Procreation, Population and Permissibility

Sarah Tilsley
B.A. (University of Sydney)

Submitted for the Degree of Master of Arts
Department of Philosophy
Monash University
17th October 2014

Copyright Notices

Notice 1

Notice 2

Table of Contents

Abstract

In this thesis I examine the permissibility of bringing a child into being from a person-affecting position. I argue that as long as a future person's life is somewhat worth living, and that their lifetime wellbeing has been maximised, they will not be wronged by being brought into existence. I further argue that it is permissible to bring about future population X, as long as those living with X have lives that are somewhat worth living. Finally I argue that in addition to the permissibility of procreation, becoming a parent creates many goods for the potential parent.

Declaration

I declare that this thesis contains no material which has previously been submitted for a degree or diploma in any university and, to the best of my knowledge and belief, this thesis contains no material which has previously been published or written by another person except when due reference is made in the text of the thesis.

Signed:

Sarah Tilsley

Acknowledgements

I would like to thank my supervisor, Toby Handfield, for his help in the development of this thesis and for being so accommodating in regards to my work/study-life balance. I would also like to thank Linda Barclay for advice and feedback on chapter three. Thanks also to my family for their patience and support.

1. Introduction

Many people reach a point in their lives, at a culturally appropriate age, where they contemplate, deeply or not, having children. There are many reasons for having children: social, cultural, religious or family expectations may weigh heavily into the decision; one may feel a duty to one's partner to fulfil their desires; one may have strong patriotic reasons or because becoming a parent is a strong part of one's identity and that particular relationship would contribute significantly to the success of one's life. There are also many reasons not to have children, primarily because there is reason to believe that the prospective parent or child or some other person or persons will be adversely affected by bringing that child into being.

The act of bringing a child into being affects many people. It affects the parents who then become responsible for the child's wellbeing and development. It affects the immediate family who then make room for the child and may have their interests compromised, or who may also become responsible for the child's upbringing. It affects the community small and large who may be responsible at some point for the child's wellbeing. It affects the remaining population to an extent who now have to share finite resources with another being. It affects future generations who will potentially feel the effects of that person's existence and their actions within the world. And last but certainly not least it affects the person who is brought into being. This is the act that causes their existence – without that act they may never have existed. In fact Parfit (1986) argues that only that particular act of conception at that particular moment in time could cause that particular person to come into being. Others argue that even a

moment's delay in conception could have resulted in a different person coming into being. As such, the act of bringing another person into being has many and serious consequences.

In this paper I examine the ethics of procreation in what is considered by some to be a heavily populated world. My examination consists of two main questions: (i) when is it wrong to bring a child into the world? (ii) is it permissible to increase the population?

To answer the first question I will look at whether we can harm a person by bringing them into being. In doing this I will focus on Roberts' Personalism approach which takes the form of a person-centered consequentialist theory.

Traditional forms of consequentialism are interested in maximising good (happiness, pleasure, wellbeing), thus acts and their outcomes are evaluated on how they affect the total good. Under an aggregative approach, permissible acts maximise the total good. Individual well-bring is important only to the extent that it is a factor in determining total good. Roberts calls this approach *totalism*.

A slightly different approach looks at the average good; under this approach, acts should aim to increase the average good of the population. This type of consequentialism still takes an overall maximising approach and does not consider the individual good any further than in how it contributes to the average good.

A personalist approach, however, specifically focuses on how acts affect the individual. Roberts argues that it still takes a maximising form, but the good to be maximised is personal good, or individual wellbeing (1998, p. 6). The basic idea of personalism according to Roberts is "among other things, that any person whose wellbeing has been maximised has not been wronged and that with

restrictions, people whose wellbeing we have failed to maximise sometimes *have* been wronged." (Roberts, 1998, p. 6)

In considering whether we can wrong someone by bringing them into being I start by looking in a broad manner at what our duties are towards future people and what set of people constitute future people. I argue that we should only take into consideration potential people, those who will actually exist, and exclude possible people.

I then move into a more focused approach specifically considering the ethics around bringing a person into existence.

I will look closely at the Non-Identity Problem, as one of the main issues in determining whether we can harm a person by bringing them into being. This raises a subset of questions such as whether we benefit or harm someone by bringing them into being, whether non-existence is ever a better option, and whether harm is absolute or comparative.

I will critically analyse Roberts' attempt to answer the non-identity problem, and argue that while her approach is flawed, I agree with her conclusion that as long as their life is somewhat worth living we do not harm someone by bringing them into being.

In a world of finite resources, increasing the population could potentially mean that others already living could have their welfare decreased by having to share the resources with yet another being. For this reason, in the third chapter I examine Parfit's Repugnant Conclusion, as a way of determining whether increasing the population is permissible. Parfit's famous conclusion stated that a world in which there was a vast population in which every person has a very low

level of wellbeing would, under an aggregative approach, be a better outcome than a smaller population in which every person has a high level of wellbeing. Parfit found this conclusion not just counter-intuitive, but repugnant. I will examine some of the responses to the Repugnant Conclusion and show that they are not satisfactory. Following a critical analysis of Roberts' response to the Repugnant Conclusion, I will conclude that as long as those in population X have a life somewhat worth living and their wellbeing has been maximised, then that population is permissible. In conjunction with my conclusion in chapter two, this allows me to conclude that we do not wrong someone by bringing them into being, and that bringing more people into being is permissible as long as their life is somewhat worth living and their wellbeing has been maximised.

Finally, I consider the effect bringing a child into the world can have on the person who wishes to become a parent. Once I have established that procreation is permissible (as long as those created have a life that is somewhat worth living and their wellbeing is maximised), I want to further establish that there are benefits to the potential parent in procreating.

2. When is procreation permissible?

When considering how our actions might affect future people, it is important to clarify who exactly we are referring to, who we think might be affected by our actions.

A person-affecting approach like personalism naturally considers those who are already in existence. When we consider the effects our actions have on others, we automatically consider those who are already in existence. The next question is how we should think about the consequences of our acts on those who are not yet in existence, future people. Most theories argue that the interests of future people should be given the same weight as those already in existence. But the question goes further, and by asking "which future people?" we differentiate between potential people, people that have not yet, but will come into existence and merely possible people, people that may not ever come into existence.

By potential people we mean people who do not yet exist, but that will come into existence. For example, if Jane is planning to have three children over the course of her life, we can consider those children potential, if all goes ahead as planned, because those children will come into existence. A merely possible person on the other hand is a person we cannot say will definitely exist. For example, if I was hypothesising about the possibility of every woman alive now giving birth to three children each, and the population this would result in, that would be speculating on possible people. It is not the case that every woman alive will have three children, and so that group of people is merely possible.

In this paper, I take the approach that Roberts (1998) takes in her account of personalism, that, for the most part, considering the effect of our actions on possible people is not relevant. (Roberts gives exception in cases where the lives of possible people would be less than one worth living.)

However, a comprehensive account of a person-affecting consequentialism must consider the effect of acts on potential people, people that will actually come into being. Either way the research question implies that we are considering the effect of our acts on potential people. So without going into the arguments for and against including possible people in a person-affecting consequentialism, I will limit the scope of this paper to potential people, and that is the group I am referring to when I use the terms future people or future persons. Furthermore, the scope of this paper excludes considerations of how our acts affect those already in existence, except where it is warranted for comparative purposes.

Personalism

A person-affecting approach by its nature determines whether an act is good or bad by its effect on persons. A person-affecting approach is a type of consequentialism that differs from traditional forms of consequentialism which take an aggregative approach (Roberts, 1998). Consequentialists generally agree that agents ought to bring about the outcomes that maximise good. It is in determining what 'good' is, that they differ.

A utilitarian for example would be concerned with maximising happiness, and they would usually consider the total sum of happiness, or an average level

of happiness across a total population; this is referred to as an aggregation of happiness (or good). A person-affecting approach however makes the distinction between good on the whole (the total or average good of a population) and personal good. A person-affecting approach is more concerned with individual person-based good, or person-based welfare (Roberts, 1998). So in taking a person-affecting approach we are concerning ourselves with how acts affect the wellbeing or welfare of individual persons. To take Roberts' person-affecting framework, which she calls personalism, personal wronging is brought about when "we can but do not increase [a person's] wellbeing at no cost to others"[1] (1998)

Principles of Personalism

In light of this, let's examine more closely the principles that make up Roberts' personalism framework, and use them as a guide for navigating our way through the landscape of maximising the welfare of future persons.

The first principle that Roberts puts forth is that we cannot wrong anyone who never exists; she formulates it this way:

N* = *s* is not wronged by agents in *X if s* never exists in *X*.

N* attempts to answer those who propose that we can wrong someone by not bringing them into existence. The establishment of N* denies the symmetry that is implied in aggregative approaches. Where someone taking an aggregative

[1] Roberts does clarify that failing to maximise is a necessary condition of personal wronging, but not always a sufficient condition.

approach would say it is wrong *not* to bring a happy *s* into existence, N* says it is *not wrong* if *s* never exists.

The second principle that Roberts puts forward states that a person whose lifetime wellbeing has been maximised has not been wronged. She states it thus:

> M* = *s* is not wronged by agents in *X if*, for each world *Y* accessible to such agents, *s* has at least as much wellbeing in *X* as *s* has in *Y*.

Roberts refers to alternative worlds (such as X and Y) for subject s, such that s could have existed in either world, and that the worlds can be ordered relative to each person in respect of how good each world is for that person. Roberts, and I, say that a world is *accessible* to a person if some person at some time could have brought about that world. For example I exist in a world where I have only one sibling, although a world where I have more than one sibling is accessible to me.

For Roberts, these two principles provide sufficient conditions for a person not being wronged.

In order to cover how persons *are* wronged, she offers the following two principles.

> D* = *s* is wronged by agents in *X if s* exists in *X* and there is some world *Y* accessible to such agents such that:
>
> (i) *s* has more wellbeing in *Y* than *s* has in *X*;
>
> (ii) for each person *s'* who ever exists in *X*, either *s'* has at least as much wellbeing in *Y* as *s'* has in *X* or *s'* never exists in *Y*; and

(iii) for each *s'* who ever exists in *Y*, *s'* exists at some time in *X*.

D* is presented in such a way as to avoid scenarios where it may appear that personal wronging has occurred, but where a theory of fairness would excuse such cases. Roberts states that personalism does not include in itself principles regarding fairness, but could very well be used in conjunction with a fairness theory. For example D* would not condone a situation where *s* would be better off in *Y* than in *X*, if there also existed in *Y* another non-*X* person whose ill-being in *Y* made *s*'s life better. The wrong can only be said to have taken place where *Y* contains the same or fewer people than *X*, and that all the people that exist in both *Y* and *X* are at least as well off in *Y* as in *X*.

The last principle that Roberts proposes regards wrongdoing. She states it thus:

P* = *X* is permissible for agents at *t if and only if* no person who exists at or after *t* is wronged in *X* at or after *t*.

This principle is important for the consideration of how our acts affect future persons, as it explicitly states that as agents we are permitted to choose X (an action or state of affairs) if no person who exists at the time we are choosing, or *in the future,* is wronged in that action or state of affairs.

This principle allows us to establish that we have a duty toward future people not to wrong them, or more explicitly to maximise their wellbeing.

The Non-Identity Problem (NIP)

In establishing her theoretical framework, Roberts makes an attempt to overcome the infamous non-identity problem (henceforth NIP). Her denial of NIP stems from her proposals that future persons can be wronged (P*, not controversial) and that there is always some alternative (accessible world) that those persons could have existed in where their wellbeing was greater or maximised (M*, more controversial). Before going any further into Roberts' account of personalism vs NIP, let's first examine the non-identity problem in more detail, and some attempts to overcome it.

The NIP is described by Parfit in *Reasons and Persons* (1984)[2], and it rests on his gametic essentialist beliefs. Gametic essentialism maintains that specific personhood or identity is dependent on the time of conception, and more particularly the two particular gametes a person develops from. On this view, any change in the timing of conception will affect the identity of the resulting person. (Parfit is generous and gives one month's lee-way, although Kavka proposes that even a five minute delay in conception would result in a different person being conceived.)

Following this we can see that many choices in life could affect who comes into being. Choosing an Arts degree over a Science degree results in meeting an entirely different group of people in life. Choosing to go backpacking in Europe or going into paid employment and choosing to get a taxi home rather than the train are both decisions that could affect the identity of a future person.

[2] Also raised independently by Kavka (1982) and Schwartz (1978).

Additionally, social and policy decisions can have the same effect. A government may decide to spend $10b upgrading public transport facilities which means everyone gets home half an hour earlier than previously, resulting in differing conception times and hence affecting the identities of future people. In light of this, Parfit claims that because identity and, as such, existence, are entirely dependent on personal, social and policy choices made prior to conception, choices that could also potentially affect the *quality* of life of future people, any claim of wrongdoing on behalf of future persons cannot be upheld.

Consider this example given by Parfit: a government is making a decision on a resource policy. One policy will involve a conservation of resources, which ensures a steady increase in quality of life over the next few centuries. The other policy involves a depletion of resources that will result in a higher quality of life for those in existence in the next two hundred years, but a lower quality of life for those living beyond that. This situation is explained in the following diagram.

