

Quaestiones Infinitae

PUBLICATIONS OF THE DEPARTMENT OF PHILOSOPHY
UTRECHT UNIVERSITY

VOLUME LVI

Copyright © 2011 by Tatjana Visak
All rights reserved

Printed by Wöhrmann Print Service, Zutphen
ISBN 9789039355657

KILLING HAPPY ANIMALS

Explorations in utilitarian ethics

Het doden van gelukkige dieren.
Verkenningen binnen de utilistische ethiek
(met een samenvatting in het Nederlands)

Darf man glückliche Tiere töten?
Eine Untersuchung zur utilitaristischen Ethik
(mit einer Zusammenfassung in deutscher Sprache)

Proefschrift

ter verkrijging van de graad van doctor aan de Universiteit Utrecht op
gezag van de rector magnificus, prof. dr. G. J. van der Zwan,
ingevolge het besluit van het college voor promoties in het openbaar
te verdedigen op dinsdag 31 mei 2011
des ochtends te 10.30 uur

door

Tatjana Visak

geboren op 12 december 1974
te Giessen, Duitsland

Promotor: Prof.dr. M. Düwell

Co-promotor: Dr. M. F. Verweij

CONTENTS

CONTENTS	v
ACKNOWLEDGEMENTS	xi
1. INTRODUCTION	2
1. Aim of this thesis	2
2. Animal-friendly animal husbandry?	8
3. Why utilitarianism?	14
4. Do animals count morally?	19
5. Outline of this thesis	21
2. UTILITARIANISM AND ANIMAL HUSBANDRY	27
1. Introduction	27
2. What is a moral theory?	28
3. Basic elements of utilitarianism	30
4. A closer look at utilitarianism	34

4.1 Defining welfare.....	35
4.2 Aggregation of welfare	40
4.3 Evaluative focal points	44
4.4 Quantity of welfare or net benefits?.....	46
4.5 Underlying Argument for welfare maximization	50
5. A controversial moral theory	53
6. Utilitarianism on animal husbandry	59
7. Conclusion.....	68
3. ARE ANIMALS HARMED LESS.....	72
BY DEATH?	72
1. Introduction	72
2. The relevance of future-oriented desires.....	74
3. Wants versus value	79
4. Does psychological connectedness matter?.....	82
5. Conclusion.....	88
4. THE REPLACEABILITY ARGUMENT	90
1. Introduction	90
2. The underlying idea of the Replaceability Argument.....	91
3. Defining the Replaceability Argument: specifying the relevant conditions.....	93
4. Practical implications for animal-friendly animal husbandry: can the conditions be met?.....	97
4.1 Do the animals have pleasant lives?	98
4.2 What about unbalanced negative (side-) effects?	101

5. The scope of the Replaceability Argument.....	102
5.1 Singer’s account of welfare: desire-satisfaction account plus moral ledger model	104
5.2 Singer’s account of welfare is pessimistic.....	109
5.3 Singer’s account of welfare is linked to a particular account of the harm of death.....	110
5.4 Singer’s account of welfare is ad-hoc in several respects.....	111
5.5 Singer’s account of welfare fails to restrict replaceability in the desired way	113
5.6 Singer’s most recent strategy.....	115
5.7 Summing up.....	120
6. The Replaceability Argument is based on a controversial view.....	121
7. Conclusion.....	123
5. TOTAL VIEW VERSUS PRIOR.....	125
EXISTENCE VIEW.....	125
1. Introduction	125
2. Whose welfare should count?.....	127
3. Prior Existence View versus Total View.....	130
4. Are both views coherent utilitarian views?.....	133
5. Conclusion.....	136
6. CAN EXISTENCE BE BETTER FOR A BEING THAN NON- EXISTENCE?.....	138
1. Introduction	138
2. The view that coming into existence can make a being better or worse off	139

3. Are existence and non-existence commensurable?	144
4. How to decide whether the absence of any value is the same as zero value?	147
5. For whom would non-existence have neutral value?	155
6. Conclusion.....	157
7. PERSON-AFFECTING RESTRICTION AND NON-IDENTITY PROBLEM.....	159
1. Introduction	159
2. The Non-Identity Problem	161
3. The Wide Person-Affecting Restriction.....	165
3.1 <i>De dicto</i> betterness.....	166
3.2 A morally relevant notion of identity.....	169
4. Possible criticisms.....	177
4.1 Individuals in the wide sense cannot have interests	177
4.2 The Wide Person-Affecting Restriction invites the Replaceability Argument	179
5. Advantages of the Wide as opposed to the Narrow Person-Affecting Restriction	180
6. Conclusion.....	185
8. REPUGNANT CONCLUSION AND.....	187
EXPECTED MISERY ARGUMENT	187
1. Introduction	187
2. The role of intuitions within utilitarianism	188
3. The Repugnant Conclusion	196
4. The Expected Misery Argument	202

5. Can Prior Existence Utilitarianism account for the idea that there is something morally wrong about having the miserable child?	205
6. Conclusion	217
9. VEGANISM VERSUS ANIMAL-FRIENDLY ANIMAL HUSBANDRY	219
1. Introduction	219
2. The Illogic of the Larder	220
2.1 Meat eaters as real animal lovers.....	221
2.2 The argument's relation with both versions of utilitarianism..	223
3. Implications concerning animal husbandry	227
4. Conclusion.....	231
10. GENERAL CONCLUSIONS.....	232
1. Conclusions of the thesis.....	232
2. Remaining questions.....	242
APPENDICES.....	244
A. Different-number choices	244
B. Childless George case.....	250
C. Prior Existence Utilitarianism and transitivity.....	252
D. Arrhenius on the relation between the Prior Existence View and the Person-Affecting Restriction.....	258
E. A more technical argument for the Wide Person-Affecting Restriction.....	262
BIBLIOGRAPHY	275
INDEX.....	294
SAMENVATTING.....	297

ZUSAMMENFASSUNG	314
CURRICULUM VITAE.....	331

ACKNOWLEDGEMENTS

Many people supported my writing of this thesis.

I thank my supervisor Marcel Verweij for his detailed comments and suggestions that made my writings undoubtedly much more accessible for the reader. I enjoyed his grasp of and enthusiasm about the project and our sometimes forgetting the time while discussing it. I also thank the director of the Ethics Institute, Marcus Düwell for backing up my plans and projects and allowing me to follow all kinds of courses and visit and organize conferences. I am grateful for his useful comments as a supervisor of this thesis.

I want to thank the members of the Practical Philosophy colloquium at Utrecht University, the members of the Dutch Research School for Ethics, and the members of the European Society for Agriculture and Food Ethics for giving me the opportunity for discussing some issues related to this thesis. Frans Brom supervised the very beginnings of what became this thesis. Gustaf Arrhenius, Dieter Birnbacher and Peter Singer have provided helpful comments on some draft chapters. My criticism of some of Singer's views in this thesis should be understood along the lines indicated by Dale Jamieson in his

preface to *Singer and his Critics*: “philosophers honor people by criticizing them”.¹

During the last ten years I increasingly enjoyed the company of my colleagues at the Ethics Institute during daily work and on many special occasions. For what concerns this thesis, particular thanks to Frederike Kaldewaij for sharing ideals and insights about animal ethics, Bernice Bovenkerk for detailed proof reading, and Jos Koole for technical support.

I am particularly grateful to my husband Arjen, for caring for our kids whenever I took days off from family life to ‘relax’ at my office, and whenever I took ‘holidays’ for visiting conferences. After biking home from my office through the forests, Arjen’s delicious animal-friendly meals and the warm welcome by two enthusiastically yelling great apes made me feel like a very lucky animal, happy indeed.

¹ Jamieson (1999): preface.

The sections [...] that have left me in the greatest philosophical uncertainty are those parts [...] that discuss whether there is some sense in which bringing into existence a new being – whether a human being or a nonhuman animal – can compensate for the death of a similar being who has been killed. That issue in turn leads to questions [...] [which] may seem arcane and far removed from [...] ‘practical ethics’ [...]. As we shall see, they can serve as an example of how our judgments of what is right and wrong need to be informed by investigations into deep and difficult philosophical issues.

(Peter Singer, Preface to the third edition of Practical Ethics, forthcoming 2011.)

1. INTRODUCTION

1. Aim of this thesis

In the Netherlands, one of the most densely populated countries with about 16 million human inhabitants, 12 million pigs are kept in animal husbandry. Even though those animals are invisible to the general public, due to their number and the way they are treated, those pigs, and farmed animals in general, account for a lot of suffering. It is broadly acknowledged in the Netherlands and in other Western countries that this suffering matters morally. Animals are considered due objects of our moral concern. Their welfare matters to them, and therefore it is broadly accepted that we may not neglect it. This acknowledgment has resulted in a political and societal striving for animal-friendly animal husbandry. Animal-friendly animal husbandry is conceived as a *moral* goal.¹

¹ I will speak about animal-friendly animal husbandry as being considered a moral goal or moral ideal. An ideal can be understood as a model of excellence or conception of perfection around which people can orient their thoughts and conduct, both on a private and on a public level. See Brownlee (2010): 242-243.

Animal-friendly animal husbandry is a moral goal that strikes me as inherently inconsistent. On the one hand it is assumed that animals deserve our moral consideration and that causing them suffering should be avoided. On the other hand, it is considered perfectly acceptable to kill animals. However, their killings are unnecessary: the consumption of animal products is not needed for human health.² How can it be justified, then, that we are not allowed to kick them, while we are allowed to kill them? In this thesis, I will explore a possible explanation for the apparent tension in the moral aim of animal-friendly animal husbandry. I will explore whether the aim can be justified from within a major theoretical account of what we morally ought to do: utilitarianism.

Historically, utilitarianism is the moral theory that has contributed most to the recognition of animal suffering as an evil.³ Utilitarianism, in its simplest form, states that the morally right action is the one that results in the greatest net balance of enjoyment over suffering. Thus, outcomes of possible actions have to be compared in terms of their overall effects on welfare. The welfare consequences that an outcome contains have to be brought together (aggregated) in order to determine the overall value of the outcome. Utilitarianism thereby takes into consideration the welfare consequences of all sentient beings, and of all equally. Thus, utilitarianism strives for the maximization of welfare for all sentient beings.

Yet, how does utilitarianism judge the killing of animals? In particular, could the goal of animal-friendly animal husbandry be justified

² For example, the American Dietetic Association, which is the world's largest organization of food and nutrition professionals, claims (2009): "It is the position of the American Dietetic Association that appropriately planned vegetarian diets, including total vegetarian or vegan diets, are healthful, nutritionally adequate and may provide health benefits in the prevention and treatment of certain diseases. Well-planned vegetarian diets are appropriate for individuals during all stages of the life cycle including pregnancy, lactation, infancy, childhood and adolescence and for athletes."

³ Nussbaum (2006): 338.

from within utilitarian moral theory? The mainstream view is that utilitarianism supports the ideal of animal-friendly animal husbandry. For instance, Peter Singer, the most famous animal ethicist and utilitarian, and one of the most influential philosophers of the world, makes it very clear that the standard form of intensive animal husbandry that causes animals a lot of suffering is morally indefensible. In contrast, Singer suggests that there *might* be nothing wrong with granting animals pleasant lives and then painlessly kill them.⁴ Typically, the utilitarian position is presented as opposing unnecessary suffering, while allowing the painless killing of animals, or at least keeping this latter issue somewhat vague.⁵

This utilitarian 'welfarist' position is commonly contrasted with the 'abolitionist' position of animal rightists.⁶ Those who strive for 'empty cages' rather than 'bigger cages' have accused utilitarianism of not taking animals seriously enough. In this thesis I will show that this is an incomplete and therefore wrong picture of utilitarianism. Whether or not utilitarianism is compatible with the ideal of animal-friendly animal husbandry depends on which version of utilitarianism one accepts or refers to. While one version, Total Utilitarianism, is compatible with the ideal of animal-friendly animal husbandry, another version, Prior Existence Utilitarianism, is not. The aim of this thesis is to explore the

⁴ Singer made this suggestion during his presentation at the Minding Animals Conference, Newcastle, Australia, July 2009 by clearly opposing animal suffering, while discussing the killing of animals only by way of questions. See also Singer (1993): 125, where Singer does not principally oppose the routine killing of at least some animals. More recently, due to new insights about the capacities of animals such as pigs and fish, Singer (forthcoming 2011: ch. 5) is more critical about the routine killing of pigs, cattle and the other animals that usually end up on people's plates.

⁵ Sandoe (2008): 80-81. See also Animal Ethics Dilemma (www.ae.imcode.com), a program for teaching animal ethics, developed by Sandoe and others, where the utilitarian position is presented as favouring animal welfare, while, unlike the 'rights view', not being opposed to the killing of animals. See also Norcross (2004), who eloquently attacks the torturing of animals in intensive animal husbandry while not mentioning whether the routine killing of animals is a moral problem.

⁶ Sandoe (2008): 81. Francione (1996), Francione (2000): ch. 6, Regan (2004): 205-211.

various implications and assumptions of both Total Utilitarianism and Prior Existence Utilitarianism.

Singer's Replaceability Argument plays a central role in the utilitarian evaluation of animal-friendly animal husbandry and therefore also in this thesis. This argument claims that the welfare loss that is caused by killing an animal that could otherwise have had a pleasant future can be compensated by bringing into existence another animal that would not otherwise have existed and whose life contains at least as much welfare as the future of the killed animal would have contained. That argument has likely been the most controversial part of Singer's whole theory and has contributed to the criticism of being overly impersonal. As Singer points out:

The replaceability argument was probably the most controversial, and widely criticized, argument in *PE* [*Practical Ethics*]. Unfortunately none of the critics have offered satisfactory alternative solutions to the underlying problems to which replaceability offers one, if not very congenial, answer.⁷

One of the versions of utilitarianism that I present in this thesis avoids the Replaceability Argument.

The Total View and the Prior Existence View are fundamentally different utilitarian views about the question: Across whom should welfare be aggregated? Those different versions of utilitarianism have different implications as to how personal or impersonal the utilitarian moral theory is. A major strand of criticism that has been brought forward against utilitarianism is that utilitarianism is not really interested in benefits and harms for individuals, but rather in welfare as an abstract

⁷ Singer (1993): 386. See also Singer (forthcoming 2011): 107.

quantity. What many people consider really wrong with utilitarianism is that instead of valuing happiness because of what it does for sentient beings, it values sentient beings for what they do for happiness.⁸ I will present a comprehensive version of utilitarianism that is concerned with harms and benefits for sentient beings, rather than with welfare as an abstract quantity. Thus, by exploring the utilitarian stance on animal husbandry, I will also address two related fundamental issues within the moral theory: the question across whom to aggregate and the criticism of being impersonal.

