as worth living as their parents', or perhaps marginally better, most people would still condemn the fact that these parents caused such children to exist, when they had full knowledge of the fact that their lives would be

miserable. In this way, the Mere Addition Paradox is a problem with real implications, unlike the Repugnant Conclusion.

Second, I question why the average wellbeing of all people, time immemorial, should be considered relevant in determining what it is morally right to do today. The temporal relativism inherent in averagist theories seems to severely undermine them. It does not seem as though the wellbeing of, for example, the Ancient Egyptians should be weighted into decisions about whether to create future people.⁴⁹ This is especially clear when we consider cases where the average level of wellbeing in the past was a result of certain large groups being made well-off through morally wrong acts. The Ancient Egyptians, for example, would likely not have experienced the high quality of life that they did had they not been a slave-holding society. In fact, it is nearly impossible to think of an extremely well-off group in human history who has not experienced high levels of wellbeing at the cost of some less-well-off group. Is it wrong, then, to bring down the average level of historical wellbeing experienced by these groups, when doing so would involve no injustices?

Relatedly, the number of people who exist continues to increase throughout time, which affects the marginal value of individuals' wellbeing. Consider a case similar to that of the small population living in euphoric bliss – that of Adam and Eve.⁵⁰ It seems glaringly obvious that their moral concerns are much weightier than our moral concerns would be today – as the world's *only* inhabitants, they would be solely responsible for the earth's decline in average utility when they cause future generations to be born with the

⁴⁹ Ibid. 420.

⁵⁰ Ibid. 420.

burden of original sin. They can predictably lower the welfare of *everyone* living in the future. We, on the other hand, have no way to rewind the clock and influence the welfare levels of prior persons. The point of a moral theory is to define “a single time-less evaluation of actions which is valid for everyone regardless of his location in time,” not to determine at what points in human history procreation has been justifiable.⁵¹

A third problem with averagist theories is that they involve a great deal of uncertainty in estimating the future average outcomes of a population. While we can predictably create a single child with mild dyslexia, the slight decrease in utility that this person suffers as a result of disability may be outweighed by other moral concerns. Most, if not all, people who intentionally create children do so in the hopes that their child will have a better life than did they – by virtue of living in a more advanced, more cultured, or wealthier society, or by virtue of being able to provide more financial support or stability than they themselves had. The welfare calculations defaulted to in each of the problem cases described do not account for uncertainties in the future; they merely assume that lives will continue to grow worse with time. In the depletion case, for example, highly effective new agricultural and water purification techniques could be discovered to solve resource shortages, or a cure for radiation poisoning could be discovered if the nuclear policy was implemented.

⁵¹ Hudson, J. L. “The Diminishing Marginal Value of Happy People.” 125.

MODIFYING TOTALISM – CRITICAL LEVEL PRINCIPLES

As it has been described thus far, totalism has faced the Repugnant Conclusion in its attempt to solve the nonidentity problem. While this is a serious issue, the number of possible worlds it generates in which harm occurs seem to be less possible than those generated by the Mere Addition Paradox. While this, independently, seems to be a convincing argument in favor of totalism over averagism, there is further justification for preferring this view.

As they have been described in this paper, totalism and averagism in their simple forms are too restrictive. Because they reject the Person-Affecting Principle, they cannot account for anything more than sheer numerical values. They have no way of distinguishing between whether 6 billion people living at a utility value of 60 or 7 billion people living at 50 is wrong, they can only tell us that one is worse than the other. However, when we think in real terms about, for example, the carrying capacity of the planet, it is difficult to say that, on either averagism or totalism, it is morally wrong that one billion of today's human beings exist.

One way of modifying totalism, called the critical level view, eliminates worries about the Repugnant Conclusion faced by the simple totalism.⁵² Although this modified theory faces its own challenges, the problems which it generates are not as crippling to the theory as those faced by the unmodified version.

The critical level view introduces prudential moral judgments into our consideration of what the appropriate limit of welfare is for a given society. This view avoids the Repugnant Conclusion by introducing a “critical level,” or a minimum positive

⁵² Kavka, Gregory S. “The Paradox of Future Individuals.” 105.

welfare threshold, into total welfare calculations. Recall that simple total consequentialism requires summing the welfare of each individual in a given population. On critical level totalism, the moral value of a life added below the critical level threshold will now count as a negative number.⁵³ An individual's total welfare is thus calculated not from the zero point, but by subtracting the critical level from his or her total amount of welfare. For a population, the results of this individual calculation are summed. The advantage of critical level totalism over simple totalism is that it allows us to consider when a life worth living would not be desirable, from a moral standpoint.⁵⁴

The greatest strength of critical level totalism is that it allows us to consider a non-Person-Affecting definition of harm which still takes into account our preferences for what is morally desirable. Merely having a life worth living, although it may be lived at a repugnantly low utility level, is no longer an appropriate justification for maximizing total wellbeing. If a future person will predictably have a quality of life valued at a level below this threshold, it becomes morally impermissible to cause him or her to exist.