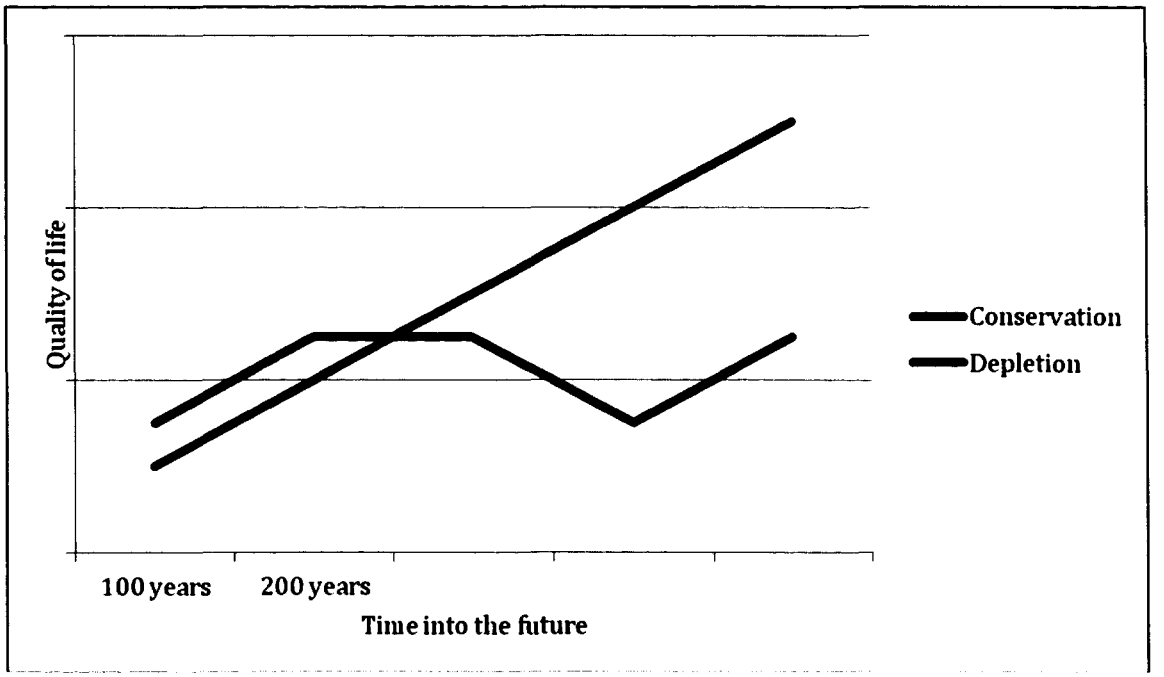

The government chooses the depletion policy, which results in a higher quality of life for those living across the next two hundred years, which in turn results in people having to work less, being able to spend more time with family and friends and in turn results in a particular set of persons being brought into existence. Those persons who live in the period beyond two hundred years will have a lower quality of life, although a life still worth living, than if the government had chosen conservation. However, their existence is counterfactually dependent on the government choosing the depletion policy.

If the government had chosen the conservation policy, an entirely different set of individuals would have been brought into being. So can we say that for those people, the choice of the depletion policy was worse? Parfit says 'no'; he argues that because the alternative for them would have been nonexistence, the flawed existence is not worse for them. On a person-affecting view what is bad must be bad for *someone*, so we cannot say that choosing the depletion policy is bad for those people. Their existence is necessarily dependent on choosing the depletion policy. This example represents a situation of environmental degradation, such as high levels of pollution, or increased use of non-renewable resources; it's a situation that many people would feel strongly about and would intuitively disagree with. Let's then look at some other examples that take a slightly different form but still present the non-identity problem.

Consider the example of Mary, a 14-year-old girl who wants to have a child.[3] If she chooses to have the child at the age of 14, she will not be able to provide the child with a very good quality of life purely because she is young and does not have the resources. If she chooses to wait a few more years, the child she produces at a later date will have a better quality of life. Mary decides not to wait, she brings a child into the world, the child's life is worth living, but it does not have the higher quality of life Mary would have been able to give a child had she waited. However, according to the non-identity problem, we cannot say that the child is worse off or harmed, because had Mary waited until she could have provided better for her child, a different child would have been born. The life that the child has, born to Mary at age 14, is the only life that child could have known. The two alternatives for that child would have been the lower-quality-of-life (flawed) existence or non-existence.

The conclusion that the lower quality of life for a particular set of people or person is not bad for them seems counterintuitive. It seems intuitive to suggest that a better quality of life should have been provided for those people, but according to NIP, no such alternative was available for them. In both the examples provided above, the government depletion policy and 14-year-old Mary, the existence of the persons was counterfactually dependent on the choices made, had it not been for those decisions, they would not have come into existence.

In summary the non-identity problem suggests that claims about a future person's harmed existence cannot be upheld because attempts to avoid the adverse effects could also affect that person's, or set of persons', conception. The

[3] Similar to the example given by Parfit in Reasons and Persons (1984)

conclusion of the non-identity problem is that for some people, an existence that contains some ill-being is the only existence that would have been possible for them, and therefore those people cannot be said to have been harmed.

There have been a number of responses to the non-identity problem and attempts to resolve it. They range from some more extreme cases such as claims that all existence is harmful (Benatar, 2006) to less extreme attempts to redefine harm. Responses to NIP typically involve discussions on whether we can benefit or harm someone by bringing them into existence, and whether the harm (or benefit) of existence is comparative or absolute and thus whether nonexistence has a value. Some responses also include views on the asymmetry of procreational duties. I will go into more detail on each of these areas as I examine some of the responses to NIP.

In Creating A Person, Can We Benefit Or Harm Them?

The conclusions of the NIP as presented by Parfit, Kavka and Schwartz seem untenable; certainly from a person-affecting view point, it follows that for some action to be deemed morally wrong it must wrong *someone*, *someone* must be harmed or worse off. But worse off than what? If we examined a case where an already existing person has been harmed, say for example Ben is hurt in a car accident and his leg has to be amputated, we would say that Ben has been harmed. In missing one leg, he is worse off than before the car accident. But when we are discussing cases involving quality of life and the creation of a person, exactly what are we comparing their current state to? Is non-existence a comparable state such that we can say someone has been benefited or harmed?

One response to this is yes, we can harm someone by bringing them into existence, and yes, nonexistence is a comparative state, and the claim of harm is a comparative claim. This is the view that Nils Holtug takes in his paper *On the Value of Coming Into Existence* (2001). Holtug puts forward what he calls *The Value of Existence View*, which is essentially the view that we can benefit or harm someone by bringing them into existence. Holtug is explicit in pointing out that he does not consider existence to be intrinsically valuable, that in every case it is better to exist than not exist, rather his view is that if life is, on balance, worth living, then it is better to exist than not. If, on balance, life is not worth living, then it is possible to say that existence is not better than nonexistence. What does he mean by "on balance"? He means that on the whole if within one's life the goods outweigh the bads, then life is worth living, and we can say that one is benefited by being brought into existence. To say that someone is harmed is to say that on the balance, the bads in their life outweigh the goods, and as such they would have been better off not existing.

David Benatar (2006) on the other hand seeks to refute this position. Benatar claims that discussions about what makes a life worth living are misguided, and that a further distinction needs to be made for those discussions to be helpful. The distinction he is referring to is between a life worth starting and a life worth continuing; he calls them the present-life sense (of a life worth living), and the future-life sense (of a life worth living). Benatar claims that it is misguided to make assessments on whether a future life is worth living using the present-life sense because the threshold for determining whether a life is worth living for someone who is already existent is much lower than for someone who is yet to come into being. To understand this better, consider the example of

Down syndrome, while we would not consider Down syndrome to be a condition so bad that it is worth ending a life we would (and do) consider it to be such that it is better not to bring into existence someone with the condition. Thus, for Benatar, weighing up goods and bads within a life and then comparing it to nonexistence is not satisfactory for determining whether someone is benefited or harmed in being created.

Is Non-Existence Comparative Or Not?

The approach that Holtug takes makes the assumption that nonexistence is a comparative state. Holtug uses a number of theories of wellbeing in order to be able to determine when existence is better (or not) than nonexistence. In doing this he assumes that nonexistence can be a state that is sometimes better for someone. The issue with this is how can we refer to someone and discuss their wellbeing if their state is nonexistence? We cannot refer to *someone* who is non-existent.

Benatar on the other hand argues that nonexistence is not a state for *someone*, he argues that it is comparable only in the sense that one can objectively compare two states: one in which person *s* exists and one is which *s* does not exist. The state in which *s* does not exist cannot be better (or worse) *for her*. However, Benatar does not take a totalist approach. His is a person-affecting approach. He explains it like this:

> "I shall not claim that the never-existent literally are better off. Instead, I shall argue that coming into existence is always bad for those who come into existence. In other words, although we

> may not be able to say of the never-existent that never existing is always good for them, we can say of the existent that existence is bad for them... Once we acknowledge that coming into existence can be a harm, we might then want to speak loosely about never coming into existence being 'better'." (Benatar, 2006, p. 4)

Thus Benatar carefully avoids a comparative approach as well as an aggregative approach by claiming both that nonexistence is not a state for someone, and that nonexistence can only be good in the sense that existence could be bad for a potential person.

Benatar's argument leads to the conclusion that we always harm someone by bringing them into existence. This rests on his interpretation of the asymmetry of pleasure and pain. He argues that the absence of pain is always a good, even if there is nobody to experience the good. While on the other hand the absence of pleasure is not necessarily bad, but it is not necessarily good either, and for the absence of pleasure to be a good there has to be someone to experience it as a good. This asymmetry in pleasure and pain leads Benatar to the conclusion that any pain in existence is a bad that outweighs any possible bad in its opposing absence of pleasure in nonexistence. More precisely he claims that nonexistence with the absence of pleasure and pain is better than existence that contains pain, and since all existence contains some pain, nonexistence is preferable (Benatar, 2006).

This conclusion is, as Benatar foresees, unpalatable to many. It is counter-intuitive in the extreme as many of us feel that our lives are worth living, even if they do contain some amounts of pain. For this reason I do not defend Benatar's

position. However, his interpretation of the asymmetry of pleasure and pain provides a good argument for why we do not have a duty to produce happy children, which is a major problem for those that take an aggregative position.

Both Benatar and Holtug maintained that we can harm someone by bringing them into existence by claiming that harm is a comparative state. That is, their arguments rested on the claim that in some circumstances the harm/pain/bad contained within existence is so bad that nonexistence would have been comparably better.

An alternative approach to this is to argue not that harm is defined by being "worse-off" than we could have been, and is therefore comparable to non-existence, but that harm is absolute, or in the case presented by Harman, comparable to a "healthy bodily state"(Harman, 2004). Taking an absolute harm approach avoids the non-identity problem by avoiding the idea that in some cases nonexistence is better, which is the crux of the non-identity problem. Harman chooses to define harm as "causing pain, bodily damage, early death or deformation" (2004, p. 92) rather than causing someone to be worse off than they would otherwise have been. Treating harm as absolute rather than comparable means we can count harms and benefits without having to account for *who* those harms and benefits are affecting.

Harman's approach is similar to Benatar's in that it rests on an asymmetry of sorts, and is careful not to take an aggregative approach. Harman's asymmetry is this: there are reasons to benefit and there are reasons against harms, in a situation where benefits and harms are bestowed, the reasons against harms outweigh the reasons to benefit. It's this balance of reasons (or

asymmetry in the balance of reasons) that forms the crux of her argument. Harman states:

> "Reasons against harms are so morally serious that the mere presence of greater benefits to those harmed is not in itself sufficient to render the harms permissible: when there is an alternative in which parallel benefits can be provided without parallel harms, the harming action is wrong." (2004, p. 93)

In the case of the depletion policy vs the conservation policy outlined above, Harman would advocate for the conservation policy on the basis that although adopting the depletion policy would both benefit (by bringing into existence) and harm (by reducing access to resources) those individuals that result from the choice, it is nevertheless morally wrong to choose that policy when there is an alternative that could confer the same benefits but avoid the harms (the conservation policy).

It appears from this that Harman is taking an aggregative approach, by treating harms and benefits as absolute, and not accounting for who they affect. However, Harman is careful to avoid an aggregative approach. At the beginning of her paper, she states that one of her goals is to "vindicate the presence and explanatory value of reasons against harm in the cases that generate the non-identity problem."(p. 90) Harman explicitly states that by achieving this goal she will show that the impersonal view (aggregative approach) is false (p. 91). So how is it possible to take an approach where benefits and harms are accounted for without accounting for who they affect, but also avoid an aggregative approach. Harman does it by treading very carefully (as does Benatar) around

explaining how the benefits are accounted for in regards to future persons. As Benatar calls it, she is taking a "loose person-affecting approach", whereby rather than saying *someone* has been wronged because they are worse off, she says it would be wrong, based on the balance of reasons, in virtue of the harms that would afflict future persons.

Although Harman's approach avoids aggregation and thereby avoids the implied symmetry, that if we have a duty not to create harmed individuals, we also have a duty to create happy individuals, she does actually advocate that we have a duty to create happy individuals. This also stems from her balance of reasons argument, that there are reasons in favour of a course of action in virtue of the benefits to the future individuals that would be created (2004, p. 98). This stance also implies that we have a duty not to create anyone who could be harmed and raises the question of when and why it is ever permissible to cause harm in creating. Harman's absolute harm approach may get around the non-identity problem, but it gets her into a different difficulty: it means she has to take the approach that any state of "harm" (pain, early death, bodily damage and deformation) is bad, and therefore causing it is wrong. She makes this clear in addressing the cases of the teenage mother and the temporary condition. She says: "In both cases, conceiving causes the resulting child to be in a bad state...Because each woman has an alternative in which she provides parallel benefits without parallel harms, her harming action is wrong." (pp. 94-95) In these cases the harms caused are emotional problems and deafness respectively, which some may consider to be minor harms. In taking this stance, Harman is claiming that we have a duty to bring children into the world *if and only if* they will be in possession of a "healthy bodily state" and not be suffering any kind of

harm caused by or prior to the conception. This seems like it could be an impossible task. However, Harman does allow for two scenarios where harming someone by bringing them into existence is permissible: i) where the alternative could lead to extinction (thus she directly opposes Benatar's position) and ii) where the harm in question is not so serious and there is no alternative in creating someone who is not similarly harmed. This second scenario is a direct reference to the permanent condition case.[4]

Although this softening in her position seems necessary to avoid any counter-intuitive ideas, it does seem to back-track slightly. If harm is absolute is it possible to have harm that is "not so serious"? This implies a ranking system, one on which 'a life worth living' could potentially sit (and hence so could nonexistence). However, although I feel that this part of her position is contradictory, I don't reject her claim that it is permissible to cause some degree of harm where there is no alternative in creating someone who is not similarly harmed.

So far I have discussed three attempts to avoid the non-identity problem and although the arguments have their merits, they have avoided taking a direct person-affecting approach. This brief examination highlights what seems to be an unavoidable implication of taking a person-affecting approach: that you are then faced with NIP. Is it possible then to take a person-affecting approach and still avoid NIP?