Furthermore, my exploration of the utilitarian stance on animal husbandry will lead to fundamental philosophical discussions beyond animal husbandry and beyond utilitarianism. Topics to be discussed in this thesis are the moral status of possible beings, and the question whether causing a being to exist can harm or benefit that being. A further topic is the question what to do if our choices determine who will exist or how many will exist. It will be explored what the harm of death consists in for humans and animals; and the relevance of some competing concepts of identity will be investigated.

This thesis, then, examines the Replaceability Argument. This argument has been used to justify the killing of animals. The focus is on the relevance of this argument for the practice of animal husbandry. It should be noted, though, that the argument is also relevant for other practices in which animals are killed and replaced. Relevant practices include:

- Meat production (An animal is killed and a new animal takes its place: that is the core of the business.)
- Milk and dairy production (A dairy cow is killed and a new one takes her place. Here the new calf already exists, as the

⁸ Jamieson (1984): 218.

dairy cow has to give birth each year in order to give milk. This calf would be killed if it were not used as a replacement of a killed cow. So, its existence does not directly depend on the killing of another cow, but its continued life does.)

- Egg production (A laying hen is killed and a new one takes her place.)
- Aquaculture (Animal husbandry with fish.)
- Animal experimentation, and in particular the breeding and selling of animals for that purpose (Animals with particular characteristics, so-called 'animal lines' are produced, and sold animals are replaced by new ones.)
- Sport fishing, where bred fish is released in the waters where the fishing takes place in order to keep the fish population on a certain level.
- Sport hunting of animals bred for this purpose (Animal populations are often maintained by providing them with food, and sometimes tame animals are released for the hunt.)

Those are all practices in which non-human animals are killed and replaced. In those practices, the killed animals are not replaced for moral reasons. Rather, the animals are replaced for practical reasons: their replacement is necessary in order to continue the practice. According to the Replaceability Argument, the fact that killed animals are replaced (for whatever reason) is relevant for the moral evaluation of the killing. The Replaceability Argument does also have implications for human beings. Those implications will be indicated as well in my exploration of the Replaceability Argument.

In the remainder of this introductory chapter, I will further elaborate the central practical question of this thesis, namely whether utilitarianism is compatible with the ideal of animal-friendly animal husbandry. I explain what I mean by 'animal-friendly animal husbandry'

(section 2) and why my focus will be on utilitarianism (section 3). Furthermore, I will introduce the question whether we do have any moral obligations towards animals and relate it to utilitarianism and the justification of animal-friendly animal husbandry (section 4). This will be followed by the outline of this thesis (section 5).

2. Animal-friendly animal husbandry?

In the last decades, there has been a growing public concern about the welfare of animals in animal husbandry. After World War II, the agricultural policy in the Western World was directed towards safeguarding the food supply. The focus on sufficient and safe food and economic efficiency resulted in an industrialization of the production process, including animal production. Since then, the aim of safeguarding the food supply in Western countries has been reached. In particular since the late 1970's, ethical concerns regarding our treatment of animals have been voiced. Those ethical concerns have influenced policy making about animal use in general, and animal production in particular. For instance, concerns about animal welfare have been translated into animal welfare laws and regulation.⁹

Many people, however, are still unsatisfied with the treatment of animals in common production systems. Recent food scares such as swine influenza, foot and mouth disease and avian flu - with the accompanying pictures in the media - have again contributed to ethical criticism of intensive animal farming. Moreover, intensive animal husbandry is broadly criticized because of its negative effects on animal welfare. For instance, in the Netherlands there has been a citizen's initiative, titled 'Stop wrong meat', which asked for – and indeed resulted

⁹ While some laws seem to acknowledge that animals are sentient creatures and are morally owed a certain treatment, animals are also still 'things' according to the law. This makes the legal status of the animal unclear. See Bordes (2010).

in - a parliamentary reconsideration of intensive animal husbandry.¹⁰ The Party for the Animals, which has entered Dutch Parliament, speaks out against intensive animal husbandry. In radio spots, prominent party-members told the public something like: “It is up to you whether or not you should eat meat, but please be aware of how it is produced and choose the animal-friendlier meat.” Other organizations on behalf of animals bring a similar message. For instance, the yearly flyer of *Varkens in Nood (Pigs in Peril)* starts with: “Don’t eat meat from factory farms at Christmas”. Many people consider intensive animal husbandry as such morally unacceptable and hold that animals should not be treated in that way. The actual striving of politicians, societal organizations, citizens and consumers for a more animal-friendly agricultural sector must be understood in that context.

European governments promote ‘animal-friendly husbandry systems’ and citizens acknowledge the moral superiority of more animal-friendly products.¹¹ For instance, the Dutch government strives for 5% of all animal sheds being ‘entirely sustainable and animal-friendly’ in 2011.¹² A growing number of producers and consumers (though still less than 5%) choose organic animal products.¹³ In the Netherlands more than 95% of consumed meat is from intensive production systems, but this fact does not prove a lack of support for animal-friendly animal husbandry. Rather, many people who do consume products from intensive production systems agree, on reflection, that animals should not be treated in that way.¹⁴ People indicate that they would rather like to

¹⁰ This does not mean that a more animal-friendly system must necessarily be an extensive one. It only means that current intensive systems are conceived as being bad for animal welfare.

¹¹ See, for instance, the coalition formation of the Dutch government, 7 February 2007.

¹² Nota Dierenwelzijn, Ministry of Agriculture, 2007.

¹³ Platform Biologica, EKO-Monitor 2002, no. 13, p. 3.

¹⁴ In the Netherlands, 86% of the respondents even indicated to be willing to pay more for meat, if it is produced in a more animal-friendly way (Intomart, May 2001).

consume products from animal-friendlier production systems and would probably do so if those products were as easily available and had the same price. Due to several reasons, those moral values are very often not followed by appropriate purchasing behavior.¹⁵ One reason seems to be that governments, rather than individual consumers, are held responsible for changing the agricultural system.¹⁶ So, even if the actual behavior seems to be motivated by other interests, the goal of animal-friendly animal husbandry is at least theoretically broadly supported.

There are roughly two groups of people who do not support animal-friendly animal husbandry. One group finds animal husbandry *unacceptable*, even if the animals have ‘happy lives’. This group considers the use and killing of animals for consumption morally unacceptable. It seems to be a small group, as is indicated by the small percentage of vegans; i.e. people who strive, as far as practically possible, to avoid animal products. The other group that does not support animal-friendly production finds animal husbandry *acceptable*, even if the animals have bad lives. This group is, I suspect, small as well. At least mammals, such as cows, pigs and sheep, are generally conceived as sentient beings, whose welfare deserves some protection. This, again, indicates the widespread support for animal-friendly animal husbandry.

“Animal husbandry” refers to the keeping of animals in order to produce animal products for human consumption. This can be meat or other animal products such as milk or eggs. Animal husbandry implies that the animals undergo a certain treatment from birth to death. The treatment, or use, differs per sort of animal, production purpose and production method. For instance, animals are kept in confinement; so-called ‘broiler chicks’ are often de-beaked; cows are milked. Besides the *use* of animals, animal husbandry, as I understand it here, does also imply the *killing* of animals. This does not only hold true in the obvious case of

¹⁵ Visak (2004). Sterrenberg (2001): 88.

¹⁶ Diederer (2003), Visak (2004).

meat production. In case of milking cows, the cows are also usually killed after several years of use, based on an economic cost-benefit calculation. Furthermore, their male calves are usually killed. Laying hens are killed after a period of use, and male chicks, which cannot be used as laying hens, are killed soon after their birth. So, killing animals is implied by animal husbandry, as I understand it here, not only in the straightforward case of meat production.

“Animal-friendly” in turn refers to husbandry systems that allow the animals more possibilities to live out their species-specific behavior. This means that animals should be given the opportunity to display what is typical behavior for their sort. For instance, cows should be allowed to graze and hens should be allowed to scratch. In current intensive animal production, the opportunities for animals to live out their species-specific behavior are very limited. Research has shown that animals that are allowed to show species-specific behavior have enhanced welfare.¹⁷ The general public seems to know this. At least, people seem to have the intuitive idea that keeping animals in more ‘natural’ ways is better for the animals. Thus, animal-friendly systems provide more adequate housing with more space. At the same time, the practice of animal husbandry as such sets practical limits to how ‘naturally’ (in the sense of ‘species specific’) the animals can be allowed to behave. Jan Brinkhorst, a former Dutch minister of agriculture, expressed the idea of granting animals more – but still limited – options for living out their species-specific behavior as follows:

“Several square centimeters extra to grub in order to alleviate the suffering of the poultry should no longer be only window-dressing, good for a misleading image on the package. No, what naturally grubs shall grub, what grazes shall graze. This will be the basic rule number

¹⁷ Bracke and Hopster (2006).

one, being aware, however, that the recreation of Eden is not our ambition.¹⁸

A more recent Dutch Minister of Agriculture, Gerda Verburg, also identifies the option of living out ‘natural behavior’ as an important component of animal welfare.¹⁹

The concern with animal welfare does also imply striving for a reduction of stress and pain by avoiding unnecessarily painful or stressful treatments. It is debated what exactly counts as ‘unnecessary’. A recent discussion in this respect resulted in the (European) ban of the castration of piglets without anesthesia. The castration routine was meant to prevent a certain odor of the meat, but it is now considered unnecessary, because preventing the odor can be achieved by less painful means. The Dutch Minister of Agriculture holds that some stress and discomfort for the animals is unavoidable. According to her, it should not count as bad animal welfare, if animals experience stress while they are packed for transport to the slaughterhouse. Chronic stress, due to bad housing, on the other hand, should count as bad animal welfare. She claims: “It does always depend on the frequency, time span and intensity [of suffering – TV], due to human interference and action.”²⁰ So, apparently, animals with ‘happy lives’ are meant to be satisfied and well-off, most of the time, while their suffering should be minimized as much as possible within the context of animal production.

The ideal of animal-friendly animal husbandry implies a remarkable combination of moral duties towards animals. The requirement to take an animal’s welfare into account suggests that this animal counts morally. Apparently, it is worthy of our moral

¹⁸ Report of the German-Dutch debate on the future of agriculture (2002): 14. For a similar comment, see Nota Dierenwelzijn 2007, p. 9.

¹⁹ Nota Dierenwelzijn 2007, p. 19.

²⁰ Nota Dierenwelzijn 2007, p. 10.

consideration. If the animal counts morally, however, why do we only have to protect its wellbeing and not its life? How can it be explained that we are not allowed to kick it, while we are allowed to kill it? This seems strange. As I indicated, in the Dutch debate about the future of agriculture, the term ‘animal-friendly animal husbandry’ is used in order to indicate the moral aim of the transformation of the sector. It is claimed that animal production can be ‘entirely animal-friendly’. Isn’t ‘animal-friendly animal husbandry’ a paradox in terms? Can the practice of routinely using and killing animals rightly be called ‘animal-friendly’ at all? Are our moral duties towards animals compatible with the routine killing of animals for food production?

This is what the initiatives striving for animal-friendly animal production suggest. Recently an overview has been published about the welfare of pigs in different husbandry systems. The title, which is *Welfare of Pigs from Birth to Slaughter*, reveals that welfare considerations are discussed *within* the practice of pig husbandry, rather than being used to question this practice.²¹ Similarly, the title and aim of the above-mentioned initiative ‘Stop wrong meat’ point to the underlying assumption that there is also ‘right meat’. With other words: animal husbandry as such is not the problem, only the unnecessary suffering of animals is. Equally, the Party for the Animals agitates against intensive animal production, but not explicitly against animal production as such. In the same vein, the folder of *Pigs in Peril* does not say: “Don’t eat meat at Christmas”. The organization condemns ‘abuses’ in slaughterhouses, such as poor anesthetization, rather than slaughterhouses as such. The general understanding of ‘avoiding unnecessary suffering’ tends to accept animal husbandry as a given. It does not question in how far animal production as such is ‘unnecessary’.

²¹ Faucitano & Schaefer (2008).

Opposing intensive animal husbandry rather than animal production as such might be a strategic choice for people who ultimately like to see an end to all animal production. Yet, it seems that many people do honestly believe that it is morally permissible to use and kill animals in animal husbandry, provided that the animals have 'happy lives'. Even if animal-friendly animal husbandry is promoted for strategic purposes, it is worth exploring whether this ideal is defensible. After all, people and organizations that make moral claims can be challenged to defend those claims. In that case, a lack of any plausible and coherent defense will cause problems and the strategy might not prove successful in the long run.

So, the ideal of animal-friendly animal husbandry seems to be broadly accepted, but there is also something conspicuous about it. This gives me a reason to take a critical look at this ideal, which is what I will do in this thesis.

3. Why utilitarianism?

At first glance, utilitarianism seems to be a promising candidate moral theory for supporting the ideal of animal-friendly animal husbandry. After all, utilitarianism holds that welfare is the only value that is ultimately relevant for ethical evaluation. Utilitarianism claims that moral agents should promote overall welfare. Therefore, what matters is how an action affects the welfare of others, including animals. Every affected being counts equally. For instance, it does not matter whether the affected being is male or female, human or non-human. So, as the maximization of welfare is central to utilitarianism, this theory naturally includes all sentient beings, including animals, as moral objects. This is expressed in two famous dictums of Bentham (1748 – 1832), one of the founding fathers of this moral theory: “The question is not, Can they

reason? nor Can they *talk?* But, Can they *suffer?*”²² And: “Each to count for one and none for more than one.” If we are to reduce suffering and promote welfare, there seems to be much to win in the field of animal husbandry. This holds all the more true if animal suffering counts equally to human suffering. So, utilitarianism justifies putting animal welfare high on the agenda.

Utilitarianism does not categorically rule out killing. According to utilitarianism, the consequences of every action, and the consequences alone, determine whether it is right or wrong. There are no categories of actions that are obligatory or forbidden as such and under all circumstances. Everything depends on how overall welfare is affected by a particular action. Therefore, killing can be justified if it is compatible with the maximization of overall welfare. Its positive effects can in principle compensate its negative effects. Considering all this, utilitarianism seems to be a promising candidate for supporting the ideal of animal-friendly animal husbandry.

One might wonder whether there are other plausible candidates. Support for the assumption must come from a moral theory that considers animal welfare a proper issue of moral concern and that does not rule out the killing of animals. Animal-rights theories are the most prominent alternative to utilitarian theories when it comes to accepting animals as moral objects. However, animal-rights theories do usually accept the *right to life* as a very basic moral right.²³ Therefore, animal rights theories are unlikely to support animal husbandry.²⁴ There might be other theories, except utilitarianism, which show concern for animal welfare, *and* allow the killing of animals.

²² See his *Introduction to the Principles of Morals*, chapter 17.

²³ Conceptually, animal-rights theories need not accept the right to life.