Since we may now take into account practical considerations about when a life adds value to the total level, this view avoids the Repugnant Conclusion. If, for example, the appropriate level of wellbeing in the depletion example were set at 30, it would, at some point, become morally impermissible for the *P* generations to continue procreating, assuming no solutions to the resource shortages had been found.

⁵³ In "The Paradox of Future Individuals," Kavka argues that below the critical level, the value of a life should be neutral, or "0." While neutral valuations do not change the overall structure of critical level principles, the argument in favor of critical levels is more robust when we default to Blackorby's calculation of negative welfare.

⁵⁴ Blackorby, Charles, Walter Bossert, and David Donaldson. "Critical-Level Utilitarianism and the Population-Ethics Dilemma." 197.

It could seem that introducing a critical level may be problematic because of the degree of subjectivity which it builds into total consequentialism. In breaking from a strict requirement to maximize total wellbeing, we lose any hard-and-fast standard as to what exactly constitutes morally right action. The problem with determining an appropriate critical level is that it must not be set so high as to prohibit the creation of lives which may still be judged “worth living.” Conversely, it must not be set so low that some would consider life above this level to not be, on the whole, not good (and thus fall susceptible to the Repugnant Conclusion). While there is no absolute measure by which to determine the appropriate level, perhaps the best way to achieve consensus, or at least an adequate assessment, is through policy, in the same way that the initial decisions in cases such as the depletion example were decided.

CONCLUSION

This thesis has covered both the nonidentity problem and the solutions proposed by different variations of consequentialist theories. The nonidentity problem claims that we cannot hold both the Person-Affecting Principle and the idea that we are obligated not to harm future persons to be true at the same time. Consequentialism rejects the Person-Affecting Principle, but also provides us with a means of justifying avoiding harm to future persons – namely, by maximizing the wellbeing of future persons.

One version of consequentialism, totalism, argues that we do this by increasing the total number of units of wellbeing, or by increasing the total number of people who experience wellbeing. This view generates the Repugnant Conclusion, which states that for every population experiencing a given level of wellbeing, there is an exponentially larger one wherein people experience lives barely worth living, but where the total level of welfare is higher. Averagism avoids the Repugnant Conclusion, and provides us with an additional means of solving the nonidentity problem, by mandating that we maximize the average amount of wellbeing experienced in a given population. However, the average view generates the Mere Addition Paradox, which states both that it is morally impermissible to have one extremely well-off child given a population that is marginally better-off, and that it is morally permissible to have one extremely miserable child given a population that is marginally more miserable.

I argue that the Mere Addition Paradox merits more worry from a practical standpoint, since it is almost impossible that the Repugnant Conclusion will ever be a problem in the future of human existence. I point out additional flaws with the averagist view, namely that it fails to account for the diminishing marginal value of wellbeing with

the passage of time, and that it fails to account for historical injustices which have led to heightened levels of wellbeing for previous generations.

Finally, I argue that a modified version of totalism, the critical level view, avoids the Repugnant Conclusion by enabling us to set a discretionary threshold at which lives contribute positively to overall wellbeing. While this view seems to be the most capable consequentialist means of solving the nonidentity problem, it faces the issue of injecting subjectivity and ambiguity into the original totalist calculation.

Different variations of consequentialism provide strong justifications as to why we should avoid predictably harming future persons. However, consequentialism rejects the Person-Affecting Principle, and thus do not fit the framework for a possible theory of morality which Derek Parfit originally laid out in *Reasons and Persons*. Whether any moral theory exists which can fit the constraints of the nonidentity problem in explaining our responsibilities towards future generations remains to be seen. Until a solution is discovered, the nonidentity problem will remain one of the most pivotal puzzles in ethics – and that, in itself, is remarkable.

WORKS CITED

- Blackorby, Charles, Walter Bossert, and David Donaldson. "Critical-Level Utilitarianism and the Population-Ethics Dilemma." *Economics and Philosophy* 13.02 (1997): 197-230. Print.
- Dell'Amore, Christine. "How a Man Produces 1,500 Sperm a Second." *National Geographic*. National Geographic Society, 18 Mar. 2010. Web. 28 Apr. 2013. <<http://news.nationalgeographic.com/news/2010/03/100318-men-sperm-1500-stem-cells-second-male-birth-control/>>.
- Kavka, Gregory S. "The Paradox of Future Individuals." *Philosophy & Public Affairs* 11.2 (1982): 93-112.
- "Kentucky QuickFacts." *United States Census Bureau*. U.S. Department of Commerce, 14 Mar. 2013. Web. 28 Apr. 2013.
- Hudson, J. L. "The Diminishing Marginal Value of Happy People." *Philosophical Studies* 51 (1987): 123-37. Web.
- Parfit, Derek. *Reasons and Persons*. Oxford: Oxford University Press, 1984.
- Roberts, Melinda A. *Child versus Childmaker: Future Persons and Present Duties in Ethics and the Law*. Lanham: Rowman & Littlefield, 1998. Print.
- Roberts, Melinda. "The Nonidentity Problem." *Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, CSLI, Stanford University, 21 July 2009. Web. 28 Apr. 2013.
- Temkin, Larry S. "Intransitivity and the Mere Addition Paradox." *Philosophy & Public Affairs* 16.2 (1987): 138-87.