---

[4] In the permanent condition case a woman has a permanent condition whereby when she conceives, the child she gives birth to will be deaf. As opposed to the temporary condition case, there is no option for this woman to wait a period of time for the condition to pass and avoid bearing a baby that will be deaf.

Bykvist attempts to do this in his paper *The Benefits of Coming into Existence* (2007a). He argues that neither existence nor nonexistence can have comparative benefits or harms, and that we benefit people by creating them only in the sense that we can affect things such that they are good or bad for them. His argument is actualist: for a person to enjoy a benefit it is not enough that a state of affairs exists (or is possible) that is good for that person, that state of affairs must *obtain* (must become actual) for the person to enjoy the benefits. By taking this approach Bykvist appeals to a person-affecting principle, and thereby avoids a totalist approach and also avoids facing the problem of the asymmetry.

Bykvist's approach is appealing because it avoids an aggregative position and rejects the idea of comparative harm.

Roberts also avoids an aggregative approach in her account of personalism (1998), however Roberts embraces a comparative notion of harm which I believe leads to the downfall of her position. But her version of comparative harm is slightly different to that of Holtug and Benatar. Roberts takes a tri-comparison approach rather than a bi-comparison approach to harm, it is in this development of tri-comparison that she attempts to avoid the non-identity problem.

Before we delve into Roberts' account of personalism and how it deals with NIP, let's briefly recall the four principles that make up personalism:

N* = *s* is not wronged by agents in *X if s* never exists in *X.*

M* = *s* is not wronged by agents in *X if,* for each world *Y* accessible to such agents, *s* has at least as much wellbeing in *X* as *s* has in *Y.*

D* = *s* is wronged by agents in *X if s* exists in *X* and there is some world *Y* accessible to such agents such that:

(i) *s* has more wellbeing in *Y* than *s* has in *X*;

(ii) for each person *s'* who ever exists in *X*, either *s'* has at least as much wellbeing in *Y* as *s'* has in *X* or *s'* never exists in *Y*; and

(iii) for each *s'* who ever exists in *Y*, *s'* exists at some time in *X*.

P* = *X* is permissible for agents at *t if and only if* no person who exists at or after *t* is wronged in *X* at or after *t*.

We should note here that Roberts' argument rests on the understanding that she is comparing possible worlds for a subject. As discussed previously, Roberts refers to alternative worlds (such as *X* and *Y*) for subject s, such that s could have existed in either world, and that the worlds can be ordered relative to each person in respect of how good each world is for that person. A world is *accessible* to a person if some person at some time could have brought about that world.

Roberts argues that we do not harm someone, even if they have some bads in their life, if we have maximised their wellbeing. Roberts' interpretation of maximising someone's wellbeing differs from Bykvist's. Roberts would agree with Bykvist that we can benefit someone by affecting things such that they are good or bad for someone once they are already in existence. This is at the heart of a person-affecting approach. But Roberts extends maximising someone's wellbeing to choosing the best possible world for them. Roberts claims that NIP relies on M*: *s* is not wronged by agents in *X if*, for each world *Y* accessible to such agents, *s* has at least as much wellbeing in *X* as *s* has in *Y*, and that NIP assumes that there are only two possible worlds, existence and nonexistence.

She calls this a "type 2-alt case", meaning that there are two alternatives for *s*: a flawed existence or nonexistence. But Roberts introduces a different type of case, she calls it a type 3-alt case, where there are three alternatives for *s*: flawed existence, an unflawed existence and nonexistence.

> "I will call "type 3-alt' (for "three-alternative") those cases in which there is at least one accessible world that is better for *s* than *s*'s actual, flawed world... In this kind of case, the fact that a better existence is available for the subject means that M* does *not* justify the agents' imposing on the subject the lesser existence, for M* implies no wrong to the subject only when the subject's wellbeing has been maximized." (1998, p. 95)

According to the person-affecting principle M* when considering a type 3-alt case, NIP is not a problem because we cannot say that *s*'s wellbeing has been maximised, and we cannot say that s has not been wronged. In fact according to D*, *s* has been wronged.

This is what I call Roberts' tri-comparison approach. Whereas Holtug, Benatar and Bykvist all considered harm as a possible comparison between a flawed existence and nonexistence, Roberts introduces a third state, a possible unflawed existence to the comparison. Roberts' position is not an absolute position like Harman's. Although Harman's position also appeals to an alternative world Y in which a possible person exists unflawed, she doesn't take into consideration *who* is affected in either world. Roberts presents a case where it is possible for the very same individual to exist in an unflawed state, thus taking an approach that is both comparative and person-affecting.

Roberts then goes further to claim that every case of existence is a type 3-alt case, and that every case where NIP has been assumed can in fact be reimagined as a type 3-alt case thus ruling out any possibility of NIP.

> "How plausible is the assumption that non-identity cases are type 2-alt cases?" (Roberts, 1998, p. 96).

She goes on to show that three well-known non-identity cases can be reimagined as type 3-alt cases. In the slave child case[5], she says it is possible to imagine that the very same child could have been conceived if the parents hadn't signed the agreement to give the child up as a slave. She applies this same re-imagining to the pleasure pill case.[6]

The depletion case too she claims can be reimagined as a type 3-alt case:

> "How plausible is it that there is no accessible world in which some members of an impoverished future population in the depletion case enjoys ample resources? The ill-being in this case is presumably a function of the excessive size of the future generation. This means only that not all members of the future generation could have existed without resources being spread very thinly across the population. But it does not mean that some member could not have existed had resources not been spread so thinly. All that is needed, to avoid the implication from M* that depletion wrongs no one, is an accessible world in which at least one member of the

---

[5] The Slave Child case is proposed by Kavka in "The Paradox of Future Individuals". Briefly described it is a scenario where two people sign an agreement whereby the child they produce post the agreement will be given up and sold into slavery. They conceive the child and it becomes a slave. The NIP component is that this child is the result of signing the agreement and thus would not have been brought into being, except in these circumstances, and so no harm can be said to have been done (as the alternative would have been nonexistence).

[6] The Pleasure Pill case is also introduced by Kavka in "The Paradox of Future Individuals", and has a similar structure to the Slave Child case, where two people pause before intercourse to take a pleasure pill, that causes the resulting child to be mildly handicapped. Had they not paused to take the pill a different child might have been conceived.

depleted world enjoys greater resources and a correspondingly higher level of wellbeing."

In appealing to type 3-alt cases, Roberts claims that the NIP fails to take into account that the agent could have improved that same numerically identical child's lot but did not, and in doing this the child's wellbeing has not been maximised. She claims that an agent must aim to bring a better-for-*this*-child world about. But it seems to me that Roberts is too quick to claim that a better-for-this-child world is *always* accessible.

It seems to me that there are two factors at play in a non-identity case, two factors that the parents may or may not have control over. The first factor (i) is the environment the child is brought into. It is sometimes the case that the parents can affect the environment such that it improves the child's wellbeing (in comparison to not affecting it). For example in the slave child case the parents could indeed have not signed the agreement, and thus brought the child up in a loving family as opposed to selling it into slavery. The second factor (ii) is the identity of the resulting child. If we accept gametic essentialism,[7] as Parfit does and I do, then the timing of conception affects the identity of the resulting child. So in one sense the parents have control over this, they can cause one identity to be born or another depending on when they conceive. However, to say that they have *control* is too strong. They cannot *choose* an identity as such, it is only the case that had they delayed conception (or brought it forward) a different identity would have come into being. In arguing that all NIP cases can be re-imagined as

[7] Gametic essentialism is the idea that identity is directly linked to the timing of conception and any change in timing will affect the identity of the person brought into being..

type 3-alt cases, Roberts is making two assumptions (that correlate to the two factors listed above): firstly, that it is possible that the numerically identical child can be brought into existence in a different environment. This assumption appeals to the first factor in NIP cases, identity. I agree with Roberts that it could be the case in some type 3-alt cases that a numerically identical child could be produced in different environments: "It is like being conceived in Maine. The very same person in fact conceived in Maine could have been conceived in Ohio instead." (1998, p. 96)

The first assumption is more problematic, by arguing that all NIP cases are type 3-alt cases, she assumes that the parents always have control over the environment the child is born into, and that there is always a better-for-*this*-child world (let's call it world *Y*) that they could have chosen. Furthermore she assumes that they can combine their control over access to world *Y*, with the ability to bring into the world a numerically identical child. Thus in her re-imagining of the slave child case, it is possible for the parents to bring the very same child into the world without signing the slavery agreement, hence they have wronged the child by not maximising its existence.

However, I argue that parents do not always have access to world *Y*, and in cases where they do have access to world *Y*, they cannot combine this access with the ability to bring into existence the very same child they would have brought into world *X*. Let's take the example of Mary, the 14-year-old girl who wants to have a baby. If Mary has the baby at the age of 14, she will not be able to give the baby its best start in life, purely because of her age. This is world *X*. Mary does have access to world *Y*, she can wait until she is 21, by which time she will be more capable of raising a happy healthy child. However, in choosing world *Y*,

Mary is then forced to bring into being a different child than would have existed in world *X*. For the child that would have existed in world X, this is clearly a type 2-alt scenario, that child would either have had a flawed existence or no existence at all.

Let's take another example: imagine a couple living in the slums of India, who live in dire poverty with no hope of rising out of it in their lifetime. If they have a child, the child will also live in dire poverty and will be subject to malnutrition and ill health. We can imagine another world for that child, where its parents don't live in the slums in India and the child has all the provisions it needs for a happy and healthy life. However this world is not accessible to the actual parents. Unlike Mary, they have no control over the environment their child will be brought into. This is also clearly a type 2-alt case.

This argument can also be applied to Roberts re-imagining of the depletion case as a type 3-alt case. In the alternative accessible world she is proposing, those choosing which policy to implement are in control of the (future) environment, and can choose world *Y*. However, this then affects factor (ii) in NIP, and will affect the identities of those eventually born. Thus it is not really a type 3-alt case. If we look at it from another perspective, parents already living in the depleted world, who are controlling the identities brought into being, do not have access to world *Y*, the conserved world, in just the same way that the parents living in the slums in India do not have access to a better-for-*this*-child world.

So in contradiction to Roberts' claim, I would argue that it is plausible that NIP cases are type 2-alt cases. Further, I would argue that it is *only* type 2-alt cases that are NIP cases.

Roberts foresees the objection to her theory that NIP cases can all be re-imagined as type 3-alt cases and she sets out to answer it by saying "Perhaps it is intended that the cases be *stipulated* to be type 2-alt cases." She goes on:

> "If the non-identity cases are stipulated to be type 2-alt cases, then M* implies the results avoided earlier and thus excuses, under P*, the kind of conduct that I earlier conceded was wrong. However, if the non-identity cases are stipulated to be type 2-alt cases, then the non-identity problem becomes vulnerable on another front. As type 2-alt cases, the non-identity cases suppose that no matter how the child's parents conducted themselves they *could not* have brought the child in fact born a slave into existence as a non-slave... since the slave child case is stipulated to be a type 2-alt case, the non-slave child in the one world cannot be identical to the slave child in the other.
>
> "However implausible, this result is just a given if the non-identity cases are taken to be type 2-alt cases. But then it should be obvious that, if the non-identity cases are taken to be type 2-alt cases, these cases are highly artificial in a way their authors nowhere expressly recognize." (Roberts, 1998, p. 97)

Roberts is mistaken in answering the potential objection. It is not the case that NIPs are *stipulated* to be type 2-alt cases. It is the case that type 2-alt cases *are*. If NIP cases can be plausibly re-imagined as type 3-alt case, then it is not legitimately a non-identity problem. If it can't be re-imagined as a type 3-alt case (as in the example of the couple living in the slums of India, or the couple living in an already depleted world) then it is a non-identity problem. I believe that all

non-identity problems are type 2-alt cases, not just that they are stipulated to be type 2-alt. This is what constitutes a NIP, it must be the case that there is no accessible world *Y* such that the agents can bring about the numerically identical child into that world. And the way to determine whether a problem is a type 2-alt case, and therefore a legitimate non-identity problem, is to consider the two factors stated above: (i) do they control the environment they are bringing the child into, or is world Y accessible to them, and (ii) can the couple plausibly bring into existence the numerically identical child in world Y.

Roberts, after having spent so much time arguing that all NIP cases could be reimagined as type 3-alt cases (and claiming that cases that are presented as type 2-alt cases are "highly artificial"), concedes that there are cases that are strictly type 2-alt cases, and cites the fourteen-year-old girl case as an example. It does seem strange that Roberts would go to great lengths to argue that all non-identity cases can be reimagined as type 3-alt cases and hence that the perceived non-identity problem can be overcome, only to do a u-turn on the issue later in her book[8]. However, I think Roberts' distinction between type 2-alt cases and type 3-alt cases is an important tool in understanding when we are actually faced with a non-identity case, and when we are faced with real options, whereby a person's quality of life can and should be maximised. The principle M* contained in Roberts' personalism theory gives us a good guide for determining when someone has been wronged, even though it is not enough to overcome NIP.

[8] Roberts is also writing for the legal community in her book, and goes on to promote the argument of wrongful life. Her argument for wrongful life only stands when it can be said that *someone* has been wronged in a case of negligence. Thus to argue successfully for wrongful life, Roberts cannot have NIP hanging around to counter that claim. This explains why she spends a great deal of time denouncing NIP.

When Is Nonexistence A Better Option?

So far I have shown that when considering harms and benefits to future persons, if one is taking a person-affecting approach, one is faced with the non-identity problem. It seems to me that the non-identity problem cannot be overcome from a person-affecting position. We must accept that there are or will be some cases whereby a person is brought into being who is harmed (as defined by Harman, where harm is comparable to a "healthy bodily state") and although they have a harmed life, if it is at least somewhat worth living, then they cannot be said to have been wronged. This conclusion implies that there may also be cases where a person is brought into being whose life is so harmed that in fact it may not be worth living. So now I will move into a discussion of when we might say that nonexistence would have been better.