²⁴ Examples are Regan (2004): ch. 9, Pluhar (1995): 271-272, Cochrane (2007).

Some have supported the goal of animal-friendly animal husbandry without accepting that we have any moral duties towards animals. According to those approaches, the duty to care well for the animals is not strictly speaking something we owe *to the animals*. Rather, not caring well for the animals is considered a shame for the human race and/ or a sign of disrespect for certain moral and spiritual realities.²⁵ Kant is known for having suggested that mistreating animals should be avoided because of its effects on humans. Those lines of argument imply that any duties we might have with regard to the treatment of animals are not duties towards the animals. Animals are not accorded moral status. This is opposed to the idea that animal welfare matters because it matters for the animals.

Others have supported the goal of animal-friendly animal husbandry on the basis of a kind of relational moral theory. Such a theory holds that our moral duties towards others are based on the relationship that we have with them. For instance, our relationship to pets is different than our relationship to laboratory animals or to wild animals. Animals that are kept in animal husbandry are considered a separate category. The diverse practices are taken as a given. Any moral duties towards animals are thought to *arise from* those practices. It is certainly true that for instance a mouse is treated differently depending on whether it is seen a pet, a laboratory animal or a pest. The question that this approach faces is whether this differential treatment can be morally justified, and if so, how. Critics might claim that such an approach is not suited for critically assessing the very practices that are taken as a given. Such an approach seems to preclude the question whether animal husbandry as such is compatible with our moral duties towards animals.

²⁵ Scruton (2006).

The approach of the Dutch minister of agriculture is illustrative. She claims:

For me, it is a given that human beings keep animals and have domesticated animals in the course of the time. This means that, with respect to the kept animals, we cannot go back to the situation with the appropriate natural behavior, in which the animals move freely in nature and are being hunted as food for humans.²⁶

This starting point of a moral evaluation of animal husbandry restricts beforehand what is accepted as a possible outcome of the moral evaluation. The alleged arguments, however, do not provide a sufficient justification for the practice of animal husbandry. It is a fact that some humans are actually keeping animals and have domesticated them in the course of history. This fact, however, does nothing to *morally justify* the (current) practice of animal husbandry. Furthermore, it is a misleading opposition to depict the situation as a choice between sending all actual farm animals into the forests or keeping them in the farm. The possibility of (gradually) stopping the production of farm animals and animal husbandry is completely neglected and kept out of the discussion. Therefore this approach is unsuited for morally scrutinizing our treatment of animals, including the practice of animal husbandry.²⁷

Yet others have defended the goal of animal-friendly animal husbandry by depicting animal husbandry as a mutually beneficial endeavor, which benefits both humans and the animals that are part of it.²⁸ The underlying argument is that the animals in question have good lives, and they wouldn't exist at all, if it weren't for our consumption.

²⁶ Verburg, G. (2007): 9.

²⁷ Another example is Swart (2006). NVBe.

²⁸ Scruton (2006).

Thus, or so it is claimed, they are benefited by those who consume them. This argument, which is also defended by some utilitarians, will be discussed later on in this thesis (in chapter 9). The argument assumes that having a happy life is better for the animal than having no life at all. I argue in chapter 6 that such a comparison is not warranted.

Finally, some approaches take the idea of mutual benefit further and claim that animal-friendly animal husbandry would be sanctioned by an imaginary contract between humans and animals.

We could conceive of the (decent!) keeping of cows, pigs and sheep for slaughter as a kind of contract between people and farm animals: human beings care for the beasts, and the beasts give us their products, such as milk and wool, and finally their lives. In exchange for good care, their drink and food, cows, pigs and sheep finally give their lives and we slaughter them for their meat. The contract between people and farm animals gives duties. Humans have to do their part: care well for the animals, don't reduce them to bio-machines, milk- and meat machines; then the animals will give their lives for our meat.²⁹

At first glance such an approach might sound plausible. People care well for the animals and get something in return. A problem with such an approach is that it would need to be shown that animals agree with a contract, which sanctions their being used and killed. Furthermore, it would need to be shown that making such a contract is ethical. If we assume that presupposing such a contract with one's possible future child would be unethical, it would have to be shown why it should be ethical in the case of non-human animals.

This brief sketch shows that there are other approaches besides utilitarianism that have considered the question whether animal-friendly

²⁹ Korthals (2002): 137 (my translation from Dutch).

animal husbandry is acceptable. In this thesis, I will focus on utilitarianism, because I want to explore whether the routine killing of happy animals is compatible with this major moral theory, which is known for taking animals seriously. I will explore whether utilitarianism is compatible with the goal of animal-friendly animal husbandry. I will explore the assumptions and implications of two versions of utilitarianism. One of them supports animal-friendly animal husbandry; the other does not.

4. Do animals count morally?

Do we have to take animals into account in our moral considerations? According to utilitarianism, the morally right thing to do is whatever maximizes welfare. Animal welfare counts on a par with human welfare. Thus, it follows that, according to utilitarianism we have to take effects on the welfare of animals into account in our moral considerations. In fact, it follows that all beings that can be affected in their welfare count morally. Effects on the welfare of all beings that 'have a welfare' have to be taken into account. Hence, it depends on one's exact definition of welfare, which beings 'have a welfare' and thus deserve moral consideration. For instance, if welfare is defined in terms of desire-satisfaction and frustration, all beings that have desires need to be included in moral considerations. If, alternatively, welfare is defined in terms of pleasure and pain, then it follows that all beings that can experience pleasure and pain count morally. As there is a strong link between sentience and the having of desires (in the sense that all and only sentient beings seem to have desires, at least the desire not to feel pain), both accounts of welfare usually come to the conclusion that all sentient beings deserve moral consideration. This is because all and only sentient beings care about what happens to them. They can be affected in their welfare. This is what is obviously morally relevant for a theory that aims at the neutral maximization of welfare.

That an entity is a due object of moral consideration is typically expressed by saying that the entity has moral status. The question about the moral status of an entity is basically the question about whether and in what way the entity has to be included in moral considerations. According to utilitarianism, as I explained, that an entity has moral status means nothing more and nothing less than that effects on the entity's welfare have to be taken equally into account in the evaluation of outcomes. As such, the fact *that* an entity is accorded moral status has no implications as to *how* this entity might be treated. As explained, utilitarianism does not accept that certain kinds of actions as such are morally right or wrong, or that there are any limits as to how an entity ought to be treated. Whether an action is right or wrong and how an entity ought to be treated depends on the consequences of the action in terms of overall welfare. It depends on nothing else.

As we have seen, utilitarianism accords equal moral status to all sentient beings. Equal moral status is compatible with unequal treatment. If different animals have different interests, treating them differently – each according to his or her interests – is compatible with granting them equal moral consideration. For instance, some utilitarians have argued that killing a normal human teenager, is worse than killing a mouse. According to that view, mice might be killed for reasons that would not justify the killing of human teenagers. If this unequal treatment is based on the idea that mice lack the interest in continued life, which human teenagers have, killing the mice rather than the teenagers is compatible with equal consideration of their interests. After all, equal consideration of interests means that *equal* interests should get equal consideration. Unequal interests can therefore justify unequal treatment, even in case of equal moral status. (The question whether death is indeed a lesser harm to non-human animals will be discussed in chapter 3.)

5. Outline of this thesis

In *chapter 2* I present the basic elements of utilitarianism. Utilitarianism requires the maximization of welfare, giving equal weight to each being's welfare. I show that utilitarianism needs a particular argument, the Replaceability Argument, in order to be compatible with the goal of animal-friendly animal husbandry. While killing a being that could otherwise have had a pleasant future counts as a welfare loss within utilitarianism, the Replaceability Argument holds, as I said, that this welfare loss can be compensated by bringing a new animal into existence that would not otherwise have existed and whose life contains at least as much welfare as the future of the killed being would have contained. This argument is controversial. So, in order to know whether utilitarianism can support animal-friendly animal husbandry, we need to know whether the Replaceability Argument can be defended within utilitarianism.

In *chapter 3* I start with the above-mentioned feature of utilitarianism that unequal treatment can be justified if the beings in question have unequal interests. I explore whether animals are less interested in their continued life, and whether this difference of interests can justify the routine killing of animals as implied by animal husbandry. I show that some accounts on the harm of death do indeed imply that animals are harmed less by death, while other accounts do not have this implication. All accounts imply that death harms the animals that usually end up on people's plates at least to some degree. Therefore, their routine killing would in any case be morally problematic. Furthermore, quite apart from the harm that death is for the animal, the death of an animal that could otherwise have had a pleasant future implies a welfare loss that utilitarians can take into account in moral considerations. Killing an animal that would otherwise have had a pleasant future is problematic because it causes a loss of overall welfare, and/or, because it causes a loss of welfare for the animal that is killed. That welfare loss must at least be compensated for the killing to be morally justified.

In *chapter 4* I further develop the Replaceability Argument, which proposes a way for compensating the welfare loss due to the killing of an animal that would otherwise have had a pleasant future. I explore the relevant conditions under which an animal is replaceable. I also explore which beings are replaceable, according to this argument. I discuss Singer's arguments for restricting the scope of the Replaceability Argument to non-persons and I point out the necessary assumptions and implications of Singer's arguments. Furthermore, I show that the Replaceability Argument only works if one takes the possible welfare of the possible newly created animal into account in the evaluation of outcomes. This possible new animal does not exist, and whether it will exist at all depends on whether the first animal will be killed. The second animal is therefore called a 'contingent' being: whether it will exist depends on how the moral choice (about whether or not to kill the first animal) is decided. It is controversial whether the possible welfare of contingent beings should be taken into account.

In *chapter 5* I introduce the dispute about whether or not to include the possible welfare of contingent beings in the aggregation of welfare. The Total View counts contingent beings in the aggregation of welfare, while the Prior Existence View does not. The question arises whether both the Total View and the Prior Existence View are really feasible views within utilitarianism. At first glance, it is not obvious that the Prior Existence View really fits within utilitarianism. After all, the Prior Existence View excludes the possible welfare of a certain kind of beings, namely contingent beings, from the aggregation of welfare. This might be considered at odds with the duty to neutrally maximize welfare. Not taking into account the welfare of a certain kind of beings is in need of justification. I show that the refusal to count the possible welfare of contingent beings in the aggregation of welfare is compatible with the duty to neutrally maximize welfare under two assumptions. With other words: in order to be a feasible view within utilitarianism the Prior Existence View depends on two assumptions. It depends on the

assumption that causing an animal to exist *cannot* harm or benefit that animal. It also depends on the assumption that utilitarianism should be concerned with harms and benefits for sentient beings, rather than with quantities of welfare as such. If what matters morally are indeed harms and benefits to sentient beings, and if causing a being to exist can indeed not harm or benefit that being, then (and only then) it makes sense to exclude the possible welfare of contingent beings from the aggregation of welfare.

In *chapter 6* I defend the assumption that causing a being to exist can indeed not harm or benefit that being. Although someone can of course have a good, bad or neutral life, it is impossible to compare the welfare of a person's life with the welfare that this person would have had, if he or she had not come into existence. I argue that the absence of value in case of non-existence is not the same as neutral welfare. Therefore, existence and non-existence are incommensurable and thus it cannot be said that existence can make a person better or worse off than he or she would otherwise have been. Coming into existence cannot be a genuine comparative benefit or harm.

In *chapter 7* I defend the assumption that utilitarianism should be concerned with harms and benefits to sentient beings rather than with the quantity of welfare as such. This assumption is known as the Person-Affecting Restriction. The Non-Identity Problem is known as the major challenge to the Person-Affecting Restriction. The Non-Identity Problem refers to the fact that in particular cases the optimal outcome in terms of welfare does not benefit any particular being. I defend a wide definition of the Person-Affecting Restriction that can avoid the Non-Identity Problem.

Given this result and that of chapter 5, both assumptions that Prior Existence Utilitarianism needs can be upheld. Therefore both Total Utilitarianism and Prior Existence Utilitarianism are feasible utilitarian

views. In the remaining chapters, I explore the implications of both Total Utilitarianism and Prior Existence Utilitarianism.

In *chapter 8* I explore the implication of Total Utilitarianism that is known as the Repugnant Conclusion. I also explore the alleged implication of Prior Existence Utilitarianism that it cannot account for the intuition that bringing into existence a miserable being should be considered morally wrong. I argue that, contrary to what has been claimed, Prior Existence Utilitarianism can account for that intuition. In that context I also reflect on the role of appeals to intuitions in utilitarian moral reasoning.

In *chapter 9* the implications of both versions of utilitarianism for the practice of animal husbandry will briefly be explored. An influential argument in favor of animal-husbandry, known as the Logic of the Larder, claims that keeping animals for consumption and granting them pleasant lives benefits those animals, because if they were not kept for consumption they would not exist at all. I point out that Total Utilitarianism might accept this argument, while Prior Existence Utilitarianism does not. Total Utilitarianism at first glance seems to be compatible with the goal of animal-friendly animal husbandry. However, this is not necessarily the case. Prior Existence Utilitarianism is more in line with the goal of vegan agriculture.

In *chapter 10* I draw conclusions concerning the implications and assumptions of both versions of utilitarianism, and make some notes about further research questions.

Let me end this introductory chapter with a note on my choice of literature. The principle behind my choice of literature in each of those chapters is as follows. Nowhere do I aim at giving a complete overview of a field of discussion. In each chapter I make an argumentative step that I deem to be necessary for the overall argument of this thesis. The literature I use in each chapter is what I considered necessary or helpful for making the specific argumentative step. This includes, ideally, an idea

of all possible counter-arguments. So, it is surely not my aim to mention everything that has ever been written on a subject, how interesting it might be. However, it is my aim to present arguments that hold up against any interesting counter-arguments that have been brought forward in the literature.

2. UTILITARIANISM AND ANIMAL HUSBANDRY

1. Introduction

In this chapter I want to assess in more detail what utilitarianism has to say about animal-friendly animal husbandry (section 6). Utilitarianism is a prominent moral theory. I will first say a few words about what a moral theory is (section 2). Then, I will introduce the basic elements of utilitarianism, which are welfarism, maximization and impartiality (section 3). Once those basics have been introduced, I will take a closer look at the theory and mention some areas in which different versions of utilitarianism can go apart (section 4). It will also be pointed out that utilitarianism is not just any moral theory. Though it is a major moral theory, and often the first to be mentioned in any introduction to ethics (or political theory, for that matter), it is a highly controversial one¹ (section 5).

¹ Kymlicka (1990).

2. What is a moral theory?

Utilitarianism is a moral theory. Indeed, it is one of the two main enlightenment moral theories, together with Kantianism. The aim of a moral theory, at least on the dominant conception thereof, is to provide “a theoretical account of the nature of right and wrong and thereby giving a moral criterion of right action.”² With other words, a moral theory aims to tell us in a systematic and coherent way what is morally right and wrong and why this is so.

As Jamieson puts it:

On the dominant conception, moral theories are abstract structures that sort agents, actions, or outcomes into appropriate categories. Proposed categories include virtuous, vicious, right, wrong, permitted, forbidden, good, bad, best, worst, supererogatory, and obligatory. Characteristically outcomes are ranked according to their goodness, actions according to their rightness, and agents according to their virtuousness. Different theories take different categories as primary.³

As we will see below, utilitarianism takes the goodness of outcomes to be primary, and evaluates actions and agents with respect to that central category. According to utilitarianism, whether an action is right or wrong depends on the goodness of its outcome. In general, all moral theories include an account of what is good and an account of what is right. The first, i.e. the theory of the good, is “a view about what is good or valuable, [...] which properties we ought to want realized in our actions or in the world more generally.”⁴ The second part of any moral theory,

² Timmons (2002): 131.

³ Jamieson (1991): 477.

⁴ Pettit (1991): 230.

i.e. the theory of the right, tells us “what individual and institutional agents should do by way of responding to valuable properties.”⁵

Since the nineteen seventies and eighties the dominant conception of the nature, status and role of moral theory has been challenged.⁶ Moral theories have been criticized for having unrealistic pretensions and for being useless for practical reasoning. Furthermore, the implicit ideals as well as the outcomes of moral theories have been criticized. Besides, there has been a lot of debate concerning the method of moral reasoning. More recently, and in reaction to that criticism, there have been explicit defenses of moral theory in general.⁷ The critics of moral theory have contributed to the reconsideration of many aspects of moral theories, but they have not established that moral theories are something we are better off without. As Jamieson puts it:

The anti-theorists remind us that people in their everyday moral practices create theory; that there are limits on what these theories can do; that their job is to help us do what is right rather than to be true. In these ways their attack on the dominant conception is important and helpful. However, what they do not succeed in showing is that we would be better off without moral theory.⁸

Not every defender of moral theories would agree with this summary of the state of the art. For instance, moral realists believe that theories should not only tell us what is morally right, but should also be true. Suffice here to note that there is controversy about the nature, aims and status of moral theory. On the dominant conception, a moral theory is a systematic account of what is morally right and wrong. Utilitarianism is a major moral theory, which will be further introduced in the following sections.

⁵ Idem.

⁶ See, for example Anscombe (1981), McDowell (1979), Stocker (1976), Williams (1985).

⁷ See, for instance Tännsjö (1998).

⁸ Jamieson (1991): 480.

3. Basic elements of utilitarianism

Utilitarianism is a consequentialist moral theory in the sense that the moral evaluation of actions depends solely on their consequences. The consequences of a particular action, and only its consequences, determine the action's rightness. There are no categories of actions that are obligatory or forbidden as such. Characteristically for utilitarianism, outcomes are evaluated in terms of their goodness or value. As already mentioned, for utilitarianism the category of the 'good' is primary. What is right is determined in terms of what brings about the most good.

This good that is ultimately valuable according to utilitarianism is itself a non-moral value. That means that the good is considered valuable in other than moral respects. It is about what is ultimately valuable for a being. As we will see, the sole ultimate value, according to utilitarianism, is thus a prudential value, namely welfare. Welfare is not considered good in moral terms (such as a 'good character' would be). Rather it is a prudential value. With other words, it concerns self-interest; it is concerned with what is good or bad *for someone*. The moral evaluation of an action depends on how an action bears on that good. In that sense, utilitarianism belongs to the teleological moral theories. Teleological theories can be defined as those theories that strive for the realization of a certain goal that is considered good in a comprehensive way. In a more narrow sense, teleological theories are defined as theories that morally evaluate actions solely in terms of how they bear on a *non-moral value*.⁹ One can consider the latter definition of teleological theories with a focus on non-moral value as a sub-category of the wider definition with a focus on the realization of value.¹⁰

⁹ Birnbacher (2006): 95-96.

¹⁰ Hübenthal (2006): 61.

The right thing to do, according to utilitarianism, is to *maximize* the good. Utilitarianism's 'theory of the right' entails that the right action is the one that maximizes the good:

Any particular (concrete) action is right if, and only if, in the situation, there was nothing the agent could have done instead such that, had the agent done it, the world, on the whole, would have been better. [...] an action is wrong, if, and only if, it is not right. And an action is obligatory if, and only if, had the agent acted in any way differently, the world, on the whole, would have been worse.¹¹

As utilitarianism is usually defined as a maximizing theory, it can be doubted whether satisficing versions, which claim that value should not always be maximized, should still be called utilitarian.¹² This is a terminological issue. I will stick to the common definition of utilitarianism as a moral theory that requires the maximization of the good.¹³ The requirement of maximization implies that outcomes of possible actions, so-called 'options', need to be compared. The option with the best outcome ought to be chosen.

The values of the consequences of actions are referred to in terms of 'utility'. Utility, more specifically, is used to refer to the net value of the outcome. Thus, if an action has both, good and bad aspects, the net utility refers to both kinds of aspects together. The guiding idea behind utilitarianism is the following: "[...] we are obligated to bring about the best state of affairs [...] that we can in the particular situation

¹¹ Tännsjö (1998): 31.

¹² See Jamieson (2009, p. 243) who makes a distinction between utilitarianism on the one hand and satisficing versions of consequentialism on the other hand. According to Jamieson utilitarianism "[...] requires the identification and pursuit of a single best outcome or one of a set of outcomes that contain the highest achievable value."

¹³ For a recent critique of satisficing utilitarianism, see Bradley (2006). For a defense, see Rogers (2010).

in which we find ourselves. In some situations, the theory will direct us to make the best of a bad situation.”¹⁴

Note that being able to determine the utility of an action in theory is sufficient. Utilitarianism does not necessarily require of an agent to calculate utilities before every action. Most likely, it does not. Instead, utilitarianism would require adopting a decision procedure that is most likely to maximize welfare.¹⁵ As will be explained in section 4.3 of this chapter, acts are not necessarily the only evaluands that a utilitarian can evaluate with respect to the maximization of welfare. Global utilitarians, for instance, propose to evaluate everything in that respect, thus naturally also decision procedures.¹⁶

We have seen that utilitarianism’s theory of the right asks us to maximize the good. This implies that central to utilitarianism is a notion of the ‘good’ or the valuable. But what is the good? In determining what is good or valuable, utilitarianism is focused on what is valuable for valuing beings: it focuses on prudential value. Their own welfare, in a very broad sense, seems to be a central value for valuing beings. The moral duty to neutrally maximize welfare is based on the acknowledgement of that prudential value. What is more, prudential value or welfare is the sole ultimate value according to utilitarianism. All else that is valuable can only be so because of its contribution to welfare. Hence, utilitarianism is a *welfarist* moral theory:

¹⁴ Timmons (2002): 105.

¹⁵ Bentham (1789, Chap. IV, Sec. VI), Mill (1861, Chap. II, Par. 19) and Sidgwick (1907, 413) have already pointed out the distinction between utilitarianism’s theory of the right and decision procedure.

¹⁶ For a detailed investigation regarding the choice of a decision procedure, see Ord (2005).

The most famous consequentialist theory is utilitarianism, which is the result of combining consequentialism with welfarism. Since consequentialism holds that an act is right if and only if it leads to the best consequences, and welfarism holds that the goodness of an outcome is ultimately a matter of the amount of individual well-being, counting everyone equally, it follows that utilitarianism is the view that an act is right if and only if it leads to the greatest total amount of well-being.¹⁷

Welfarism is central to utilitarianism. If welfarism were replaced by another theory of the good, holding that other things besides welfare or other things instead of welfare are ultimately valuable, the resulting theory would arguably no longer be utilitarianism. It would rather be another consequentialist theory.¹⁸ Again, that is an issue of terminology, and that is how it is usually defined.

Utilitarianism is both *universalist* and *impartialist* with regard to the issue of whose welfare counts in determining the value of outcomes.¹⁹ Universalism refers to the requirement that the welfare consequences for *everyone that is affected* by an action must be taken into account. The typical utilitarian understanding of impartiality is that ‘each counts for one and none for more than one’. That means that every affected being counts *equally*: the consequences of an action in terms of welfare for each affected being are taken equally into account.

¹⁷ Kagan (1998): 61. Kagan speaks here about the ‘total amount’ of welfare. Taking the total and taking the average, as we will see below, are alternative ways of aggregating welfare.

¹⁸ Kagan (2002): 61. However, this usage is not uniform, as even non-welfarist theories are sometimes called ‘utilitarian’.

¹⁹ Timmons (2002): 104.

To sum up, utilitarianism's basic elements are maximization, welfarism and universality/ impartiality. Utilitarianism requires the maximization of welfare, giving equal weight to each affected being's welfare. Even critics of utilitarianism have acknowledged a certain attractiveness of those features:

Utilitarianism captures three intuitively plausible ideas about morality. First, its commitment to *welfarism* accommodates our sense that morality has to do with human well-being. Second, it is based on a very plausible view of practical rationality. When it comes to rationality in the realm of choice and action, the idea that we ought to *bring about as much good as possible* seems irresistible. And finally, utilitarianism captures the idea that *impartiality* is at the heart of morality.²⁰

After this first sketch of utilitarianism and of welfarism, maximization and impartiality as its defining elements, let us take a closer look at that moral theory.

4. A closer look at utilitarianism

In what follows, I will further introduce the utilitarian moral theory. As utilitarianism is a welfarist moral theory, different versions of utilitarianism can be distinguished on the basis of different accounts of welfare (section 4.1). In order to determine which outcome scores best in terms of welfare, outcomes need to be compared. Their value in terms of welfare needs to be assessed. In order to determine the value of an outcome, utilitarians 'gather' (aggregate) all relevant welfare consequences in this outcome (section 4.2). In the business of maximizing welfare, utilitarians need not be focused on actions alone.

²⁰ Timmons (2002): 147. Italics mine. See also Kymlicka (1990): 12.

They might also evaluate rules or other evaluative focal points, such as motives or character (section 4.3). An outcome might be evaluated in terms of its quantity of welfare as such, or alternatively in terms of how much it benefits or harms sentient beings in terms of welfare (section 4.4). Last but not least, there are different views on why utilitarianism asks us to maximize welfare (section 5.6).

4.1 Defining welfare

As welfare is so central within utilitarianism, much of the plausibility of utilitarianism as a moral theory depends on what exactly is meant by ‘welfare’. It is not surprising that this crucial element of the utilitarian theory has received a lot of attention by both utilitarians and their critics. Some critics have found utilitarianism implausible, because they did not agree with the crucial role of what was presented as welfare. Utilitarians and others, who attach relevance to welfare, have brought forward various definitions of welfare in order to accommodate the critiques and make the account of welfare and thus the whole moral theory more plausible.

So, how should the crucial concept of welfare be defined? It is common to distinguish objectivist from subjectivist accounts of welfare. According to subjectivist accounts, welfare is at least partly a matter of mental states. According to objectivist accounts, welfare can be defined without reference to mental states.²¹ Objectivist accounts of welfare hold that whether and in how far someone fares well can be fully determined objectively, from the outside, so to speak, without any recourse to the perspective of the subject itself. For instance, according to Moore, welfare is determined by the possession of some things, which he considers intrinsically good. The degree to which one possesses those

²¹ See, for example Sumner (1996).

things determines one's level of welfare. The things on that 'objective list' might be health, wealth, friendship, or whatever. Usually at least what are considered primary goods are on the list. There are problems determining what counts as a primary good in an objective way. Even if one arrives at a list of criteria, the question remains whether this is a definition of welfare, or rather a list of criteria that are usually contributing to welfare. An important standard critique against objective accounts of welfare is that if someone has the goods on the list, such as health, wealth, and meaningful relationships, it is still an open question whether this person fares well, from her own perspective.

As Sumner explains,

Welfare assessments concern what we may call the prudential value of a life, namely how well it is going *for the individual whose life it is*. [...] [H]owever valuable something may be in itself, it can promote my well-being only if it is also good or beneficial *for me*.²²

Sumner's point that subject-relativity is and should be crucial for our concept of welfare is convincing, because welfare is about prudential value, hence it is about what is good for the subject. However, even if subject-relativity is accepted as being crucial, this needs not automatically lead to subjective accounts of welfare. Bogner has argued that Sumner equivocates on the concepts of subject-relativity and perspective: that welfare is about what is good *for a person* does not necessarily mean that welfare is about what is good *from that person's perspective*. There might be other ways of accounting for subject relativity. Furthermore, if perspective is taken to be crucial, it need not be a person's *actual*

²² Sumner (1996): 20.

perspective.²³ It might also be, for instance, a person's well-informed perspective, rather than his or her actual perspective.

It is an open question whether anyone who fulfils the criteria on the objective list is actually faring well, from her own perspective. According to Sumner, it is hard to see how any plausible account of welfare can be offered, without recourse to the subjective point of view. Some alleged objective accounts of welfare are really about something else: they do not offer an account of *prudential* value. For instance, a definition of welfare in terms of living according to one's excellences conflates prudential value and perfectionist value. Perfectionist value is about how much one lives up to certain standards of perfection.²⁴ Thus, while welfare is indeed a subject-relative concept, it is controversial what the role of the subjective perspective should be in the definition of welfare. Objective theories of welfare, while being subject-relative, define welfare without recourse to the subject's perspective.

Subjective theories define welfare with recourse to the subject's perspective. As far as subjectivist theories are concerned, it is common to distinguish between theories according to which someone's welfare is *solely* a matter of her states of mind and theories according to which someone's welfare is *additionally* a matter of some states of the world.²⁵ The most common accounts of welfare within utilitarianism belong either to *hedonism*, which is the classical utilitarian account of welfare, or to the *desire theory*.²⁶ Both are subjective accounts of welfare. Hedonism is a mental state theory that defines welfare in terms of subjectively experienced pleasure and pain or, rather more broadly, enjoyment and

²³ Bognar (2010).

²⁴ Sumner (1996): 78.

²⁵ Sumner (1996): 82.

²⁶ Hedonism, as a definition of welfare, is distinct from psychological hedonism, which is the claim that human beings always act in pursuit of what they think will give them the greatest balance of pleasure over pain.

suffering.²⁷ More precisely, substantive hedonism holds that welfare should be defined in terms of enjoyment and suffering, while explanatory hedonism holds that it is exactly the pleasantness of the pleasure and the painfulness of pain that make pleasure and pain good and bad respectively.²⁸ The desire theory defines welfare in terms of satisfied desires. The desire theory is not a mental state theory, because a state of the world is included in the definition: the desired state of affairs must *actually have been realized*. As the desire theory is about desires, it is still a subjective theory. After all, desires are attitudes of the subject.²⁹ The *present desire theory* defines welfare in terms of the fulfillment of a being's current desires. The *comprehensive desire theory* defines welfare in terms of a being's overall level of desire satisfaction in his or her life as a whole. This distinction is relevant with respect to the harm of death, and will come back in the following two chapters.