I will use Roberts' personalism as a guide for this discussion as I believe that she puts forward a solid account of how we might decide that a life is not worth living. In any case, the other arguments presented in this paper so far either don't address the issue or don't have room for the concept that some lives might be worth living, and in other cases nonexistence might have been the better option. Benatar, for example, thinks all lives are harmed and therefore nonexistence is better in every case; Harman, although not as explicit, argues along the same lines, while Bykvist rules out nonexistence as a comparable state altogether. Bykvist's stance on this is based on his actualist position: for a state to be better *for* someone, it must obtain for that person, since nonexistence cannot obtain for any person, nonexistence cannot be better (or worse) for them.

We can imagine cases where a person's life is so bad and full of suffering, that we would want to say it would have been better if they had not been brought into being.

Roberts' personalism theory recognises the following factors in wellbeing: happiness, pleasure, income, resources, autonomy, and capability (not necessarily exclusively) (1998, p. 138). Conversely, personalism recognises these as factors of ill being: unhappiness, pain, indebtedness, dependence and incapacity.

This gives us a good starting point for recognising when someone's life could be said to be in a state of harm. We now need to determine what level the level of harm should be before we can say that a life is not worth living. Note here that Roberts states that making the claim that a person would have been better off not having been born is not the same as making the claim that they should be euthanased (1998, p. 148). Benatar explains this difference well, he makes a distinction between a life worth starting (which he calls the future-sense of a life worth living) and a life worth continuing (the present-sense of a life worth living). Benatar claims that often the present-sense of a life worth living is applied to future-sense cases.

> "However, quite different standards apply in the two kinds of case. The judgement that an impairment is so bad that it makes life not worth continuing is usually made at a much higher threshold than the judgement that an impairment is sufficiently bad to make life not worth beginning." (2006, p. 22)

Down syndrome is an example of this. For someone that is already in existence, we would say that the condition of Down's syndrome does not make

life *not* worth living. However, if we were considering someone that was not yet in existence, we would say that it is a condition that affects quality of life in a way that it is not worth beginning (and hence the reason for testing for it so early in pregnancy).

To determine the level of harm needed to claim life is not worth starting, personalism uses a netting process: netting the wellbeing in someone's life against the ill-being to determine overall lifetime wellbeing. Roberts is careful here not to fall into an aggregative trap. She explains that personalism does not value aggregative efficiency as a primary value, rather it values efficiency as a non-aggregative value. By this she means that wellbeing must be maximised in every case. M* implies that when agents have maximised a person's wellbeing, then that person cannot be said to have been wronged. And if agents can, at no cost to others, maximise a person's wellbeing but don't, then D* implies that person has been wronged (1998, p. 140).

Further to the netting process, personalism allows that nonexistence has a neutral value, or a zero value. As Roberts has explained her theory in terms of accessible worlds, which are ranked in terms of how good they are relative to subject *s*, then we can understand subject *s*'s wellbeing as having a value of zero in any world in which she does not exist.

Following from this, personalism would say that a life would have been better not to have been started when it is harmed to the point that overall lifetime wellbeing is less than zero (taking into account surrounding agents having made every attempt to maximise *s*'s wellbeing). We should note here that Roberts' evaluation that a life was not worth starting is being made counter-factually, as Roberts' specifies that nonexistence is evaluated from *s*'s point of

view in that state of affairs in which *s* already exists. This is because Roberts is arguing for when someone can be said to have been wronged or harmed in a wrongful life case, and for that to be done, the subject needs to exist. For her purposes, she does not need to assign a value to s's nonexistence if s never "*not* exists". However, for the purpose of this paper, we do. In order to make a judgement on whether it is right to bring someone into the world, we need to be able to make a judgement on what their life might be like were they to exist and how that would compare to nonexistence. For example if we knew that two prospective parents both carried the gene for a serious defective disease such as Tay Sachs, and therefore a child of theirs would be born with the disease, we would say that life for that child, even though they don't yet exist, will be worse than nonexistence.

Thus personalism tells us the following in regards to whether nonexistence is a better option or not:

a) If the subject's lifetime wellbeing is below zero, then we can say nonexistence would have been a better option and they should not have been brought into being (and they can be said to have been wronged by being brought into being).
b) If the subject's wellbeing is above zero, even if only slightly, and their wellbeing has been maximised by the agents around them, then existence is a better option than nonexistence for the subject (and they cannot be said to have been wronged by being brought into existence).

Personalism suggests that the baseline for determining whether someone has been wronged by being brought into existence is their level of wellbeing once it has been maximised. Once the subject's wellbeing has been maximised, if their wellbeing is in the positive range, then they cannot be said to have been wronged.

> "If there is no accessible world in which the subject's wellbeing is higher, then the subject has not been wronged; but if there is such an accessible world, we must, in effect, justify the departure from it (by appealing to, e.g., considerations of fairness) or else be held accountable." (Roberts, 1998, p. 160)

Conclusion

From this chapter we have seen that in considering the wellbeing of future persons, and how our acts can affect future people, we are often faced with the non-identity problem. The non-identity problem suggests that there are some cases where a harmed life is the only life possible for that particular person, any attempt to avoid the harm would mean that that particular person would not have come into being. Following a review of some responses to the non-identity problem, I concluded that given a person-affecting approach the non-identity problem cannot be avoided, and thus we must accept that it is permissible to cause some degree of harm in creating someone, where there is no alternative in creating someone who is not similarly harmed. Further, in a review of personalism's approach to nonexistence, I showed that as long as a subject's wellbeing has been maximised, and their wellbeing is above zero, then

we cannot say that they have been wronged. This leads us to the conclusion that we do not wrong someone by bringing them into existence, even if they have a "harmed" existence, as long as their wellbeing is being maximised and their life is somewhat worth living.

3. The Repugnant Conclusion

In the previous chapter I concluded that, in regards to bringing a person into the world, we do not wrong a person if:

a) Their life has been maximised, i.e., that there is no accessible world where their wellbeing could have been higher, and
b) That their wellbeing is above zero, that is that their life is considered at least somewhat worth living.

Conversely: a person is wronged if their wellbeing has not been maximised and/or their wellbeing is below zero (that is their life is not worth living). This has provided a baseline for determining when someone has been wronged in regards to existence. This leads to the second question this paper attempts to answer: if we have determined that we do not wrong someone by bringing them into being as long as their lives are at least somewhat worth living, does this mean we can bring many people into existence as long as their lives are somewhat worth living? Can the population be extremely large and general wellbeing low? Or is a smaller population, with higher wellbeing better?

Many discussions in population ethics revolve around population evaluation and ordering populations by their goodness, using relations such as "better than" or "as good as". The most famous in this area is Parfit's discussion of such evaluations and the resulting "Repugnant Conclusion"(Parfit, 1984). Parfit's discussion starts around the questions 'How many people should there be? Can there be too many people?' He compares two populations, population *A*

has a relatively high level of wellbeing for all individuals (where wellbeing can refer to level of happiness, quality of life or share of resources per person, or a combination of all three where they correlate), while population *B* has a slightly lower level of wellbeing for all individuals, but contains more people, all with lives well worth living. Which of these two populations is *better*? Can a decrease in wellbeing be outweighed by an increase in the number of people living?

Parfit answers the question with reference to an aggregative principle: "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."(1984, p. 387)

In light of this, and referring to Figure 1 below, we can conclude that population *B* is better than population *A*; although the wellbeing in *B* is slightly lower than that of *A*, there are more people living in *B* than *A*, and therefore more people enjoying the things in life that make life worth living. If we are evaluating the populations on aggregative wellbeing, *B* is better than *A* (and is therefore permissible).

Following from this *C* is also better than *B*, again even though *C* has a slightly lower level of wellbeing, there are more people living a life that is still well worthwhile and so, on aggregation, *C* is better then *B*. We can continue on, *D* is better than *C*, *E* is better than *D* and so on until we reach population *Z*, such that there is a vast number of people within a population whose lives are only barely worth living. The lives in *Z* contain very little wellbeing, but due to the extremely large population this is the population with the largest sum of happiness. Assuming transitivity of 'better than', *Z* is better than all the previous populations, and as such is not only permissible, but preferable.

fig 1

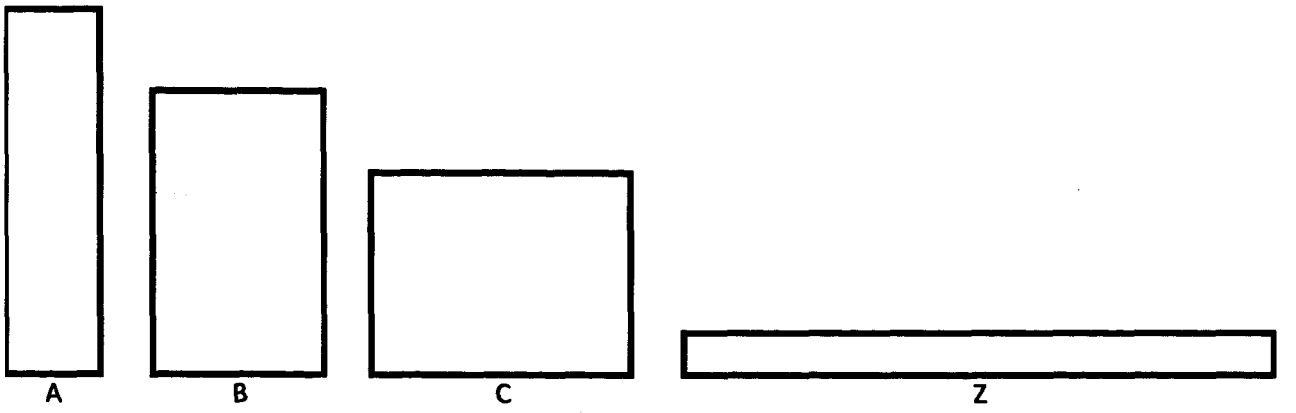

Parfit expresses it in this way:

> "For any possible population of at least ten billion people, all with a very high quality of life, there must be some much larger imaginable population whose existence, if other things are equal, would be better, even though its members have lives that are barely worth living." (1984, p. 388)

This is the conclusion that Parfit finds repugnant, as it seems counterintuitive that the best possible outcome would be the world that contained vast numbers of people all living lives that are only barely worth living. Yet if we evaluate populations based on how much of "the thing that makes life worth living" they contain, or how much wellbeing they contain, then we eventually arrive at the repugnant conclusion. Any morally relevant factor can be taken into account in the value of outcomes. It doesn't matter which values you choose to consider (wellbeing, happiness, quality of life, the amount of wine one is allowed to drink in a lifetime), in evaluating and ordering populations that contain these values any consequentialist theory will also arrive

at a version of the Repugnant Conclusion – it is better if more people get to enjoy some of what makes life worth living than less.

However, Parfit thinks that the Repugnant Conclusion is *intrinsically* repugnant, and that *all things considered A* is better than *B*, *B* is better than *C* and so on. On reflection, Parfit says the mere addition of extra lives does not make the outcome better: "If these lives are worth living, they have personal value. But the fact that such lives are lived does not make the outcome better" (1984, p. 412). However, neither can the addition of extra lives worth living, to an already existing population, make the outcome worse.

Consider the diagram below; *A* is a small population with a high level of wellbeing. *A+* is population *A* with the addition of a group of people who all have lives worth living, although with wellbeing below those in *A*, and who affect no one else. All things considered, we can say that *A+* is not worse than *A*.

fig 2.

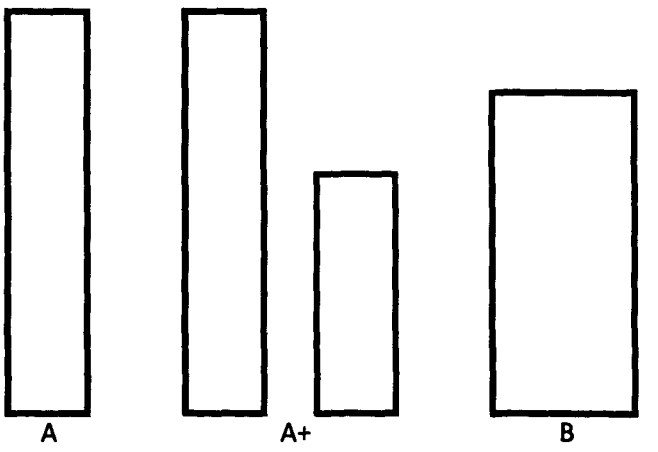

"Is *A+* worse than *A*? Note that I am not asking whether it is *better*. This we have already implicitly denied, since we have denied that extra lives...have intrinsic moral value. [Still] it seems harder...to believe that *A+* is *worse* than *A*. This implies that it would have been better if the extra group had never existed.

If their lives are worth living and they affect no-one else, why is it bad that these people are alive?" (Parfit, 1982)

Now look at *B*, we can say that *B* is better than *A+* because it involves the worse-off gaining more than the better-off lose; on egalitarian and utilitarian grounds, *B* is better than *A+*. So now we have plausibly claimed that *A+* is not worse than *A*, and that *B* is better than *A+*. These two beliefs would imply that *B* is not worse than *A* (since it is better than something else, *A+*, that is not worse than *A*). However this conclusion leads us in the direction of the Repugnant Conclusion, and we have already claimed that intuitively, we are inclined to believe that *A* is better than *B*. This is what is known as the Mere Addition Paradox.

In considering populations, number of people living and wellbeing, both the Repugnant Conclusion and the Mere Addition Paradox imply that a very large population, where wellbeing is only just at a level where life is worth living is permissible and even preferable. This directly answers the question posed at the start of this chapter. This also implies, contrary to popular belief, that there is no such thing as overpopulation (as long as wellbeing is above the neutral level), if we consider the term 'overpopulation' to have negative connotations and as referring to an undesirable state. But both proposals seem implausible and as their names suggest are repugnant or paradoxical.