It is worth pointing out in what way the desire theory and hedonism coincide and in what way they differ. Substantially, a desire theorist and a hedonist might agree on what makes life good for a being: pleasurable experiences. However, the explanation why a pleasant life is good for a being differs for both accounts of welfare. According to hedonism, pleasantness is the good-maker, while the desire theorist will have to point to the fact that pleasantness is desired. This gives way to a major criticism that has been brought forward against the desire account of welfare. It has been pointed out that we desire things because we

²⁷ For recent defenses of hedonism see Bradley (2009): 8-18, Crisp (2006), Tännsjö (1998): ch. 5. For a recent critique, see Haybron (2008): ch. 4. Haybron defends a mental state account, according to which happiness is not only a matter of momentary feelings, but also, and on a deeper level, of moods and mood propensities. Haybron's mental state account of welfare might also be incorporated into a utilitarian moral theory.

²⁸ Crisp (2008).

²⁹ See, for instance Birnbacher (2005), Gesang (2000) and Crisp (2006) for defenses of subjective accounts of welfare.

consider those things independently good rather than considering things to be good because we desire them.³⁰

According to Sumner, subjectivism and objectivism are exclusive categories: if a theory makes welfare dependent (also) on the subject's perspective, it is subjective. If it does not, it is objective.³¹ Others have rejected this dualism and have argued that more fine-grained distinctions can be made. For instance, moderate subjectivism can be distinguished from pure subjectivism. Moderate subjectivism holds that in order count as positive welfare, an agent must not only have a positive attitude about something, but the condition that is the object of the positive attitude must also *obtain*. For instance, someone's perceived good health does not count as welfare, when the person, unbeknownst to him, is in fact very ill. Pure subjectivism rejects this condition. It seems that theories which accept the above-mentioned condition that the state of affairs that is valued must also actually obtain - the so-called 'reality condition' - are less subjective. There seems to be an objective element involved. Therefore, it has been argued, against Sumner, that the distinction between subjectivist and objectivist theories is in fact gradual.

Some 'subjective' accounts define welfare as the fulfillment of informed desires. These are not the desires an agent actually has, but the desires the agent would have, if she were well informed. It has been debated whether those are still rightly called 'subjective' accounts, or whether they assume an account of welfare that is rather objective.³² Bogner has argued that the distinction between subjectivist and objectivist accounts of welfare is not helpful at all. He has suggested that one should rather focus on the requirements that different theories of welfare do or do not accept, such as the above-mentioned 'reality

³⁰ Crisp (2008). See also Aristotle (1984 [C4BCE], *Metaphysics* 1072a, tr. Ross): "[...] desire is consequent on opinion rather than opinion on desire."

³¹ Sumner (1996): 82.

³² Sobel (2009).

requirement'.³³ Another example of a requirement is the so-called 'information requirement', as when it is required that a preference must not have been based on wrong information if its satisfaction is to count as contributing to welfare.

Some utilitarians have defended pluralist accounts of welfare. For instance, according to David Brink there are three main components of welfare. These are (1) reflective pursuit of one's reasonable projects, (2) realizations of those projects, and (3) certain personal and social relationships. Furthermore, according to Brink, all three components must be morally constrained in the sense that they only count as being part of welfare if they respect persons. Such a pluralist account of welfare is less simple and less determinate. The question that is relevant here is whether a consequentialist theory with such a pluralist account of welfare should still be called 'utilitarian'. In as far as such a theory would still define what is morally right in terms of impartial maximization of welfare, it seems to be a utilitarian moral theory. However, one can wonder, as for any definition of welfare, whether this one, including the three aspects plus the moral constraint, is a plausible definition of welfare.

4.2 Aggregation of welfare

From prudential goodness, the step must be made to the goodness of states of affairs: axiology. Utilitarianism determines the goodness of an outcome by bringing together the changes in prudential value that the outcome brings about. This is called 'aggregation'. Aggregation can take place *across individuals*. That means that the welfare consequences for all individuals are brought together (inter-personal aggregation). Aggregation can also take place *across time*, as when we want to determine

³³ Bognar (2010).

how the welfare that a person has at different times in her life determines how well her life goes as a whole (intra-personal aggregation). Aggregation can also take place across the dimension of states, if we want to know how the values of each of an act's various possible results (one for each state of nature) together determine the value of the overall outcome that the act leads to.³⁴

The most straightforward way of aggregating is summing up. Thus, for instance, the welfare of a whole life consists of the sum of all instances of positive and negative welfare. The value of an outcome consists of the sum of all positive and negative welfare consequences that the outcome contains. This method of aggregation is called the *total view*, because all relevant instances of welfare are simply brought together. An alternative method of aggregation is the *average view*. This method consists in summing up and then taking the average. So, the average welfare level of the person's life would indicate the value of the life for the person. The value of an outcome can be calculated by calculating the total amount of welfare and dividing it through the number of beings that are involved.

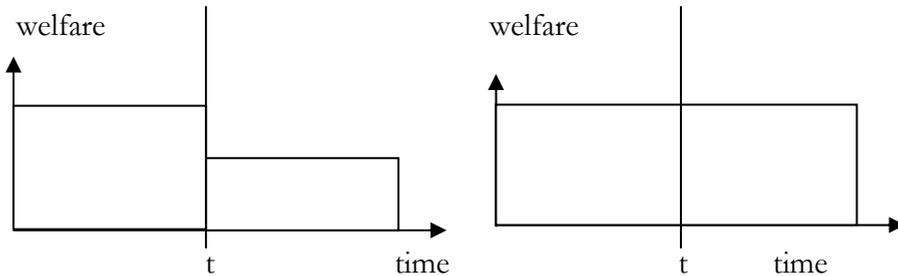
Let me illustrate how welfare calculations can be made, at least in theory. Every action leads to a certain distribution of welfare. For instance, the effect on the welfare of a woman of the choice of whether or not to have a child can be depicted schematically, as beneath.³⁵ The horizontal arrow depicts the time; the vertical arrow depicts the level of welfare. The vertical line depicts the moment of choice. The welfare to the right of the vertical line would be the consequence of the choice for the welfare of the woman. We can see that having a child slightly reduces the subsequent welfare level of the mother, maybe because caring for the child interferes with what she prefers to do. Thus the mother would be better off not having a child. In general, thus, this is how one can depict

³⁴ Broome (2004): 16-17.

³⁵ The above scheme is inspired by Broome (2004): 5.

a person's lifetime wellbeing in a graph. The graph begins as soon as the person has any wellbeing. Depending on one's definition of welfare, this might be when interests or sentience emerge. It ends with death, or as soon as the having of interests or sentience disappears.

Now, what is the value of the state of affairs in which the woman has the child as compared to the state of affairs in which she does not have the child? In order to assess that value, we need to know the welfare consequences for all that are concerned by that choice. Let us assume that these are only the parents and the possible child. The distribution of welfare depicts how much welfare each being has at each time. This can be depicted schematically:



Outcome A: the woman has a child at t .

Outcome B: the woman does not have a child at t

Figure 1: Woman's choice about having a child.

The choice can also be depicted in numbers. If being p exists at time t , the condition of p at t will be p 's welfare and it will be expressed as a real number: positive, negative or zero. If p does not exist at t , the condition of p at t will have some arbitrary non-numerical value, say Ω .³⁶ Imagine, following our example, that both parents would be slightly less happy when having the child, because they would have less time for their favored occupations. The child would have a happy life; it would be well

³⁶ Broome (2004): 25.

cared for and grow up under favorable circumstances. Here is the distribution of welfare for both outcomes:

Outcome A (not having the child):

		Time					
		1	2	3	4	5	6
	A	5	5	5	5	5	5
Person	B	5	5	5	5	5	5
	C	Ω	Ω	Ω	Ω	Ω	Ω

Outcome B (having the child):

		Time					
		1	2	3	4	5	6
	A	5	5	4	4	4	4
Person	B	5	5	4	4	4	4
	C	Ω	Ω	5	5	5	5

In the aggregation of welfare, it is assumed that welfare can be treated as an arithmetic quantity. It is quantifiable and is thus measurable. Lifetime welfare and temporal welfare are comparable between individuals and temporal welfare is comparable between times.³⁷

Note that within utilitarianism, the *distribution* of utility does not matter as such. So, it does not matter how exactly the welfare that is entailed by an outcome is distributed across persons and times. For instance, applying the total view as a method of aggregation, it does not matter whether a distribution of welfare among persons A and B is (3/3) or (1/5). The same holds for the average view. The distribution affects neither the total, nor the average. So, it is not taken into account as such. (It is taken into account though in as far as it affects welfare, for instance

³⁷ Broome (2004): 78

when people actually suffer from an unequal distribution. Then, this comes back in the numbers.)

Summing up, the value of states of affairs is defined in terms of their aggregate welfare. There are different definitions of welfare within utilitarianism. In this thesis, I will sometimes specify welfare in terms of ‘happiness’ and speak about ‘enjoyment’ versus ‘suffering’. I do this for reasons of style: it reads less abstractly and it fits the common way of talking about animal-friendly animal husbandry. In that public discourse, reference is made to ‘happy’ animals, ‘animal suffering’ and animals with ‘pleasant’ lives. Readers should keep in mind that adherence to a particular account of welfare is not necessary for the main argument in this thesis. My argument holds for *all* utilitarian accounts of welfare. Whenever it will be relevant, I will indicate the implications of different accounts of welfare.

4.3 Evaluative focal points

As we have seen, utilitarianism’s ‘theory of the right’ requires maximizing welfare. Now, the question rises, how, more specifically, should we go about maximizing welfare? What should be our focus? Should we always choose the *act* that maximizes welfare? Should we follow the *rules* that generally tend to maximize welfare? Should we adopt *motives* that are likely to maximize welfare? Should we develop a *character*, which is likely to help us maximize welfare? Should we do all of these?

In order to get a clear idea of what utilitarianism requires, we must be clear about the ‘evaluative focal points’. The term ‘evaluative focal points’ is from Kagan. It designates the focus of moral evaluation. ‘Objects of moral evaluation’ or ‘evaluative focal points’ are things on which we focus when trying to maximize welfare. Examples are acts,

rules or motives.³⁸ Evaluative focal points can be evaluated either directly or indirectly. Evaluating the focal point directly means to apply the principle of rightness directly to whatever focal point is to be evaluated. There are many possible forms of direct utilitarianism, depending on what is taken as legitimate evaluative focal points. There are direct utilitarian theories that accept only one evaluative focal point, for instance acts. There are also theories that accept several direct evaluative focal points. Indirect utilitarianism applies the theory of the right only to one central evaluative focal point. Other evaluative focal points are evaluated only indirectly in terms of their relation to the central one. Rule utilitarianism is an example of an indirect utilitarian theory. Rules are evaluated directly in terms of the goal of maximizing welfare. (“Choose the rule that maximizes welfare.”) Acts are judged right whenever they conform to the right rules.

Depending on what is considered a relevant evaluative focal point and depending on whether it is evaluated directly or indirectly, different versions of utilitarianism can be distinguished.³⁹ At this stage, it is not necessary to subscribe to one particular version of utilitarianism. I only want to show that utilitarianism is richer than many people are inclined to think. Simple act utilitarianism and rule utilitarianism are not the only respectable versions of utilitarianism. To the contrary, many counter-intuitive judgments, which are ascribed to utilitarianism, might disappear if other evaluative focal points are accepted along with acts.⁴⁰

³⁸ Louise (2006): 65.

³⁹ See Ord (2008) for an overview on this issue.

⁴⁰ Consider, for instance, a variation of the famous case in which a runaway trolley heads towards three people on the track. The only way to prevent those three people from being killed by the trolley would be to throw a huge person from a footbridge, which would thereby be killed and stop the trolley. Now, the right act might be to go for the maximum welfare and save the three people. However, “not wanting to kill anyone” might be the motive that prevents you from throwing the person from the bridge. This is a motive, which is generally good, in utilitarian terms; it generally helps to maximize welfare. So, a right utilitarian motive can lead to a wrong utilitarian act. If

4.4 Quantity of welfare or net benefits?

There are two rival views about the evaluation of outcomes: the Person-Affecting Restriction and the Impersonal View. Both views evaluate outcomes in terms of welfare, but in different ways. The opposition between these views on the evaluation of outcomes will play a central role later on in my evaluation of Prior Existence Utilitarianism. The Impersonal View focuses on the quantity of welfare that an outcome contains. The Person-Affecting Restriction focuses on the question in how far the outcome affects the welfare of sentient beings. The focus is on the question in how far sentient beings are made better or worse-off in an outcome. Thus, according to this view on the evaluation of outcomes, an outcome must always be evaluated as compared to another situation. It cannot be evaluated simply by its intrinsic features.⁴¹ The Person-Affecting Restriction might better be called the ‘Sentient Beings-Affecting Restriction’, because it is not only concerned with harms and benefits to persons, but to sentient beings in general. Yet, I will stick to the common label. The Person-Affecting Restriction implies that an outcome can only be better (worse) if it is better (worse) for one or more sentient beings.

An example can illustrate this distinction between both views on the evaluation of outcomes. Imagine that I contemplate whether it would be morally better for me to give a book that I have already read to a friend or rather to keep it on my shelf. If I give her the book, she will enjoy reading it, and she will learn something useful from it that will benefit her for the rest of her life. Giving it to her would cost me the

both, acts and motives, are accepted as legitimate direct evaluative focal points this situation would show evaluative conflict. This has been depicted as a weakness and as a reason for sticking to one direct evaluative focal point. But evaluative conflict is exactly what we experience in these kind of dilemmas. It has been suggested (Louise, 2006), and I agree, that evaluative conflict is not so unacceptable after all. It can be a strong point of a moral theory to acknowledge this and to be able to show exactly where it comes from.

⁴¹ Holtug (2003): 5.

minor inconvenience of carrying it with me, but it would give me the enjoyment of knowing that she would benefit from it. If I do not give it to her, she will do something less pleasurable instead of reading it and she will lack the good things that the new insights would have brought to her. I would lack the enjoyment of benefiting her, and I would save myself from the minor inconvenience of carrying the book with me. In this example, giving the book to my friend would bring about the outcome with the highest quantity of welfare: the welfare scores of all concerned would be higher. Giving the book to my friend would also bring about the outcome with the largest net benefits. Giving it to her would benefit her: she would be better off having the book. It would provide a benefit to myself, giving me the pleasant idea that I help her. Whether one evaluates the outcomes in terms of the quantity of welfare that they contain or in terms of net benefits does not make any difference in this case. Yet, these are two different ways of evaluating outcomes and in other cases they do make a difference, as we will see soon.

An implication of the Person-Affecting Restriction is that “One outcome cannot be worse (or better) than another *in any respect* if there is no one for whom it is worse (or better) *in any respect*.”⁴² This claim, also known as “the slogan”, has been described as the modern-day Ockham’s razor of moral reasoning.⁴³ Implicit adherence to the Person-Affecting Restriction has been attributed to different moral thinkers, such as Rawls, Nozick, Locke and Scanlon.⁴⁴ Some utilitarians have explicitly affirmed the Person-Affecting Restriction as a basic assumption. In that case, the

⁴² Temkin (1993a): 248. See also Arrhenius (2003): 186. This is usually interpreted in the narrow way, meaning that the outcome must be better or worse for a particular individual. However, it can only be interpreted in a wider way, meaning that the outcome must be better or worse for sentient beings, whoever they are. This distinction will be further explored in chapter 7.

⁴³ Temkin (2000): 133.

⁴⁴ Temkin (1993a) and (1993b).

Person-Affecting Restriction is meant to restrict the domain of what utilitarianism should be about. It is presented as a “restriction on the utilitarian principle”.⁴⁵ Maybe instead of a ‘restriction’ it should be understood as a particular interpretation of the utilitarian principle. This results in Person-Affecting Utilitarianism. The alternative to Person-Affecting Utilitarianism is Impersonal Utilitarianism.⁴⁶ Impersonal Utilitarianism strives for the outcome that contains the maximum quantity of welfare, rather than the maximum aggregate net benefit. Impersonal Utilitarianism evaluates outcomes by taking into account the welfare that is contained in an outcome, no matter whether sentient beings are harmed or benefited by it.

The Person-Affecting Restriction and the Impersonal View are fundamentally different views on what matters in the evaluation of outcomes. Should we be concerned with harms and benefits to sentient beings, or with maximizing an abstract value as such? The Impersonal View is an intrinsic aspect view, according to which the goodness of an outcome depends only on its intrinsic aspects and not on its relation with alternative outcomes. One only needs to count the welfare that is contained in the outcome. The Person-Affecting Restriction, in contrast, holds that “whether B is better than A depends on the precise relation between them, in particular, on who their members are or how they have come about.”⁴⁷ According to the Person-Affecting Restriction, one has to know whether and in how far sentient beings benefit or are harmed in an outcome. As I already mentioned, that necessitates a comparison. Therefore it cannot be decided on the basis of the outcome’s intrinsic aspects, i.e. its amount of welfare, alone.

One might wonder what the difference of those views is in practice. After all, the value to be maximized, according to Impersonal

⁴⁵ Glover (1977): 66.

⁴⁶ Parfit (1984): 387.

⁴⁷ Temkin (1999): 782.

Utilitarianism, is welfare. And welfare never comes in a free-floating way. It is always the welfare of someone. The practical differences come to the fore in ‘different people choices’. In different people choices, the Person-Affecting Restriction and the Impersonal View differ. Let me explain what characterizes those choices. In comparing any two acts, we can ask: Would all and only the same people ever live in both outcomes? If yes, we are faced with a same people choice. If no, we are dealing with a different people choice.

Let me illustrate the difference between the Impersonal View and the Person-Affecting Restriction in practice. Consider the question which of two outcomes should be brought about. Outcome A consists of a population with a very high welfare level (GREATER). Outcome B consists of an equally large population and an even larger population, both with a slightly lower welfare level (GREAT and ALSO GREAT).

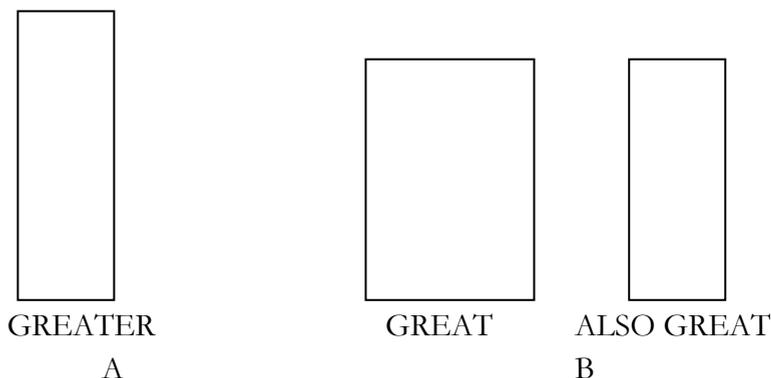


Figure 2: Greater versus Great and Also Great

If one considers only the intrinsic aspects of the outcomes, i.e. the total amount of welfare which they contain, B is better than A. On the intrinsic aspect view, B should be brought about rather than A. If one accepts the Person-Affecting Restriction, it depends. If, for instance, the question is whether the population in GREATER should give up welfare, resulting in ALSO GREAT, in order to have additional children

(GREAT), thereby bringing about B, the answer would be ‘no’. According to the Person-Affecting Restriction, B would not be preferable to A in that case. This is because it matters that the population GREAT consists of contingent beings and those would not be made ‘better off’ by being brought into existence.⁴⁸ In contrast, the GREATER population would be made worse-off by being required to transform into ALSO GREAT. This is a different people choice, because outcomes A and B do not consist of exactly the same people.

4.5 Underlying Argument for welfare maximization

We have seen that utilitarianism requires the maximization of welfare. Now, one can further inquire: “*Why* should welfare be maximized?” This is an important question to ask, in order to understand the driving force of utilitarianism as a moral theory. Different variations of utilitarianism can be more or less coherent with different rationales of what utilitarianism is all about. The rationale of utilitarianism can, for instance, indicate whether utilitarianism is ultimately concerned with individuals, or rather with welfare as such. This is relevant with respect to the criticism mentioned in chapter 1 (and also to be mentioned later on in this chapter), that utilitarianism is overly impersonal.

It is common to distinguish two underlying rationales for the duty of neutral welfare maximization.⁴⁹ One of those rationales is called ‘*equal consideration of interests*’. This means that the basic goal is to give all equally strong interests equal consideration, no matter whose interests these are, and no matter what the interests are. If this is taken to be the underlying idea of utilitarianism, then maximization “arises as a by-product of a standard that is intended to aggregate people’s preferences

⁴⁸ This is because the non-existent cannot be better or worse off. I will have to say more on this issue in chapter 6.

⁴⁹ Kymlicka (2002): 32.

fairly”.⁵⁰ Thus, the duty to neutrally maximize welfare can be conceived as arising from the underlying concern with equal consideration. Arguably, this basis of utilitarianism is implicitly affirmed by Mill and explicitly embraced by many contemporary thinkers, for instance Harsanyi, Griffin, Singer and Hare.⁵¹ The alternative possible rationale underlying the duty of neutral welfare maximization holds that the underlying concern is to bring about valuable states of affairs.⁵² The idea, then, is that welfare is intrinsically good and moral agents ought to maximize what is intrinsically good. G. E. Moore’s *Ethics* can be mentioned as a rare example of such a theory that is purely focused on an abstract non-moral good, without appealing at all to the ideal of equal respect for sentient beings.⁵³

Here, my concern is not an exegetical one, nor will I comment on the question of which rationale for utilitarianism is the more plausible one. I do only want to point out that the underlying rationale for utilitarianism does not need to have the same structure as the moral theory itself. Utilitarianism is commonly described as a teleological normative theory, because, as explained above, it defines what is morally right in relation to a non-moral good, namely welfare. Although utilitarianism is a teleological normative theory, it needs not necessarily have a teleological rationale.

At the factorial level, [...] a teleological theory is one that holds that the only basic normative factor is the good – whether the individual good (as in ethical egoism) or the overall good (as in consequentialism). Note that, in and of itself, teleology at the factorial level takes no stand

⁵⁰ Kymlicka (2002): 33.

⁵¹ Kymlicka (2002): 33. Kymlicka calls them ‘utilitarians’, however, it can be doubted whether and in how far these are all utilitarians. I will not determine this here.

⁵² Idem.

⁵³ Kymlicka (2002): 36.

whatsoever as to what it is that grounds or explains the relevance of this single factor.⁵⁴

The ground or explanation of the obligation to maximize welfare can be either teleological, i.e. focused on a non-moral good, or not.⁵⁵ Underlying rationales that are not focused on a non-moral good are usually called ‘deontological’. However, that is confusing, because rationales that do not focus on a non-moral good might still focus on the realization of a good and thus be ‘teleologic’ in a wider sense.⁵⁶ In order to avoid that confusion, I will from now on talk about the rationale of ‘equal consideration’ versus the rationale of ‘welfare maximization’ without using any overarching labels.

The distinction between different possible rationales is relevant with respect to the reproach that utilitarianism only cares for an abstract value rather than about real people. Utilitarians have already pointed out that welfare never comes in a free-floating way and always is the welfare of someone. Utilitarians that accept the rationale of equal consideration might argue that giving equal interests equal consideration is what it means to respect people (or sentient beings generally). The idea, then, would be that neutral welfare maximization happens to be the best way of giving the welfare of all sentient beings equal consideration and thus of showing them equal respect. This shows that equality might be a core value of utilitarianism, albeit not at the level of outcome. When outcomes are evaluated utilitarianism counts only quantity of welfare; the

⁵⁴ Kagan (1998): 191.

⁵⁵ Some might not like the term ‘foundation’ here, thinking that a foundation or foundational theory must be more than just an underlying argument or rationale. Kagan has a different, less pretending, understanding of ‘foundation’. I will however, unlike Kagan, rather speak about ‘basic idea’ or ‘underlying argument’ or ‘rationale’ in order to avoid confusion with different understandings of what a ‘foundation’ of a moral theory must be.

⁵⁶ Hübenthal (2006): 61.

way it is distributed across people does not matter. Proponents of the rationale of equal consideration can claim that equality is accounted for on a very fundamental level, as it provides the underlying argument for the duty to maximize welfare.

The distinction between different rationales of the duty to neutrally maximize welfare is important. It tells us why welfare should be maximized. As utilitarians strive to maximize welfare, it is often assumed that they must consider welfare valuable as such and this is the ultimate reason for maximizing it. However, there are other possible rationales for maximizing welfare. As we have seen, one might accept “equal respect for all sentient beings” as a rationale on which the duty of neutral welfare maximization is based. Later on we will see that one of the versions of utilitarianism that I will distinguish fits particularly well with the rationale of equal consideration.

5. A controversial moral theory

This brief introduction of utilitarianism shows that many different versions of this theory have been brought forward. Very often, new versions of utilitarianism have been proposed as refinements and improvements of that moral theory. This has usually been done in reaction to criticism. As we have seen, sometimes it is hard to tell which among that broad array of versions are still genuine utilitarian moral theories, and which are not. As indicated, this thesis will be concerned with utilitarianism understood as a theory that requires the *maximization* of *welfare*, giving *equal weight* to *each being's* welfare. More precisely, as will become clear later on, the focus will be on two different utilitarian views regarding the question whose welfare to count in the aggregation of welfare. In particular, the views to be discussed in this thesis divert on the question whether or not to take into account the possible welfare of contingent beings in the evaluation of outcomes. This controversy will be shown to be crucial for determining whether or not utilitarianism

should permit the killing and replacement of sentient beings under specified conditions.

Among all the criticisms that have been brought forward against utilitarianism, let me here only briefly discuss two lines of critique that are particularly relevant for the central issue of this thesis: One is the problem of knowing the long-term effects of actions, in particular the ‘cluelessness objection’, the other is the criticism of utilitarianism being overly impersonal, in particular the ‘separateness of persons’ criticism.

The ‘cluelessness objection’ dismisses the plausibility of utilitarianism, and indeed of any consequentialist moral theory as a guide to what we ought to do. The idea is that what one does is likely to have, or at least might very well have far reaching consequences that profoundly influence the future state of affairs in the world. Actions can be directly identity determining, as when one decides whether and when to have a child or to euthanize, abort, or kill somebody. Such actions determine not only whether a particular person will live, which is likely to have all kinds of impact on the future of the world. The actions also influence the possibility of identity-determining actions of that person, which in turn bears on future identity-determining actions, and so on. Besides, there are also *indirectly* identity-determining actions. For instance, many actions and policies influence who is where and when and thus, indirectly, who is having children together, and when. Next to directly or indirectly determining who will live, for generations to come, actions, even seemingly insignificant ones, might have huge and unforeseeable effects on the course of the world. Think about the famous ‘butterfly effect’, first noticed by Edward Lorenz, who found an extreme sensitivity of weather prediction models to initial conditions. Given all that, it seems that only a fraction of an action’s consequences are knowable to us, either *ex ante* or *ex post*. Before the action we simply cannot know all its invisible consequences. And even when we act on the

basis of past experiences, this will not help, because we cannot know the invisible consequences of like actions that have been done in the past.⁵⁷ This is because invisible consequences are per definition invisible. Unlike merely unpredictable consequences, they remain invisible after the action.

This cluelessness does not only regard the choice of actions, but also the choice of rules, character, decision-procedures, and other evaluative focal points.⁵⁸ Furthermore, in response to that epistemological problem it would not help to focus on intended or foreseen consequences (subjective utilitarianism) rather than on actual consequences (objective utilitarianism). Neither would it help to focus only on proximate causes of actions (as is done in law) or to adopt any other decision procedure.⁵⁹ This is because in the absence of any idea about the relevant effects of an action it is impossible to make a link between a decision procedure and those effects.⁶⁰

Both views to be discussed in this thesis take into account effects on the welfare of beings that do not yet exist when the moral choice in question is made. Thinking about effects of our actions on those who will exist or might exist in the future might bring worries along the lines of the cluelessness objection to mind. How do we know how our actions affect the (far) future? Therefore, I will say a few words about how utilitarians and others can deal with, and have actually dealt with that objection. This will show more precisely what is at stake.

The issue is not about the effects that are knowable, either *ex ante* or *ex post*. The issue is precisely about effects that are unknowable,

⁵⁷ Lenman (2000) has most forcefully brought forward the cluelessness objection.

⁵⁸ Lang (2008).

⁵⁹ Concerning the focus on proximate causes, see Hart and Honoré (1985). The idea is, briefly, that “[...] when voluntary acts and coincidences intervene in certain causal chains, then the results are not seen as caused by the acts further back in the chain of necessary conditions (Sinnot-Armstrong, 2006: 14).