Proposed Solutions To The Repugnant Conclusion And The Mere Addition Paradox

There are a number of proposed solutions to the Repugnant Conclusion and the Mere Addition Paradox, and there is a healthy literature around this topic. Here I will give a brief overview of a couple of the proposed solutions.[9]

Perfectionism And Discontinuity

Parfit proposes Perfectionism as a solution; Perfectionism is the idea that we should aim to enjoy or maximise the best things in life. Parfit considers what he calls "the best things in life", these are "the best kinds of creative activity and aesthetic experience, the best relationships between different people, and other things which do most to make life worth living." (Parfit, 1986, p. 161). Parfit argues that some of "the best things in life" are lost in each step taken in the Mere Addition Paradox. So, to use Parfit's example, the step from *A+* to *B* might see the loss of Mozart's music, the step from *B+* to *C* might see the loss of Haydn's. And so on, until finally in world *Z*, we are left with Muzak. And thus he argues, even if some change brings a great net benefit to those affected, the loss of some of the best things in life means it is a change for the worse (1986, p. 19). And when, he asks, should we make this claim? He argues that it would not be plausible to make those claims when we are considering the lives of those living

[9] I will note here one response to the Repugnant Conclusion and Mere Addition Paradox that I didn't feel warranted full explanation in the body of the paper, namely that they are just examples of Sorites paradoxes (Mulgan, 2004) and should be referred to metaphysicians for a solution. Briefly, in a Sorites paradox a series of steps is taken, each of which makes no great difference, though taken together they make a big difference. This argument has been rejected on the ground that the steps taken in the Repugnant Conclusion and Mere Addition Paradox do make a big difference and are relevant in many ways (S. Rachels, 2004).

in the *Z* world; for this claim to have any plausibility it must be made at the start, "[w]e must reject the change in which the music of Mozart is lost." (Parfit, 1986, p. 19) It is this argument that leads Parfit to the idea of Perfectionism in populations: that we are not only concerned about the value of each life within a given population, but also the kinds of experiences and activities that make life worth living. What we should be aiming for is a world or population that maximises the best things in life, and in the terms that Parfit has framed it, this leads us away from *Z* back towards *A*, where "the best things in life" are at their peak.

Griffin (1988) makes a similar argument in his proposal of discontinuity. Griffin argues that a single additive scale for measuring wellbeing is not plausible; that as we constantly reduce the level of wellbeing across a population, we might reach a point where "people's capacity to appreciate beauty, to form deep loving relationships, to accomplish something with their lives beyond just staying alive...all disappear." (Griffin, 1988, p. 340). The claim of discontinuity is that no amount of years lived at a quality of life, which merely involves the enjoyment of potatoes and muzak, or to use Griffin's example "kicks of kitsch" outranks some number of years lived at a quality of life, which involves the appreciation of beauty.

Both Parfit and Griffin attempt to counter the conclusion of the Repugnant Conclusion by showing that a loss in quality of life enjoyed by those living within a population cannot be outweighed by aggregative gains. Perfectionism is rejected by Rachels (2004), Ryberg (1996) and Qizilbash (2005) as implausible, and is even questioned by Parfit as a plausible solution. Parfit himself says that Perfectionism "conflicts with the preferences most us of would

have about our own futures" (Parfit, 1986, p. 20). By this he means that Perfectionism entails that one would prefer a future that consisted of a brief experience of ecstasy to a future that consisted of a long duration of pleasure that is only slightly less intense. It seems implausible. In regards to populations, and in particular the Mere Addition Paradox, Perfectionism implies that *A+* is better than *B*, even though in *B* the worst-off have had their wellbeing improved, and a great many more people can enjoy the pleasures of Haydn with the only loss being Mozart. This seems very implausible

Additionally, Rachels raises the problem of *Reverse Perfectionism*: "even if some change brings about a great net harm to those affected, it is a change for the better if it involves the loss of one of the worst things in life." (2004, p. 178) This would imply that the shortest duration of agony is worse than the longest duration of agony almost as bad. As Rachels puts it, this is nuts.

Thus neither Perfectionism nor discontinuity provide a solution to the Repugnant Conclusion and Mere Addition Paradox.

Accepting The Repugnant Conclusion

In contrast to Perfectionism, Tännsjö (2004) and Ryberg (2004) both suggest that we may actually already be living in a *Z* world. Although the *Z* world implies that lives are only just worth living, it is possible that the level of neutrality in our world is quite high. Tännsjö argues that when we consider the Repugnant Conclusion, we think that we are currently living in a state of high-level wellbeing, and that an enormous population of people living lives only barely worth living must look something like a huge concentration camp (2004,

p. 223). But in reality we may already be living lives that are only just worth living. Tännsjö argues that even if only our basic needs are satisfied, then most of us are capable of living lives that are worth experiencing. Even if we manage to acquire more goods into our lives, we rarely reach beyond the level of a life just worth living – a great improvement on our current lives is not really possible.

Broome (2004) takes a similar view. While Broome accepts that the Repugnant Conclusion is unintuitive, he argues that, understood in terms of goodness, it must be right, as everyone who exists in that population has a life that is worth living. Broome refers to a level of neutrality, which is the level of lifetime wellbeing that distinguishes when life is worth living, and when life is not worth living. If lifetime wellbeing is above the level of neutrality, then life is worth living. But, according to Broome, where the level of neutrality lies could affect our view of the Repugnant Conclusion: "If that was at the level of a mediocre life, the repugnant conclusion might be unattractive. But it need not be that level...It might be a lot better. It might be a reasonably good level of life." (Broome, 2004, p. 212) In this case, even is we are living just above the level of neutrality, we may have a reasonably good quality of life.

Tännsjö adds that there are moments in our lives where our quality of life might drop below the line of that worth living, such as suffering a major illness; there are likely even people living lives that are not worth living. If it is correct to say that the *Z*-world is a world where everybody lives roughly at our standard level of wellbeing, then it transpires that the repugnant conclusion is not repugnant after all (Tännsjö, 2004, p. 224).

Accepting the Repugnant Conclusion has its own problems, but we may also have to accept what Arrhenius calls the Sadistic Conclusion (2000, p. 251). Consider a reverse case or mirror image of the Repugnant Conclusion. In this case there are two similar but reverse populations, let's call them *RZ* and *RA* (Fig. 3).

fig 3

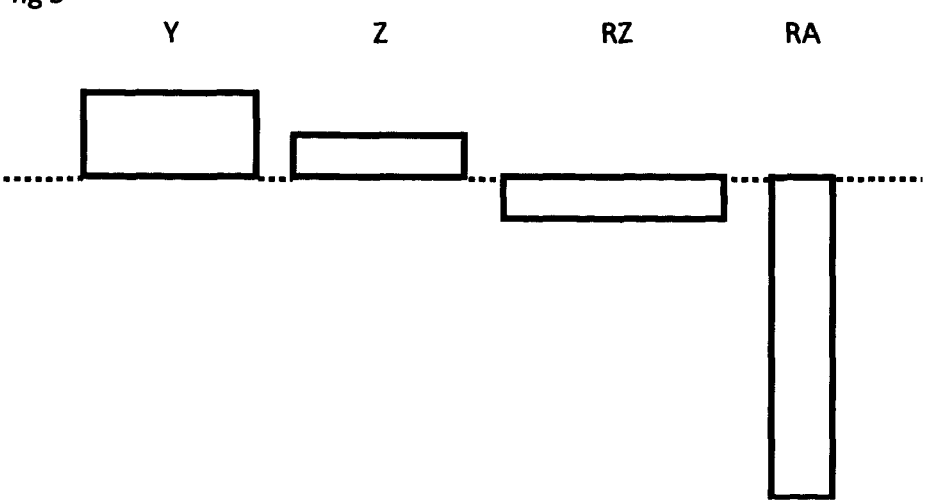

RZ is just below the level of neutrality (represented by the dotted line), we can imagine that this world is like ours, but where each person suffers a small amount of agony throughout their life; *RA* has fewer people but their average level of suffering is higher. An aggregative principle here would suggest that *RA* is better than *RZ*, because although the suffering is greater, there are fewer people to suffer it; the aggregation of suffering in *RZ* is greater and thus it would be worse if more people lived with less suffering. Further, for *any* population like *RZ*, there must exist a population like *RA*, in which average suffering is higher across a smaller population, which is better.

This seems wrong in itself, but Bykvist (2007b, p. 99) makes a further point referencing Broome's view point. If it is the case, as Broome suggests, that the level of neutrality is reasonably high, so that even if we are living just above the level of neutrality we have a reasonably good quality of life, then the reverse Repugnant Conclusion becomes even more bizarre. Consider Fig 4. below.

fig 4.

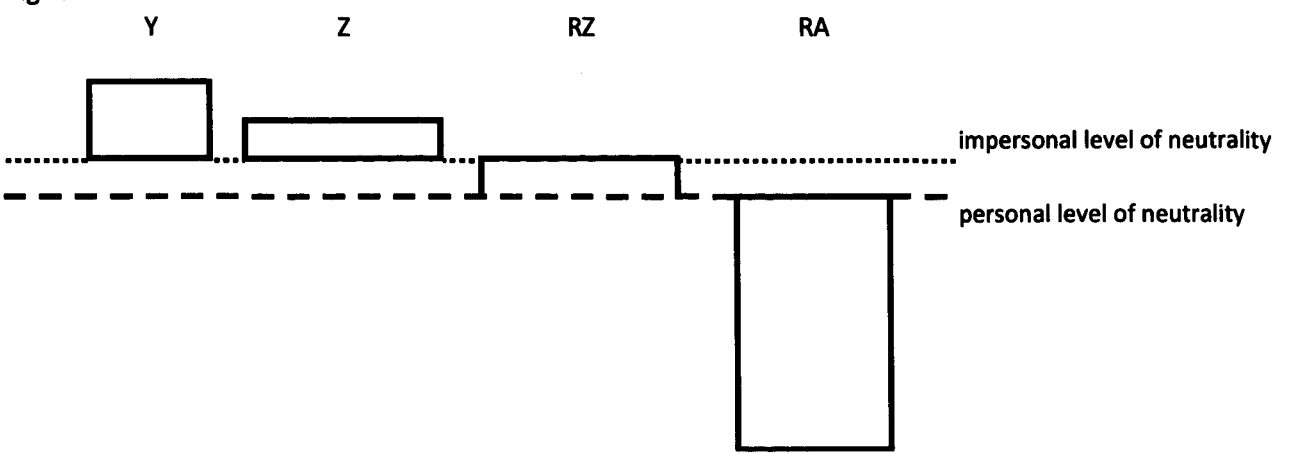

Z is above the level of neutrality, which for Broome is an impersonal level,[10] represented by the dotted line. Below that, represented by the dashed line, is the level where lives become truly miserable (this may be considered the personal level of neutrality, where below that line life may not be worth living on a personal level). We now have a case where, using an aggregative principle, a small population, *RA*, below the personal level of neutrality, where lives are not worth living, is a better outcome than a large population, *RZ*, above the personal level of neutrality, but below the impersonal level of neutrality. These are lives that are still worth living, but because more people are experiencing a quality of life below the level of impersonal neutrality, on aggregation it is not the

[10] For Broome, the level of neutrality is not the personal level of neutrality, but rather the level where creating that life or not makes no difference to the value of the population. Thus it is possible for the level of neutrality to be above the personal level of neutrality, where below that level, life might not be worth living *for that person.*

preferred outcome. Creating a level of neutrality that is quite high, at least higher than the personal level of neutrality makes this conclusion even more absurd.

Thus, pushing aside our intuitions and accepting the Repugnant Conclusion presents further problems such as the sadistic conclusion in the reverse example of the Repugnant Conclusion, and further the possibility that using an aggregative principle implies that a small and miserable population would be preferable to a larger population living lives that are somewhat worth living (in Broome's example). Both of these conclusions seem equally as unintuitive as the Repugnant Conclusion did in the first place.

The Question Of Transitivity

There exists a debate on the plausibility of denying the transitivity of the 'better than' relation. Denying transitivity means that while *B* is better than *A*, and *C* is better than *B*, in some cases it does not follow that *C* is also better than *A*. Denying the transitivity of 'better than' means that, in regards to both the Repugnant Conclusion and the Mere Addition Paradox, we don't necessarily accept the conclusion that *Z* is better than *A*. If we deny transitivity in these cases, we can accept that *B* is better than *A*, *C* is better than *B*, *D* is better than *C*, and so on down to *Z*, but argue that it doesn't follow that *Z* is also better than *A*.

Rachels (2004) proposes that there are two reasons why transitivity isn't necessarily true:

i) the factors that govern how *A* and *B* compare might not be identical to the factors that govern how *A* and *B* compare to *C;*

ii) the factors that govern how *A* and *B* compare may differ in significance when compared to *C*.

Rachels thinks transitivity fails when the factors involved in the three comparisons differ, or differ in significance.

Temkin argues that often we have disparate categories that are subject to different standards of evaluation, and the thing we are trying to evaluate might be a member of more than one category (2012, p. 173). To use the example that Temkin gives, John might belong to the categories of human, husband, father, philosopher and tennis player. The factors that are relevant and significant for evaluating husbands are different from those for evaluating tennis players. So it's possible that John might be a great husband but poor tennis player. Temkin points out that if we know that Tim is a better tennis player than John, and that John is a better husband than Ted, what does this tell us about how Tim and Ted compare? Nothing! This is because the factors that are relevant and significant in comparing John and Tim and John and Ted are different (Temkin, 2012, p. 173) And thus we can't say that Tim is better than John, and John is better than Ted, therefore Tim is better than Ted. While Rachels claims that in a case like this transitivity fails, Temkin argues that it is not a failure of transitivity, rather it is a case where transitivity simply doesn't apply (2012, p. 174).

Returning to the case of the Repugnant Conclusion, where in its simplest form we are comparing levels of wellbeing and numbers of people, it could be argued that the factors don't differ greatly between the alternatives presented. In comparing *A* to *B*, and *B* to *C*, *C* to *D* and so on, it is pleasure intensity that is relevant; as it is in comparing *A* to *Z*. However, as Rachels argues (2004, p. 180) its significance seems dramatically greater in comparing *A* to *Z*. While we can

accept that a small difference in pleasure intensity can be outweighed by a greater amount of pleasure (or a greater number of people experiencing pleasure) in the comparison of *A* to *B*, *B* to *C* and so on, the difference in pleasure between *A* and *Z* is significantly larger, and so can't be outweighed in the same way by the number of people experiencing it. Therefore *Z* is not better than *A*. This is an example of reason (ii) given above for the denial of transitivity. On Rachels' proposal, we must accept that pleasures can differ lexically, that ecstasy is lexically better than the pleasures of muzak and potatoes for example. Arguing for lexicality follows the same reasoning as the argument for discontinuity, but rather than arguing that discontinuity counteracts the aggregation of wellbeing in the Repugnant Conclusion, Rachels uses this to argue for the rejection of the transitivity of 'better than' (2004, pp. 181-182).