⁶⁰ Lenman (2000): 359-362.

so-called ‘invisible effects’. These are the possible effects of an action about which one virtually has no clue. For instance, when I donate 100 Euros to Oxfam International, I might know that at least part of that money will be put into a particular project, which is likely to benefit some people, but doing so might also have all kinds of invisible effects which nobody is able to foresee. Maybe it will save the life of a child a far descendent of which will become a cruel dictator. That dictator might kill many people, which might or might not turn out to be bad on the long run, depending on what would have happened otherwise, which is impossible to know. In response to that criticism, it has been pointed out that there seems to be no reason to assume that the invisible effects of my sending 100 Euros to Oxfam International will be worse than the invisible effects of my doing something else instead. Remember that we are talking here about the invisible effects, i.e. those effects about which by definition one lacks *any clue*.⁶¹ It has been suggested, and it actually seems plausible, that the invisible effects of different options cancel each other out.⁶² For instance, in the absence of any clue to the contrary, it seems just as likely that one option will bring about a dictator 500 years from now, than that the other option will. This is called the ‘balancing out hypothesis’.

Even though the balancing-out hypothesis seems plausible, skepticism about it is possible. The cluelessness objection against consequentialist moral theories is built on the skeptical scenario that the balancing-out hypothesis is false. The epistemological skepticism concerning the balancing-out hypothesis claims that even though it seems plausible, it might be false, and there is no way of knowing this.

Skepticism about the balancing-out hypothesis parallels the epistemological skepticism about metaphysical realism, which claims that even though the external world appears to us in a certain way, that

⁶¹ Dorsey (2010).

⁶² Hare (2008): 19-20.

appearance might be false. We might, for instance, just be brains in a vat, and there is, by definition, no way of knowing it. It has been argued that all ways of dealing with metaphysical skepticism are equally available for dealing with skepticism about the balancing-out hypothesis. In fact, it has been argued that “[i]f one rejects consequentialism on the basis of skeptical scenario’s, one has no ground to retain metaphysical realism in the face of skeptical scenarios.”⁶³

I will not explore this argument here, but I hope to have placed the cluelessness objection in an interesting perspective. It seems that the routes that are open in the light of possible skepticism about the balancing-out hypothesis parallel the routes that are open with regard to skepticism about the existence of the outer world. In case of consequentialism, choosing one of those routes for dealing with skepticism amounts to doing what consequentialists are used to do: judging the moral quality of actions and other evaluands on the basis of their visible consequences.⁶⁴

The second often rehearsed critique against utilitarianism that I want to mention here is known as the ‘seperateness of persons’ criticism. That criticism is usually traced back to John Rawls, who in *A Theory of Justice* brought forward the following concern:

The most natural way, then, of arriving at utilitarianism [...] is to adopt for society as a whole the principle of rational choice for one man. [...] Utilitarianism does not take seriously the distinction between persons.⁶⁵

Many others, among whom Rawl’s philosophical antagonist Robert Nozick have made the same point:

⁶³ Dorsey (2010): 26.

⁶⁴ Dorsey (2010) presents this argument in some detail.

⁶⁵ Rawls (1971): 26-27. See also Gauthier (1962, p. 126) and Nagel (1970, p. 134).

There is no social entity with a good that undergoes some sacrifice for its own good. There are only individual people, different individual people, with their own individual lives. Using one of these people for the benefit of others uses him and benefits the others. Nothing more. [...] Talk of an overall social good covers this up. (Intentionally?) To use a person in this way does not sufficiently respect and take account of the fact that he is a separate person, that his is the only life he has.⁶⁶

The concern that underlies the ‘separateness of persons’ objection seems to be that individuals have rights, which at least sometimes trump utility calculations and/ or that their interests should not simply be traded off against each other. The concern seems to be that utilitarianism does not seem to sufficiently respect persons.⁶⁷

The Replaceability Argument, which will be explored in this thesis, has been considered one of the most abhorrent implications of utilitarianism, from the perspective of the ‘separateness of persons’ objection. The Replaceability Argument allows the killing and replacement of a being, provided that the overall amount of welfare does not diminish. This flies in the face of those who are concerned about utilitarianism being overly impersonal. Treating sentient beings as replaceable seems to prove a concern with welfare as an abstract quantity, rather than with sentient beings. The being that is killed gains nothing from the fact that another being will live in its place. In the words of Nozick, by being killed the being loses the only life it has.

In the field of animal ethics, the Replaceability Argument has been criticized along the lines of the ‘separateness of persons’ objection. For instance, Tom Regan criticizes that utilitarianism considers animals

⁶⁶ Nozick (1974): 32-33.

⁶⁷ See Norcross (2008) for an analysis of and a reply to this objection.

to be merely “replaceable receptacles”.⁶⁸ Regan’s major criticism against Singer’s utilitarianism is that it views an animal as “only a receptacle of what has value (i.e. preference satisfaction), lacking any independent value of his own”.⁶⁹ Also many others have criticized the “inadequate individual protection” and the “killing of innocents”, in particular with regard to the Replaceability Argument, where the killing is not justified in order to prevent a major disaster, but where the “gain would be quite small”.⁷⁰

The central discussion in this thesis will have implications for the question how impersonal utilitarianism is and ought to be. I will present a version of utilitarianism that avoids the Replaceability Argument.

6. Utilitarianism on animal husbandry

After this introduction of utilitarianism, I will now embark on my exploration of utilitarianism’s support of animal-friendly animal husbandry.

As we have seen, utilitarianism implies that we have moral duties towards animals. We have the duty to take their welfare equally into account in determining what we morally ought to do. In that sense, utilitarianism grants moral status to animals: animals matter morally. For instance, what we do unto a dog is considered not only relevant in as far as it affects the owner of the dog; it is also considered relevant because of its effects on the dog him or herself. Thus, according to utilitarianism, it matters morally how our actions affect sentient beings, including animals.

⁶⁸ Regan (2004): 210.

⁶⁹ Regan (2004): 210.

⁷⁰ Pluhar (1995): 182

As explained in chapter 1, section 4, according to utilitarianism saying that a being has to be included in our moral consideration does not mean that the being is insured of a particular treatment. It is, for instance, not the case that anybody, neither human nor nonhuman, is accorded any 'unalienable rights', such as the right to life or the right to bodily integrity. In fact, what might be done unto anybody depends on what maximizes overall welfare. Including a being in moral considerations can mean completely different things according to different moral theories. In case of utilitarianism, what is granted is only that the being's welfare is taken equally into account in the evaluation of outcomes. This fact sets no limit whatsoever as to how such a being might be treated. How a being ought to be treated depends on what maximizes welfare. This holds at least at the critical level of the moral theory. In daily life, utilitarians might accept the recognition of rules and rights that determine what one might do unto others, simply because doing so maximizes welfare in all but exceptional cases.⁷¹

As far as the use of animals is concerned, its permissibility depends on the consequences in terms of welfare, as compared to alternatives. I have already introduced as one of utilitarianism's central assumptions that all sentient beings count equally as moral objects. According to utilitarianism animal welfare counts as much as human welfare. Thus, for example, the seriousness of suffering is typically determined by the intensity and duration of the suffering.⁷² It is not determined by who the suffering being is.

Let me illustrate the considerations that are necessary from a utilitarian perspective in order to morally evaluate the practice of animal

⁷¹ Here, I refer to the distinction between utilitarianism as a theory of the right on the one hand, and what would be the best utilitarian decision procedure on the other hand. In this regard, Hare (1981) distinguishes the critical level of moral thinking from the intuitive level of moral thinking. On this latter level, one should follow rules of thumb, which tend generally to maximize welfare.

⁷² Singer (1995), see also Bentham (1963).

husbandry. Imagine that a legislator had the power to choose whether or not to allow intensive animal husbandry. Allowing it would cause a lot of animal suffering. It would also cause some suffering of humans who are opposed to intensive animal husbandry. Furthermore, it would cause some enjoyment for humans who like to eat the animal products. It is an open question, how the balance between all positive and negative welfare consequences of the action would be struck. This outcome, whatever it is, would need to be compared to the outcomes of alternative options. For instance, the legislator could choose to rule out intensive animal husbandry and allow more animal-friendly animal husbandry. Possibly, many positive welfare consequences would be safeguarded, while many negative ones would be avoided. If this would be the case, so-called animal-friendly animal husbandry would be the morally better option.⁷³

On the basis of these sorts of considerations, it has been argued that at least intensive animal husbandry is a major source of avoidable suffering.⁷⁴ There are alternatives that score better in terms of overall welfare. Therefore, utilitarianism generally condemns intensive animal husbandry. So-called animal-friendly animal husbandry seems to be preferable, from a utilitarian perspective. After all, roughly the same benefits can be realized with less suffering.⁷⁵ A crucial question would be how much suffering this form of animal husbandry still implies, as compared to alternatives, including plant-based ways of feeding people.

What about the killing of animals? We have seen that causing an animal to suffer is morally problematic. But what if the killing does not cause any suffering? How should killing be evaluated, apart from the eventual suffering that it might cause? The question whether and in how

⁷³ Of course, there would be more than just the above-mentioned consequences to be taken into account. In addition, more than just those two above-mentioned options might exist.

⁷⁴ Rollin (1995).

⁷⁵ A full exploration of the choice would be complex, but the idea should be clear, I suppose.

far an animal is harmed by death will be explored in the following chapter. Let us assume for a moment that death does not harm the animal. Does this mean that killing an animal is morally unproblematic? That depends on which view on the evaluation of outcomes one assumes. According to the Impersonal View, which has been introduced in section 4.4 above, an outcome should be evaluated in terms of the quantity of welfare, which it contains. According to the Impersonal View, killing an animal that would otherwise have had positive welfare is morally problematic, even if it does not harm the animal. It is morally problematic, because it reduces the overall amount of welfare that the outcome contains. Quite simply, letting the animal live out its pleasant future would result in more overall welfare than killing the animal. In contrast, on the Person-Affecting View on the evaluation of outcomes, outcomes should be evaluated in terms of harms and benefits to sentient beings. On that view, if death would not harm an animal, killing that animal would not be morally problematic. After all, as the killing would harm nobody, the outcome in which the animal was killed would not contain more harm. Let us now assume for a moment that death harms an animal. In that case, it would not matter whether one accepts the Impersonal View or the Person-Affecting View about the evaluation of outcomes. Both views would consider killing that animal morally problematic: the first because, again, killing reduces the overall amount of welfare in the outcome, and the second because killing results in an outcome which contains more harm.

So, killing is not morally wrong as such. It depends on the effect of the killing on the value of the outcome. Now, the question is how the death of a being affects the overall value of the outcome. This depends on which method of aggregation is used: the total view (summing up) or the average view (taking the average).⁷⁶ In general, if the number of beings is the same in both outcomes, the outcome with the greatest total

⁷⁶ Those views have briefly been introduced in section 3.1 of this chapter.

will also be the one with the greatest average. However, both methods of aggregation can yield different results when the number of beings differs in both outcomes. This is the case when the choice is about killing a being.

Let us consider killing an animal from the perspective of the average view on aggregation. If the welfare of the killed animal is beneath average, killing the animal raises the average and would, all else being equal, be morally required. Thus, it becomes crucial what the average is. For instance, is it the average welfare level of all sentient beings? If that is the case, is it the average welfare level of all presently existing sentient beings, or also of those that existed until now, or even that of future beings? An implication of the average view is that the best outcome is the one in which all are dead, except for the happiest being. Such an implication seems at odds with the basic idea of utilitarianism, be it the rationale that welfare should be maximized because it is intrinsically good, or the rationale that welfare should be maximized because everybody's welfare matters equally. Therefore, I will not be concerned with the average view on aggregation any further in this thesis.

Let us now consider killing from the perspective of the total view as a method of aggregation. Killing an animal which otherwise would have had a pleasant future, is likely to count as a loss of welfare, either because it reduces the quantity of welfare in the outcome, or because it adds to the harm in the outcome, or both. If the killing of an animal makes the outcome worse in one of those ways, then, given the total view as a method of aggregation, utilitarianism forbids the killing of animals that would otherwise have had pleasant lives. That is, utilitarianism rules out killing those animals, *unless* the welfare loss can be compensated. This 'unless' is important and I will say more about it soon. First, I want to discuss the killing of animals that would otherwise have had unpleasant lives.

If an animal dies that would otherwise have had an unpleasant future, the animal's death avoids negative welfare. The death of the animal in that case would, *ceteris paribus*, have positive consequences in terms of welfare. (The *ceteris paribus* clause implies, for instance, that negative effects on others need to be taken into account.) The death of the animal would amount to the prevention of a certain amount of uncompensated suffering. The killing of such an animal might yield the best consequences. This does, of course, not mean that people are permitted to kill animals in animal husbandry, if only they make their lives miserable enough. People should, *ceteris paribus*, not cause animals to suffer, because suffering diminishes welfare. Furthermore, as we have seen, people should not kill an animal, unless the animal would *unavoidably* experience uncompensated suffering. If possible, people should avoid making the animal suffer, for instance by treating it in a better way.

How can the welfare loss that is caused by the killing of an animal that would otherwise have had a pleasant future be compensated? As explained, killing an animal is required if this is the best or only way to prevent greater welfare loss. This holds, no matter whether the welfare loss would be constituted by the future suffering of the animal itself or by the suffering of others. For instance, killing a cat would, *ceteris paribus*, be required if it were the only or best possibility to save the lives of two other cats with together at least an equal amount of welfare. Killing an animal is also allowed if the welfare loss is compensated by the positive consequences that the killing has for others. If others would gain a lot of welfare from the killing, that welfare can outweigh the welfare loss for the animal. However, as always, this only holds if there are no alternative options that score better in terms of welfare.

What does this imply, with respect to animal husbandry? We have seen that animal suffering puts negative weight on the utilitarian scales. But even the painless killing of happy animals, as intended in animal-friendly animal husbandry, has the effect that the outcome

contains less welfare than it would have contained if the animal had continued to live a pleasant life. True, painless killing does not inflict suffering on the animal. Whether it causes any harm to animals at all will, as I already mentioned, be discussed in the next chapter. In any case, killing an animal that would have had a pleasant future implies a welfare loss *as compared to the option of letting the animal live on*. The welfare that the animal would otherwise have experienced is not realized. Now, the crucial question becomes whether this welfare loss can be *compensated*. I have mentioned two ways of compensating a welfare loss, caused by killing an animal that would otherwise have had a pleasant future:

- a) The killing is the only or best way to achieve greater positive consequences on overall welfare.
- b) The killing is the only or best way to prevent greater negative consequences on overall welfare.

Is one of these things the case in animal-friendly animal husbandry?