In one of his earlier papers *Intransitivity and the Mere Addition Paradox* (1987),[11] Temkin follows a different route in denying transitivity, he develops the notion of complete and incomplete concepts and argues that only complete concepts (concepts that allow full comparability) are fully transitive;[12] terms such as "at least as good as" and "not worse than" are only valid for complete concepts.

---

[11] I recognise that this is not Temkin's most recent work on intransitivity of the "better than" relation. Indeed he has written a number of papers and books addressing the issue, and has changed his approach to a degree over the years. I have chosen to reference this paper in particular because it shows an example of how an argument against intransitivity might take shape, and it is specifically directed at the Mere Addition Paradox. In his most recent work, *Rethinking the Good* (2012, p. 9) Temkin states that he no longer claims to have arguments for the conclusion that "better than" is not transitive. And although he claims he is now not willing to deny the transitivity of "better than", he still holds that the transitivity of "better than", while deeply compelling, is incompatible with a number of other views that are also deeply compelling. He is also not willing to give up any of these other views that are inconsistent with each other (2012, p. 10).

[12] Temkin recognises that it is usually *relations* that are transitive, but here consciously refers to concepts being transitive

To explain this, Temkin sets out the following inference scheme regarding the "better than" relation (1987, pp. 144-145):

(1) Given any concept c, for all x, y and z to which c is appropriately applied, if x is c-better than y, and y is c-better than z, then x is c-better than z.

(2) Given any concept c, for all x, y and z to which c is appropriately applied, if x is as c-good as, or c-equivalent to y, then however x c-compares to z, that is how y c-compares to z.

(3) If x is at least as c-good as y, and y is at least as c-good as z, then x is at least as c-good as z.

(4) If x is c-better than y, and y is at least as c-good as z, then x is c-better than z.

(5) If x is c-better than y, and y is not c-worse than z, then x is not c-worse than z.

(6) Given any concept c, for all x, y and z to which c is appropriately applied, if x is not c-worse than y, and y is not c-worse than z, then x is not c-worse than z and

(7) Given any concept c, for all x, y and z to which c is appropriately applied, if x is not c-worse than z, then x is c-better than z.

Temkin argues that (1) is the belief that "better than" is always transitive, and (2) he calls the principle of substitution (which he equates to the principle of substitution in algebra). (1) and (2) entail (3), (4) and (5). Temkin calls belief in (3) the belief that "at least as good as" is always transitive. He then argues that equating "at least as good as" with "not worse than" means that (3) and (4) then entail (6) and (7). But it is principles (6) and (7) which are only valid for complete concepts, concepts that allow for full comparability. But, he argues,

many concepts are not complete and therefore do not allow for full comparability, or are not fully transitive. Incomplete concepts only allow partial comparability, they only allow us to make some comparisons. "The inability to make precise comparisons in all cases need not be due to ignorance, but rather to the roughness or complexity intrinsically involved in that notion"(1987, p. 145). He argues further that incompleteness may be an intrinsic feature of many notions, and that once we recognise that certain notions are incomplete we can see that "not worse than" can be used such that it is *not* a transitive relation.

An example of a concept that is complete and hence, in Temkin's terminology, fully transitive is height, principles (1) to (7) are all valid for the concept of height. To illustrate a concept that is not complete, and therefore not fully transitive, Temkin uses Parfit's example of "great philosopher". The concept is not such that we can say of Plato and Aristotle that, as philosophers, one is greater than the other, or even that they are exactly equally as great. But we do sometimes think that some philosophers are greater than others, by more or less. Of this example, Temkin says: "It allows partial comparability in the sense that it enables us to make some comparisons between philosophers as to which is great but not others"(1987, p. 145). This inability to make a precise comparison is due to the "roughness" or intrinsic complexity of the concept. In cases like these, Parfit (1982) suggests that we can say of Plato and Aristotle that one is not worse than the other, without meaning that each is exactly as great as the other. Further, we might say that Plato is not worse than Aristotle, Aristotle not worse than Kant, and Kant not worse than Descartes, but we might think Plato *is* worse than Descartes. "...it might be that our notion of great philosophers permits precision in the latter comparisons, but not in the former

ones" (Temkin, 1987, p. 146). Thus we have an example of a concept that is "rough" or not fully transitive, whereby the application of the term 'not worse than' is valid, though not transitive.

Temkin goes further to distinguish concepts that are fully transitive, partially transitive and deeply intransitive. Referring back to the inference scheme set out by Temkin, a concept is fully transitive if principles (1) to (7) are valid for that concept; a concept is only partially transitive if some of principles (1) to (7) are valid, but others are not; a concept is deeply intransitive if one or more of principles (1) to (5) are not valid.

Returning now to the example used in the Mere Addition Paradox, Temkin argues not only that the concept of inequality is not fully transitive, but that it is deeply intransitive. This is based on Temkin's assertion that inequality is *essentially pairwise comparative,* meaning that assessing the moral implications of equality or inequality between populations can only be done on a pair-by-pair basis. "It appears then that on EC [the concept that equality is comparative] the relevant and significant factors for comparing *A* and *A+* regarding inequality differ from those for comparing *A+* and *B* in a sense connected to inequality being *essentially pairwise comparative*"(Temkin, 1987, p. 149). And therefore, in regards to the concept of equality, the principle of substitution for equivalence must be rejected, and hence the relations "not worse than", "better than" and "at least as good as" should not be considered transitive.

Temkin goes on to show, in a similar fashion, that Maximin and Utility are likewise deeply intransitive. This provides his argument the foundation to show that Mere Addition Paradox is only paradoxical if we are firm in our belief that preferability is transitive. But having shown that the concepts of equality, utility

and maximin are deeply intransitive, he argues that the following statements are *not* in fact inconsistent:

(1) *A* is better than *B*, and

(2) *B* is better than *A*+, and

(3) *A*+ is not worse than *A*

According to Temkin the reason that Mere Addition Paradox appears to be a paradox is that statements (1)-(3) all appear to be plausible, yet they contradict the belief that preferability is transitive. Temkin's solution to Mere Addition Paradox is to accept that preferability is *not* transitive; Temkin argues that because equality, maximin and utility are all deeply intransitive, this supports his claim that statements (1)-(3) are not inconsistent and thus rejects the transitivity of preferability.

The claim that preferability is not transitive is a serious claim indeed. Temkin himself recognises that it will attract criticism as it threatens not only large parts of morality, but also practical reasoning on a general level, as our concept of rationality is tied to the notion that preferability must be transitive.

Broome (2004) argues that the transitivity of betterness is not an issue in ethics, rather it is an issue in semantics, viz., that 'better than' is the dyadic predicate of 'good'; it is a comparative relation (that one thing has more 'goodness' than the other) between two options (or in Broome's case, distributions), and that a comparative relation is necessarily transitive.

He suggests that the relation 'judged *by me* to be at least as good as' may be intransitive, since it is not the comparative of any predicate.

"When I make judgements of betterness, different criteria may come to my mind as I make different comparisons. The results may be intransitive

judgements...But because 'better than' is the comparative of the predicate 'good', betterness must be transitive. Therefore, my judgements of betterness cannot possibly be *correct* unless they are transitive...So long as my judgements are intransitive, I know I still have work to do in sorting them out" (Broome, 2004, p. 61).

Vagueness And Incommensurability

Qizilbash (2005) argues that what appears to be intransitivity is actually vagueness of the predicates involved; that the predicates 'good' or 'awful' or other examples used in similar cases such as 'serious (illness)' are vague, the concepts have rough borderlines. The vagueness of the predicate means it is hard to tell when it is that a situation is good, or following a slight change in circumstances, has now become 'not good'. In regards to the Repugnant Conclusion, Qizilbash argues:

"...it is plausible that there is no precise transition from levels of living which constitute a satisfactory life to those which do not...one can argue that the borderline between lives which are, and are not satisfactory is imprecise. There is, rather, a range of levels of wellbeing which neither definitely constitute satisfactory lives, nor definitely constitute non-satisfactory lives. These fall into the vague zone between those lives which are definitely satisfactory and those which are definitely not satisfactory" (2005, p. 128).

Handfield (forthcoming) follows a similar line of argument in proposing that goods can be incommensurate; he calls vagueness a special case of incommensurability. Handfield defines incommensurability: "two goods are

incommensurate (or incomparable) if and only if it is not the case that one is better than the other, nor is it the case that they are equal in value" (forthcoming, p. 14. n 16). He argues that when two goods are incommensurate, a rational agent might struggle to choose which one is better. This struggle to choose which of two options one prefers is a more acceptable way of explaining the apparent contradiction, or paradox, in the premises of cases such as the Repugnant Conclusion, rather than rejecting transitivity, a vital component of practical reasoning and our fundamental moral framework.

While Temkin and Rachels feel that the denial of transitivity is a suitable solution for the problems posed by the Repugnant Conclusion (and other similar spectrums with varying degrees of pleasure intensity and duration), denying transitivity seems to have other dire consequences for practical reasoning and making moral judgements. There are a number of arguments against Temkin and Rachels' denial of transitivity, and while they take slightly different approaches, the arguments tend to centre around the idea that either the relative qualitative predicates are vague, or that as human agents we don't have the ability to make sound judgements and choose correctly between the available options. However, this does not mean that we should reject transitivity altogether, we would be better to reject one of the other premises in the paradox than to threaten our moral and rational framework.

Personalism And The Repugnant Conclusion

Roberts (2004) puts forward a distinctive way of thinking about the Repugnant Conclusion, and since I have spent quite a bit of time discussing her

theories in relation to the non-identity problem, it is worth considering in detail her response to the Repugnant Conclusion.

In establishing her theory of *personalism* (her person-based consequentialism), based on what she calls the person-affecting intuition, Roberts states that personalism denies both maximising forms of consequentialism (what she calls *totalism* and *averagism*) (1998, p. 6).

Further, Roberts is critical of the aggregative approach taken in the Repugnant Conclusion, or as she calls it a maximising consequentialism, because, by virtue of its focus on aggregative good, it fails to take into account "the obligations that agents have in respect of *persons as individuals* and in particular the obligations that agents have in some circumstances to create additional goods for such persons." (2004, p. 99)

Roberts argues that its failure to take into account persons as individuals, or the maximisation of personal wellbeing, is precisely what leads us to the Repugnant Conclusion.

In response, Roberts puts forward her person-based consequentialism (from here referred to as PBC, also known as *personalism*). Roberts argues that PBC is a consequentialist interpretation of the person-affecting intuition (the foundation of her overarching theory, as discussed in part I of this paper), and while it rejects traditional forms of aggregation, it does incorporate a maximising approach:

"According to PBC agents have the obligation, *for each existing and future person as individual,* to create whenever they can additional wellbeing *for that person.*" (2004, p. 99)

To support her claim that PBC is a maximising consequentialism, Roberts claims that it is founded on various maximising principles that determine when a given person has been wronged. It will be useful here to recap the basic principles of her theory. Referring to possible populations (worlds/futures/distributions) *X* and *Y*, the four principles are:

N* = *p* is not wronged by agents in *X if p* never exists in *X*.

M* = *p* is not wronged by agents in *X if*, for each world *Y* accessible to such agents, *p* has at least as much wellbeing in *X* as *p* has in *Y*.

D* = *p* is wronged by agents in *X if p* exists in *X* and there is some world *Y* accessible to such agents such that:

(iv) *p* has more wellbeing in *Y* than *p* has in *X*;

(v) for each person *p'* who ever exists in *X*, either *p'* has at least as much wellbeing in *Y* as *p'* has in *X* or *p'* never exists in *Y*; and

(vi) for each *p'* who ever exists in *Y*, *p'* exists at some time in *X*.

P* = *X* is permissible for agents at *t if and only if* no person who exists at or after *t* is wronged in *X* at or after *t*.

Roberts argues that these principles form a maximising consequentialism because they focus on maximising an individual's wellbeing by taking the following into consideration:

1. Whether *p*'s own level of wellbeing at *X* has been maximised, or whether there exists an alternative *Y* such that *p* has more

wellbeing at *Y* than at *X*; if no such alternative exists, *p* is not wronged at *X*

2. Whether additional wellbeing could have been created for *p* at *X* at no cost to others
3. Whether in the case where there exists an alternative *Y* such that *X* and *Y* contain the same number of people and *p* has more wellbeing at *Y* than *p* has at *X* but only at a cost to someone else *q* who exists (or will exist) at *Y*, *p*'s level of wellbeing at *X* is lower than *q*'s level of wellbeing at *Y*; depending on other facts, if it is lower, then *p* may be wronged at *X*
4. Whether in the case where there exists an alternative *Y* such that *X* and *Y* contain the same number of people and *p* has more wellbeing at *Y* than *p* has at *X* but only at a cost to someone else *q* who exists (or will exist) at *X*, and *p*'s level of wellbeing at *X* is equal to *q*'s wellbeing in *Y*, it is also the case that more people exist (or will exist) in *q*'s bad circumstance at *Y* than people who do or will exist in *p*'s bad circumstance at *X*; depending on other facts, *p* may not be wronged at *X*
5. Whether *p* does or will ever exist at *X*; if *p* never exists at *X*, *p* cannot be wronged at *X*

(Roberts, 2004, pp. 102-103).

Roberts believes that shifting the focus from aggregative wellbeing to individual wellbeing allows PBC to "de-emphasise" the significance of aggregative wellbeing, and thus sufficiently address the Repugnant Conclusion.

According to PBC, a particular future *X* will only be permissible if existing or future persons at *X* are treated in a morally permissible way, i.e., they are not wronged. If we recall some detail from Chapter 2 of this paper, according to Roberts, a person is wronged if their wellbeing has not been maximised, i.e., if there is an *accessible* world where their wellbeing would have been higher, then there is a duty to the agents around them (parents, govt, etc.) to bring that world about, otherwise we can say that that person has been wronged. This is the part of her theory that Roberts claims contains maximising elements. PBC does not aggregate wellbeing across given populations in order to determine permissibility, rather it focuses on the maximisation of personal wellbeing. PBC claims that agents ought to do or create the most good or wellbeing that they can for persons (2004, p. 100).

An important element of Roberts' overarching theory, and for her solution to the Repugnant Conclusion, is its consideration of possible and actual future people. According to PBC "those that could have existed according to a possible future, but do not and never will, do not count for moral purposes" (2004, p. 101). PBC rejects a "wide" person-affecting approach, considering only those that *do* and *will* exist (and considers persons in each group as counting equally).

Roberts regards the *Z* world in the Repugnant Conclusion as a situation of extreme overpopulation: *Z* is vastly overpopulated but according the traditional

aggregative approaches is deemed better than *A*, which comprises a significantly lower level of population.

Roberts attributes the 'repugnance' of the Repugnant Conclusion to the *intuition* that lives in *Z* have not been maximised; she argues that there are many, many lives in Z whose lives have been "gratuitously rendered only barely worth living" (2004, p. 105) when they might have had lives well worth living. In Roberts' view, the overpopulation in *Z* world means that resources are scarce and that this is the cause of the lowered level of wellbeing. Roberts argues that there is some alternative world where their lifetime wellbeing could have been higher, at least for some of them. According to principle 5 of PBC above, bringing about a population with fewer people i.e., actively not bringing people into existence, would not be wrong, as a person cannot be wronged by not being brought into existence. Maintaining a lower level of population correlates with an increase in the level of available per capita resources, thus improving the wellbeing of those in existence within the population.

Therefore, under Roberts' proposal, when comparing populations *Y* and *Z*, those who exist in *Y* who would have also existed in *Z* have had their wellbeing improved, while those that existed in *Z* but don't exist in *Y* have not been wronged by not being brought into existence, so *Y* is better than *Z*. We can make the same conclusions when comparing *Y* and *X*, *X* is better than *Y*; and the same again when comparing *X* and *W*, *W* is better than *X* and so on back up the scale of populations. In this way Roberts argues that *Z* is not better than *A*, and that reducing the level of wellbeing in favour of a larger population is not permissible.

"To avoid treating so many so badly, agents ought, we think simply to have chosen a more conservative population policy rather than overpopulation. That

way we think *A* - or some similarly small well-off population – in which each person's life is well worth living would have been produced rather than *Z*." (2004, p. 105)

Roberts' PBC solution to the repugnant conclusion focuses *not* on the fact that some people who exist at *Z* would not have existed at all had the overpopulation choice not been made, but rather the fact that at least some person *p* who exists and suffers at *Z* would also exist in a smaller, alternative population, say *M*, where their wellbeing is greatly increased. Roberts' person-based consequentialism tells us that we should maximise the lives of those that *will* exist if we are able to do so. According to Roberts, "there once existed some alternative future *M* that is the same size as *A*, contains persons who, as at *A*, "enjoy a very high quality of life" and is such that *p* exists at *M*" (2004, p. 107).

Roberts argues that after considering all alternative futures, including *M*, we can conclude that agents who choose *Z* as a permissible future population will be doing something wrong (2004, p. 108).

There are two prominent issues with Roberts' solution. The first is the issue of the non-identity problem. As covered in Chapter 2 of this paper, Roberts herself proposes that a person *p* cannot be wronged by being brought into existence, as long as their wellbeing is within the range of a life worth living. In the case of *Z* world, although the wellbeing of the lives within that population is significantly reduced compared to other alternatives, their lives are still worth living, and so living them in *Z*-conditions is *not worse for them* than not having been brought into existence at all. So according to Roberts' own theory, the people living in *Z* have not been wronged by being brought into existence.

However, in addressing the Repugnant Conclusion with her person-based consequentialism approach, she says that *p* in *Z* is wronged because their life hasn't been maximised. Roberts recognises this as a problem, and argues that PBC manages to "steer a middle course" between considering individual wellbeing by applying a person-based approach and having to consider the identities of persons in choosing or not choosing a particular population (2004, p. 106). This is possible she says because "PBC considers normatively relevant not the fact that people who exist at *Z* would not have existed at all had the overpopulation choice not been made but rather the fact that at least some person p who exists and suffers at *Z* is such that agents could have acted...to create additional wellbeing for p" (2004, p. 107). It's a very tight rope that she is walking here. On the one hand her personalism theory allows that if *p* is brought into existence in *X*, where *p's* existence is harmed or compromised, but there is no accessible world *Y* where *p* exists, and agents in *X* are maximising *p's* wellbeing, then *p* is not wronged by being brought into *X* (this is Roberts' response to the non-identity problem (1998)). In the case of the Repugnant Conclusion, she is arguing that there is an alternative accessible world, *M*, that *p* could have existed in, and so by being brought into existence into *Z*, where *p's* wellbeing has not been maximised, *p* has been wronged. However, this assumes that there is at least one person, *p*, who exists both in *Z* and in *M*. It is possible that *p* can only exist in Z, and that if M had been brought about, there would be no person *p* who could have existed in both *Z* and *M*. And so the viability of Roberts' PBC (personalism) theory rests on the assumption that there is at least one person *p* who could have existed in both *Z* and *M*. If this is not the case, then her argument against the Repugnant Conclusion is not sound.

The other problem with Roberts' solution is that it seems to suggest a reverse Repugnant Conclusion. If we follow the path we started above, where we reduce the population level in favour of an increase in wellbeing, we not only arrive at world *A* from *Z*, but we could go further. We could imagine a situation in which only one person ever lived who had an amazing quality of life. This would be the alternative where *p*'s life had been fully maximised This would be at no cost to others, as no one else would have be wronged by not being brought into existence. She in fact goes on to argue that there is no obligation to procreate and that even if this led to a state of extinction, it would not be wrong. From a consequentialist point of view, if humans were to go extinct in a way in which no one was harmed (i.e. the last generation of humans just decided they did not want to procreate), then this is not a bad thing, as it is not bad or harmful for anybody. However, it does seem counter-intuitive that extinction could be the *preferred* option of all available alternatives. While I do not intend to go into a lengthy discussion on the ethics of extinction, it is my inclination to believe that while we might intuitively favour a lower population level with higher wellbeing over a larger population with lower wellbeing, there would come a point where lowering the population further does not increase wellbeing any more. It might be critical to one's wellbeing that they at least have family and some friends around to maintain quality of life, or that co-habitation with others increases the pleasure of living as the responsibility of providing food and shelter is shared with others, or that one needs a succeeding generation to carry on one's goals and projects to completion for one to be truly happy, and so on. I do believe that in discussing preferences in populations, it seems just as counter-intuitive for

extinction to be considered preferable to a small population, as it does for Z population to be considered preferable to A population.

Furthermore, and of particular interest for this paper, Roberts' theory states that there is a population *M* that is preferable to *Z*, and that at some point *M* existed as a genuine alternative to agents (as opposed to *Z*), and that agents, in choosing *Z* when *M* was available have done something wrong. If *M* is a smaller population than *Z*, then this suggests that at some point, agents, whoever they may be, should have chosen to reduce, or slow down the rate of population growth. While Roberts goes on to consider whether there is an obligation to procreate (leading to her discussion on extinction), she does not consider whether there is an obligation to *not* procreate. Additionally, she does not state who the agents are that would be wrong in choosing *Z* over *M* (she may not have considered it within the scope). Does she believe governments have a duty to future people to limit population by putting into place population policies? Or is she referring to individual people, procreators, and suggesting that they should be limiting the number of children they each have? While she argues that we do not have an obligation to procreate, she does not give any prescription in regards to refraining, even though she has argued that *Z* world would be a state of overpopulation and that agents would be wrong in choosing that population.

Referring back to the discussions on Roberts' response to the non-identity problem, I would suggest that if the population level has passed *M*, then agents are in a situation representative of a type 2-alt case. That is, while population *M* may be a preferable alternative for the wellbeing of future people, *M* may no longer be accessible, and thus the only alternative for those wanting to bring

children into the world is to bring them into this world, let's say *X* population, or not to bring them into existence at all. And as long as p's wellbeing is maximised within *X* population, *p* has not been wronged, and thus bringing *p* into *X* population is permissible. Further, according to Roberts, future population *X* is permissible if those within *X* have had their wellbeing maximised. If we do consider this to be a type-2-alt scenario, then as long as agents maximise *p's* wellbeing in *X*, they are not wronged, and we can consider bringing about *X* permissible.

Conclusion

In reference now to the question posed at the start of this chapter, the Repugnant Conclusion and the Mere Addition Paradox would both imply that a very large population with a low level of wellbeing (but above the neutral level) is permissible, even preferable. However, many find these conclusions unintuitive, and as we have seen, accepting these conclusions presents its own problems. But, neither do any of the proposed solutions seem satisfactory. Additionally, most of these conclusions rest on an aggregative approach, while in this paper I have focused on a person-affecting approach. Roberts attempts to avoid both the Repugnant Conclusion and Mere Addition Paradox, by finding some middle ground with her PBC theory. She concludes that *Z* represents overpopulation and that in choosing this population agents would be wrong. However, I argue that her PBC theory, while it may appear to counter the Repugnant Conclusion, leads to other issues, namely the reverse Repugnant Conclusion. Additionally Roberts' attempt to answer the Repugnant Conclusion lies in contradiction to her own views on the non-identity problem. Thus I argue

that no response has satisfactorily resolved the Repugnant Conclusion. I do argue however, that as long as we maximise the lives of those living within *X* population, *X* is permissible.

4. Children and Special Goods

This paper has to this point been concerned with whether procreation is permissible, and I have examined closely the effect bringing a child into the world might have on the child itself and on the population as a whole. Ethics is concerned with how our actions affect others and from this perspective can give us guidance in our actions and behaviour. Having determined in the previous two chapters that the potential child cannot be harmed by being brought into existence, and that as long as lives within *X* population are worth living and lifetime wellbeing has been maximised for each person living in *X*, I feel that I have established that, under normal circumstances, considering the effect on others, it is permissible to bring a child into the world.

Now I want to change the focus – what about the parent-to-be? I want to examine how becoming a parent can hold its own value for the prospective parent. My intention in this chapter is to show that in addition to the previous arguments, not only is it not wrong to bring a child into the world, but that it can bring about many goods. This may not show that it is ever obligatory to have a child, but it does show that it is at least sometimes permissible.

As opposed to examining the effects one's actions can have on others, I am examining here the effects one's actions can have on oneself. I will start by considering the importance of relationships and will argue that the parent-child relationship provides special goods. I will also look at how having children can be an extension of identity and play an important part in living a fully flourishing life.

Special Goods And The Parent-Child Relationship

Intimate relationships are an important part of human life. We form intimate relationships with various others in our lives: with our parents, our children, our siblings, our partners and our close friends. An intimate relationship involves caring for someone and having that care reciprocated. Intimate and loving relationships are personal goods of great importance: to love other people and be loved in return is part of what is involved in having a rich and satisfying human life. Further, to be loved is to have one's own value affirmed, it is a source of self-esteem (J. Rachels, 1997, p. 223). Rachels and Schoeman (1980) both argue that intimate relationships give meaning to our lives, that loving relationships provide us with things to value and with a sense of being valued.

> "For most people, not only are such unions central for defining who one is, but human existence would have little or no meaning if cut off from all possibility of maintaining or re-establishing such relationships." (Schoeman, 1980, p. 14)

We develop intimate relationships with numerous people in our lives, members of our family, close friends and lovers. These relationships are all capable of providing value in our lives, and give meaning in the way that Rachels and Schoeman suggest. Brighouse and Swift (2014) argue that the different possible relationships are not substitutable for each other. They argue that we need more than one type of relationship in our lives: romantic relationships allow us to express our raw emotions and share ourselves sexually; we need

close friendships, those that we can confide in and go to for support, and whom we can offer support in return; we also need less intimate relationships, relationships based on mutual recognition and respect, such as those in the workplace or other social settings. Brighouse and Swift argue that a successful life is a life with a variety of successful relationships, including intimate relationships (2014, p. 14).

Keller (2007) argues that the goods that exist in a parent-child relationship are special goods: those personal goods that can only be given to the parent from the child, or from the parent to the child (2007, pp. 124-125). He distinguishes special goods from generic goods, which are those personal goods that could be received from people that aren't a parent or child. Generic goods can arise from multiple relationship types: a close friend providing practical assistance when needed, the resulting success from working collaboratively with a colleague at work, the happiness one feels in catching up with a friend and sharing funny stories. Goods of this type can be gained from many or all of the different relationships we form with other people. Special goods on the other hand are goods that can only be gained from the parent-child relationship, such as the joy one might feel in playing a part in the development of a person from birth through to adulthood, of being responsible for protecting the child's interests and meeting their needs, of being directly involved in the flourishing of their lives and encouraging them to become autonomous adults.

Brighouse and Swift (2014) also argue that the parent-child relationship provides personal goods that cannot be gotten anywhere else except from the parent-child relationship. They propose four features of the parent-child

relationship that set it apart from other intimate relationships: i) parents have power over children that is not reciprocated; (ii) children do not have the power to exit the relationship; (iii) the quality of intimacy is different as children love their parents in a spontaneous and unconditional way; and (iv) the moral quality is different as the parent has fiduciary responsibilities.

While (i) and (ii) outline the shape of the relationship, (iii) and (iv) help to explain the value of the relationship to the parents as well as the children. Brighouse and Swift claim that it is the quality of the intimacy in a parent-child relationship that contributes to the flourishing of the parent's life and provides value for the parents.

> "The love one receives from one's children, again especially in the early years, is spontaneous and unconditional and, in particular, outside the rational control of the child. She shares herself unselfconsciously with the parent, revealing her enthusiasms and aversions, fears and anxieties, in an uncontrolled manner. She trusts the parent until the parent betrays the trust, and her trust must be betrayed consistently and frequently before it will be completely undermined. Adults do not share themselves with each other in this way: intimacy requires a considerable act of will on the part of adults interacting together." (Brighouse & Swift, 2014, p. 17)

Receiving love in a relationship has moral significance, but the spontaneous and unconditional love a parent receives from a child holds additional value that cannot be gained from other relationships. In contrast to what is typically seen as a vertical relationship, where children are seen as the

receivers of love and parents seen as the givers, Mullin argues that children *also* give love and that in itself has great value for the parents receiving it. The love received from one's child is another good that is special to the parent-child relationship.

> "When we consider the impact of loving relationships on parents, we should consider not only the value of developing virtues associated with giving love and providing care for another, but also the value of being loved by children, and the value of having children contribute to meeting their needs." (Mullin, 2007, pp. 390-391)

Additionally, it is the fiduciary aspect of the parental role that Brighouse and Swift argue holds particular value for parents. The fiduciary role that the parent plays contributes to a parent's flourishing in a way that other intimate relationships don't: "The parent is charged with responsibility of both the immediate wellbeing of the child and the development of the child's capabilities" (Brighouse & Swift, 2006, p. 94). The fiduciary obligations include guaranteeing the child's wellbeing and overseeing her cognitive, emotional, physical and moral development. Brighouse and Swift argue that it is not only the child that has an interest in the fiduciary aspect of the parental role, but that the parents do also. They argue that playing the fiduciary role of parent enhances one's life experiences, and as such there is a non-fiduciary interest for the [prospective] parents in playing that role.

> "Parents have an interest in being in a relationship of this sort... The role enables them to exercise and develop capacities the development and exercise of which are, for many (though not, certainly, for all), crucial to

> their living fully flourishing lives. Through exercising these capacities in the specific context of the intimately loving parent-child relationship, a parent comes to learn more about herself, she comes to develop as a person, and she derives satisfactions that otherwise would be unavailable. The successful exercise of this role contributes to, and its unsuccessful exercise detracts from, the success of her own life as a whole." (Brighouse & Swift, 2006, p. 95)

The parent-child relationship is one instance of a range of intimate relationships available within a life, which are necessary for a successful life. However, the parent-child relationship provides additional goods that cannot be gained from other intimate relationships. These kinds of goods, these special goods are goods that can only be given to a parent from a child.[13]

Parental Identity

We have seen now the value that becoming a parent can hold for an individual. While intimate relationships are important for a flourishing life, the parent-child relationship can hold a special value, and provide special goods to those involved, namely the parent and the child. But, as Keller (2007) points out, the desire to become a parent is not a straightforward altruistic desire to help another, or a desire to ensure care in later life or even a desire to do a moral or political duty. Rather the desire to become a parent is more likely to be a desire

[13] I believe that special goods can exist in an adoptive parent-child relationship; the parent and child do not have to be biologically related for special goods to exist. I will visit this argument in more detail further into the chapter.

about the shape of your own life, the desire that your plan of life includes bringing up a child and all the experiences that go along with that (Keller, 2007, p. 122) Our identity gives structure to our life (Appiah, 2005, p. 16); not only does it provide us with a vision for how our life will pan out, but it helps us to choose from the many options for action presented to us within our lives. One such option that nearly all of us will face is whether to parent children or not. For some people, becoming a parent is a part of their identity, of how they see their lives playing out.

While playing the role of parent can form a part of one's life plan, children can also contribute to one's identity in a different way, by becoming a continuation of the parent's identity. Reshef (2013) endorses Brighouse and Swift's arguments for the parent-centered value in parenting. However, he argues that it is not the fiduciary role that is distinctly valuable for parents. In addition to the four features Brighouse and Swift say are specific to the parent-child relationship, Reshef argues there is one further element to the relationship that gives the relationship value for the parents and provides justification for the family from the parent's perspective. He calls this aspect of the parent-child relationship the *element of identity*. Reshef defines the element of identity as: "a strong sense of interconnectedness and continuity between the parent's and child's identities that is established during childhood by a process of reproducing some of the parent's characteristics in the child"(2013, p. 132).

Some of these characteristics are biological, and Reshef considers those characteristics important in the continuing development of the parent's identity, that you see some of yourself, literally, in your child.

But, I believe more importantly, he is also talking of less tangible characteristics that you can pass on to your child: your beliefs, ideals, traditions, and so on. In passing on these characteristics to your child you create a continuity of your own identity, that lives on after you have died. Children carry on your beliefs, traditions and ideals into the next generation. This interconnectedness between the parent's identity and the child's identity is what Reshef argues is the distinctive value of a parent-child relationship.

> "Through the intimate process of upbringing, parents can bequeath their cultural, national, and religious horizons to their children. Children acquire their parents' language, they are raised according to their parents' values and beliefs, and they follow their parents' practices. Some of the parents' more personal characteristics also pass on to their children, such as favourite dishes, leisure activities, hobbies, body language and outward 'look'."(2013, pp. 140-141)

The parent's identity helps to shape the child's identity, and the development of the child's identity creates a continuation of the parent's identity. Reshef claims that the element of identity explains why the parent-child relationship makes a distinctive contribution to the flourishing of adults (2013, p. 141).

What About Adoption?

One might ask whether adoptive parent-child relationships can also provide special goods or feature the element of identity. If it is the case, then the argument may follow that we have a duty to adopt an already existing child who needs care, rather than bringing a new child into the world.

I believe that special goods, as described above, can be received from an adopted child. Special goods are gained from a particular type of relationship (parent-child relationship) and are not dependent on biological or genetic links. The parent-child relationship contains special goods because of the long-term, ongoing special care involved in that relationship. All the goods that Keller describes as being special to a parent-child relationship can be gained without, for example, biological links, as being responsible for and influencing the child's development and having someone who is especially committed to meeting your needs and protecting your interests. Keller does sometimes refer specifically to genetic links as enhancing special goods (2007, p. 125), for example, a birth parent may have traits in common with the child, and as such see a younger version of herself in the child. Or from the child's perspective, understanding who your birth parents are can enhance your understanding of yourself (it may explain certain characteristics about yourself). However, Keller doesn't rule out the adoptive parent-child relationship as one that cannot provide special goods. An adoptive parent can still identify with and have a special understanding of the child, can take great enjoyment in being involved in the child's development and can be especially committed to protecting the child's interests. The adopted child can in return provide special goods to the parent.

Some argue that the adoptive relationship provides further goods that cannot be gained from a biological or genetic relationship: that the relationship is defined solely through a history of love and care, rather than through biological connections (Rulli, 2014).

In both cases it is the nature of the parent-child relationship that gives rise to special goods, and as such these goods can be given and received in an adoptive relationship.

Some argue that not only does adoption have moral value, but also that adoption is superior to procreation (Rulli, 2014, p. 10). Indeed if we accept that special goods can be gained from an adoptive relationship, then we may have no special reason to have our own biological children. All the goods that can be gained from procreating can also be gained from adopting a child, *as well* additional goods can be gained. The arguments for choosing adoption over procreation include helping an already existing child in need of a family, meeting existing needs rather than creating new ones, and the moral good of extending help, love and care to someone whom one is not biologically or genetically linked. I am not intending in this paper to make the claim that procreation is equivalent to, better than or worse than adoption. I agree that special goods can be given and received in an adoptive parent-child relationship, and also that biological and non-biological parent-child relationships may hold goods that the other doesn't. I don't intend to weigh those goods up against each other here. Whether adoption is *better* than procreation, or whether we have a *duty* to adopt are both complex questions and require detailed examination. One might argue that we have a duty to adopt every child in need, and maximise their wellbeing.

Taking an aggregative approach could lead to an argument that we must adopt children in need of a family *as well as* bringing into existence as many "happy" children as possible to increase the average level of wellbeing within a population. I do not feel that this paper's scope extends to that part of reproductive ethics. Rather this paper is concerned with whether procreation is permissible. From a person-affecting view, as long as the future person's lifetime wellbeing is maximised, and their life is somewhat worth living, then they are not wronged by being brought into existence. I argue that it follows from this, that under normal circumstances, it is *prima facie* permissible to bring a child into being.

Conclusion

I have argued here that, in addition to not harming anyone, bringing a child into the world holds its own goods and values for the parent. Becoming a parent is, for many people, part of living a fully flourishing life.

The parent-child relationship provides special goods for both and holds value on a personal and social level. This provides good reason *for* people to pursue becoming a parent if that is what they wish to do.

5. Conclusion

The main aim of this paper has been to explore the ethics of procreation with a specific focus on population. I have been interested in whether it is permissible to bring a child into the world and have discussed when we can say that a person has *not* been wronged by being brought into existence. I have shown that we do not wrong someone by bringing them into existence if we maximise their wellbeing and that their life is somewhat worth living. This would imply then that procreation under normal circumstances is permissible.

Second, I explored the issue of population and wellbeing. It is commonly assumed that an increase in population leads to a decrease in wellbeing and as such we should aim to at least not increase population. Following an in depth discussion of the repugnant conclusion and some attempts to answer it, I concluded that a satisfactory response to the Repugnant Conclusion has not been found. However, I argue that as long as those living within population *X* have their wellbeing maximised, they have not been wronged and *X* is permissible.

In addition to showing that it is permissible to bring a child into the world, I also showed that procreation provides goods for potential parents. Thus bringing a child into the world is not only permissible in regards to the potential person, but also in regards to the potential parent.

Bibliography

Appiah, A. (2005). *The Ethics of Identity.* Princeton, NJ: Princeton University Press.

Arrhenius, G. (2000). An Impossibility Theorem for Welfarist Axiologies. *Economics and Philosophy, 16*(2), 247-266.

Benatar, D. (2006). *Better Never to Have Been: the harm of coming into existence.* New York: Oxford University Press.

Brighouse, H., & Swift, A. (2006). Parents' rights and the value of the family. *Ethics, 117*(1), 80-108.

Brighouse, H., & Swift, A. (2014). The Goods of Parenting *Family-Making: Contemporary Ethical Challenges.* Oxford: Oxford University Press.

Broome, J. (2004). *Weighing Lives.* Oxford: Oxford University Press.

Bykvist, K. (2007a). The benefits of coming into existence. *Philosophical Studies: An International Journal for Philosophy in the Analytic Tradition, 135*(3), 335-362.

Bykvist, K. (2007b). The good, the bad and the ethically neutral., *Economics and Philosophy* (Vol. 23, pp. 97-105).

Griffin, J. (1988). *Well-Being: Its Meaning, Measurement and Moral Importance.* Oxford: Oxford University Press.

Handfield, T. (forthcoming). Rational choice and the transitivity of betterness. *Philosophy and Phenomenological Research.*

Harman, E. (2004). Can We Harm and Benefit in Creating? *Philosophical Perspectives, 18,* 89-113.

Holtug, N. (2001). On the Value of Coming into Existence. *The Journal of Ethics, 5*(4), 361-384.

Kavka, G. S. (1982). The Paradox of Future Individuals. *Philosophy and Public Affairs, 11*(2), 93-112.

Keller, S. (2007). *The Limits of Loyalty.* Cambridge, UK: Cambridge University Press.

Mulgan, T. (2004). Two Parfit Puzzles. In J. R. a. T. Tännsjö (Ed.), *The Repugnant Conclusion: Essays on Population Ethics.* Dordrecht: Kluwer.

Mullin, A. (2007). Giving As Well As Receiving: Love, Children, and Parents. *Symposium: Canadian Journal of Continental Philosophy (Revue canadienne de philosophie continentale), 11*(2), 383-395.

Parfit, D. (1982). Future Generations: Further Problems. *Philosophy and Public Affairs, 11*(2), 113-172.

Parfit, D. (1984). *Reasons and Persons.* Oxford: Clarendon Press.

Parfit, D. (1986). Overpopulation and the Quality of Life. In P. Singer (Ed.), *Applied Ethics.* Oxford: Oxford University Press.

Qizilbash, M. (2005). Transitivity and Vagueness. *Economics and Philosophy, 21*(1), 109-131.

Rachels, J. (1997). *Can Ethics Provide Answers?: and Other Essays in Moral Philosophy.* Lanham, Md.: Rowman & Littlefield.

Rachels, S. (2004). Repugnance or Intransitivity: A Repugnant but Forced Choice. In J. R. a. T. Tännsjö (Ed.), *The Repugnant Conclusion: Essays on Population Ethics.* Dordrecht: Kluwer.

Reshef, Y. (2013). Rethinking the value of families. *Critical Review of International Social and Political Philosophy, 16*(1), 130-150.

Roberts, M. A. (1998). *Child Versus Childmaker: future persons and present duties in ethics and the law.* Lanham: Rowman & Littlefield.

Roberts, M. A. (2004). Person-Based Consequentialism and the Procreation Obligation. In J. R. a. T. Tännsjö (Ed.), *The Repugnant Conclusion: Essays on Population Ethics.* Dordrecht: Kluwer.

Rulli, T. (2014). The Unique Value of Adoption *Family-Making: Contemporary Ethical Challenges.* Oxford: Oxford University Press.

Ryberg, J. (1996). Parfit's Repugnant Conclusion. *The Philosophical Quarterly, 46*(183), 202-213.

Ryberg, J. (2004). The Repugnant Conclusion and Worthwhile Living. In J. R. a. T. Tännsjö (Ed.), *The Repugnant Conclusion: Essays on Population Ethics.* Dordrecht: Kluwer.

Schoeman, F. (1980). Rights of Children, Rights of Parents, and the Moral Basis of the Family. *Ethics, 91*(1), 6-19.

Schwartz, T. (1978). Obligations to Posterity. In R. I. Sikora & B. M. Barry (Eds.), *Obligations to Future Generations.* Philadelphia: Temple University Press.

Tännsjö, T. (2004). Why We Ought to Accept the Repugnant Conclusion. In J. R. a. T. Tännsjö (Ed.), *The Repugnant Conclusion: Essays on Population Ethics.* Dordrecht: Kluwer.

Temkin, L. S. (1987). Intransitivity and the Mere Addition Paradox. *Philosophy & Public Affairs, 16*(2), 138-187.

Temkin, L. S. (2012). *Rethinking the Good: Moral Ideals and the Nature of Practical Reasoning.* Oxford: Oxford University Press.