Let me check for the first above-mentioned justification of killing a happy animal, justification (a). Evaluating all consequences of the practice of animal husbandry as compared to alternatives, in terms of welfare, would be an important empirical task. Much has already been explored about all kinds of negative consequences of animal husbandry, for instance with regard to the environment, human health, and world food distribution. Here, it should be noted that in most parts of the world people do not use and kill animals for *food*, strictly speaking. Whenever and in as far as there are fully nutritious and safe alternatives for the consumption of animal products, animals are not used and killed for food, but for pleasure. Those who counter that food is more than nutrients alone, should at least admit that in as far as the nutrient part of food can be provided by non-animal means, the choice for animal

products does not serve the nutrient aspect of food, but rather the taste, i.e. the pleasure aspect of food. Interestingly, many of those who have erased animal products from their diets have gradually developed their tastes accordingly, and/ or have been compensated for loss of taste-induced pleasure by pleasures related to better health and ‘conscience’.⁷⁷ That is not to say that nobody would experience a vegan diet as a loss of personal pleasure. However, it is unlikely that those losses, if they persist, would outweigh the overall benefits of abandoning animal husbandry. More about this will be discussed in chapter 9.

Another relevant issue here concerns the link between the cost benefit analysis of animal husbandry and its alternatives on the one hand and the choice of a single moral agent on the other hand. After all, the options for a single moral agent do usually fall short of deciding about the continuation or eradication of animal husbandry. A single farmer, consumer or citizen can at most make a humble contribution to those states of affairs, if at all. Questions about individual responsibility in the face of limited and unclear individual influence occur in many domains. They are particularly pressing for utilitarians, who morally evaluate individual actions in terms of the difference that they make. Here, however, I will focus on the more principled issue of whether there is any compensation for the welfare loss that is caused by killing animals that could otherwise have continued to lead pleasant lives.

Let me now turn to justification (b), which states that the utility loss can be compensated if the killing is the only or best way to prevent greater or at least equally great negative effects on overall welfare. In animal husbandry, animals are not killed in order to avoid any greater

⁷⁷ See, for example Mohandas Gandhi’s “Diet and Morality”: “[...] there is something much higher which calls us to vegetarianism [...] I would just emphasize the moral basis of vegetarianism. And I would say that I have found from my own experience, and the experience of thousands of friends and companions, that they find satisfaction, so far as vegetarianism is concerned, from the moral basis they have chosen for sustaining vegetarianism.”

harm for themselves or for others. Neither does killing the animals prevent any unavoidable suffering. In intensive animal husbandry, the animals might have miserable lives, but then this is due to how they are treated and not unavoidable. In animal-friendly animal husbandry, it is assumed that the animals have pleasant lives. So, killing the animal would not prevent welfare loss but rather make impossible the further experience of positive welfare.

So, let us assume that both justifications for killing animals do not hold in the case of animal husbandry. I have explained that killing animals is not categorically ruled out. Complex calculations are necessary in order to determine whether killing animals is justified. I have suggested that the welfare benefits that are caused by killing animals in animal husbandry cannot outweigh the welfare loss that killing the animals implies. Killing the animal makes impossible the experience of welfare that the animal would otherwise have experienced. Therefore, killing the animal causes a welfare loss, as compared to the alternative option of letting the animal live. Furthermore the animals are not killed in order to prevent greater negative consequences on overall welfare. If this is right, we need another utilitarian justification for killing animals in animal husbandry, if we want utilitarianism to support this practice.

What can this justification be? For an action (including omission) to be morally permissible, it must at least score equally well in terms of welfare consequences, as alternative options. Killing a happy animal is a welfare loss, as compared to not killing it. So, it is impermissible, unless the loss is at least compensated. Now, we have seen that in the case of animal husbandry, possible ways of compensation as discussed under (a) and (b) might not do the job. Let us assume that this is indeed the case: the welfare loss due to the killing of happy animals is not compensated by any sufficiently high welfare gain, nor is it necessary in order to prevent an even higher welfare loss. We seem to be stuck with a welfare loss. Are we?

Another way of compensating for the welfare loss that is caused by killing happy animals in animal-friendly animal husbandry seems to be available. We have learned that utilitarianism is concerned with welfare, and it does not matter *whose* welfare it is. The welfare loss of a being can be compensated by welfare gains *of another being*. Now, in animal husbandry, a new animal usually replaces every animal that is killed. The existence of this new animal in a way depends on the killing of the first. The whole idea of animal production is to turn animals into animal products. This holds obviously in the case of meat production. Also in case of the production of dairy and eggs, animals are used to produce animal products, and then killed. It is only because dead animals leave the process, that new animals get in. The killing of each animal goes together with the production of a new animal, which would not otherwise have existed. It has been suggested by utilitarians that adding this new happy animal can be seen as an addition of a certain amount of welfare, which can compensate for the welfare loss that is caused by the killing of the former animal.⁷⁸

This argument is called the Replaceability Argument. It presents an alternative way of compensating for the welfare loss caused by the killing of a happy animal. The Replaceability Argument will be explored in chapter 4. The following chapter will explore the question that has been left unanswered above. This is the question whether and in how far animals are harmed by death.

7. Conclusion

Utilitarianism is a controversial moral theory. It does not accept that any action as such can be wrong: all depends on the consequences. Utilitarianism accepts no *a priori* limits of what one can do onto others.

⁷⁸ Singer (1993): ch. 5.

Whatever maximizes welfare is allowed, and even required. This has led to the criticism that utilitarianism is not really concerned with individuals, but rather with an abstract notion of welfare as such. The central theoretical discussion in this thesis is relevant for this issue. The two versions of utilitarianism that will be explored in this thesis differ in terms of how impersonal they are.

The sole ultimate good or value, according to utilitarianism, is welfare. There are different accounts of welfare commonly distinguished as ‘subjective’ and ‘objective accounts’. In this thesis I will understand welfare as a state of mind. I will speak about welfare in terms of ‘happiness’, and talk about ‘suffering’ versus ‘enjoyment’. I do this for reasons of style and personal preference. The argumentation in this thesis does not hinge on a particular account of welfare. The step has to be made from the welfare of the individuals to the value of the state of affairs. Thus, the welfare consequences for all concerned have to be put together, aggregated, in order to determine the value of an outcome. Two different methods of aggregation are the total view, which simply sums up the welfare of all concerned, and the average view, which sums up the welfare that is contained in an outcome and then takes the average by dividing it through the number of affected beings.

Utilitarianism’s theory of the right requires that the aggregate of prudential value should be maximized. As we have defined it:

Any particular (concrete) action is right if, and only if, in the situation, there was nothing the agent could have done instead such that, had the agent done it, the world, on the whole, would have been better. [...] an action is wrong, if, and only if, it is not right. And an action is obligatory if, and only if, had the agent acted in any way differently, the world, on the whole, would have been worse.⁷⁹

⁷⁹ Tännsjö (1998): 31.

Utilitarianism can evaluate other focal points, besides actions, in terms of whether they maximize welfare. For instance, it can also evaluate rules, motives or character. Any evaluative focal point can be either directly evaluated in terms of welfare maximization, as in “Choose the act that maximizes welfare!” It can also indirectly be evaluated, as in “Choose the act which is in line with the best character, which is the one that maximizes welfare.” In that example, character would be a direct evaluative focal point, and acts an indirect one. It is possible to accept several direct evaluative focal points, such as acts, motives and character.

Utilitarianism’s rationale explains *why* welfare ought to be maximized. Utilitarianism, as far as its theory of the right is concerned, is a teleological theory, because it defines what is morally right solely in relation to a non-moral good. However, this does not mean that the *underlying rationale* for welfare maximization must have a teleological structure as well. A possible rationale might be fairness, equality or respect. According to such a rationale, welfare ought to be neutrally maximized, because that is the best way to give equal consideration to equal interests.

With respect to animal husbandry utilitarianism strongly supports the idea that animal welfare matters. Causing suffering puts negative weight on the utilitarian scales, no matter whether the suffering being is a human or an animal. Choosing the option that causes most welfare for all concerned is the right thing to do. Therefore, intensive animal husbandry would not be sanctioned by utilitarianism.

What about the killing of animals? Utilitarianism does not categorically condemn the killing of animals. The average view as a method of aggregation would require killing all sentient beings, except for the happiest one, as this raises the average level of happiness. From the perspective of the total view killing an animal that would have had a pleasant life is a loss of welfare. This is because the welfare that the animal would have experienced is lost. In case of animal husbandry the

enjoyment that is gained from using the animals cannot compensate for the welfare loss that is inflicted on the animals. Alternatives are available that score better in terms of both animal and human welfare. Neither is the killing necessary in order to avoid greater welfare loss to the animal or others.

There is an argument that specifies conditions under which the welfare loss due to the killing can be compensated and the killing can be justified. We have seen that the welfare loss caused by the killing can in principle be compensated by the utility that is gained by other individuals. In animal husbandry, a new animal that would not otherwise have existed replaces every killed animal. The idea is that this can compensate for the welfare loss due to the killing. This argument is known as the Replaceability Argument.

If my argumentation so far is sound, utilitarianism's support of animal-friendly animal husbandry rests on the plausibility of the Replaceability Argument. It is therefore worthwhile to explore the Replaceability Argument in some detail. In chapter 4, I will present the Replaceability Argument and I will show that it depends on a view that is controversial within utilitarianism.

3. ARE ANIMALS HARMED LESS BY DEATH?

1. Introduction

In chapter 1, I have mentioned that even if animals are accorded equal moral status, treating them differently can be justified if they have different interests. Within utilitarianism, granting all sentient beings equal moral status means that their welfare is taken equally into account. Imagine, for example, a woman that has a child and a dog. Taking their welfare equally into account means to provide them with the food and shelter they need. Taking their welfare equally into account is compatible with sending the child to school, but not the dog. The reason is that the dog would not benefit from being sent to school. So, different interests can justify different treatment.

Can this fact about utilitarianism provide the clue for a defense of animal-friendly animal husbandry? One might argue that both humans

and animals benefit from not being caused pain. However, one might continue to argue, painlessly being killed harms humans, but not animals. If that is the case, that might explain why we are not allowed to kick an animal, while we are allowed to (painlessly) kill it. In this chapter, I will explore the utilitarian stance on the question whether animals are indeed harmed less by death. We will see that it is controversial. I will discuss the main utilitarian positions on the harm of death and explore what those positions imply for the question whether animal-friendly animal husbandry can be justified.

We are concerned with the harm that death is for the being that dies. Effects on others are not considered. In order to capture the harm of death as such, one should imagine that the being is killed painlessly while asleep. In that way, fear of death and other otherwise important side effects that death might have are left out of consideration. Of course, those things count in a complete utilitarian calculus, but not in an account of the harm of death. Another thing that should be noted is that a view about the harm of death is not the same as a view about the wrongness of killing. When considering the wrongness of killing, considerations other than the harm of death need to be taken into account. For instance, killing can have effects on others. Furthermore, for laws about killing, certain clear-cut rules rather than gradual distinctions might be preferable on utilitarian grounds, even if the harm of death is a gradual matter. Therefore, any view concerning the harm of death does not translate directly into a view about the wrongness of killing, nor does it imply any legal position on how killing should be regulated or punished.

A prominent view is that death is a lesser harm for animals, because animals do not have many desires for the (far) future. Animals, according to that view, lack a conception of their identity over time and therefore they do not have any desire to go on living. Hence, on this view, death is a lesser harm for animals (section 2). Rather than defining the harm of death in terms of the loss of what a being *wants*, it can be

defined in terms of the loss of *what would have been valuable for the being* (section 3). Arguing that the loss of future value for the being must be discounted for the lack of psychological connectedness between the being at the time of death and its future self at the time the value would have occurred to it can modify this latter view on the harm of death. This modification is controversial (section 4).

I will conclude that none of those major views on the harm of death that are compatible with utilitarianism implies that death is harmless for the animals that usually end up on people's plates. However, different accounts yield very different conclusions about the harm that death is for those animals. Utilitarians can deny that animals are harmed less by death. Even if they accept that animals are harmed less by death, death is still harmful for animals to some degree (section 6).

2. The relevance of future-oriented desires

One influential position among the desire-fulfillment conceptions of welfare is that a being is harmed more by death if more of its desires remain unsatisfied. The more future-oriented desires a being has when it dies, the more it is harmed by death.¹ This is because harm is defined as the frustration of desires. If no desires are frustrated, a being is not harmed. The relevant desires can be specified in different ways. In general, the desires whose satisfaction is relevant for welfare need not be consciously held.

Beings that have a conception of their own existence over time usually do have plans and projects for the nearer and further future and desire to go on living. Beings that lack a conception of their own existence over time lack a desire for continued life. They only have immediate and short-term desires, such as the desire to escape

¹ Singer (1993).

frightening or painful situations, the desire to eat or the desire to rest. Thus animals, according to this view, can be harmed by death to different degrees. For normal adult humans death is a great harm, in as far as they strongly wish to live on and have all kinds of plans for the nearer and further future. Great apes appear to have future-oriented projects as well and might desire their continued existence.² However, they seem to live more by the day, as do most other non-human animals. So, it seems that non-human animals are harmed significantly less by death than normal adult humans. Unborn and newborn babies and young children, as well as some mentally handicapped humans are comparable to non-human animals in this respect. They live more by the day and seem to have only a very limited grasp of their own existence over time. Like non-human animals, they are harmed significantly less by death, according to that view.

On this view, the harm of death consists in existing preferences being left unsatisfied. Unsatisfied preferences, according to that view, have a negative value for the individual (or they can even be considered to have negative value simply as such, as it were, ‘for the world’).

Even those who define welfare in terms of desire satisfaction acknowledge that getting what one desires is not always beneficial and not getting what one wants is not always harmful. If one wants to marry a person and does it, this can turn out not to contribute positively to one’s welfare after all. Likewise, if a suicidal teenager wants to die and is prevented from it, not getting what she wants at this point might actually be good for her, if the desire is based on mistaken beliefs. Therefore proponents of the desire account usually stipulate that only informed and rational desires count. Here a dilemma comes in. If the desire account defines welfare in terms of the satisfaction of actual desires, it might get the wrong results, as in the cases mentioned above. If, in contrast, it

² Singer (forthcoming 2011): 98-99.

counts only desires about what is really good for a being, the view will not focus on what is desired but on what is valuable for a being.³

So, the question arises how far one goes in focusing on what would be rational and informed desires rather than actual desires. Some want to go as far as arguing that a fetus would have the desire to stay alive if she were fully informed and rational.⁴ Thus, a fetus would be harmed by death, even though it may not possess any future-oriented desires. However, others have argued that this interpretation of the informed desire account stretches things too far:

Adjusting a person's actual desires for errors is one thing; attributing a wholly new desire to a being *that is not capable of having any desires at all*, or any desires of the relevant kind is something else altogether and something for which there is no obvious motivation. Preference utilitarianism should be formulated to cover only the former, as follows: We should satisfy, to the greatest extent possible, the preferences a being has, except that we should not satisfy a preference that results from errors of reasoning or errors about matters of fact.⁵

So, adjusting actual desires for errors is sanctioned by the informed desire account. Attributing desires to a being that is not capable of having any desires, however, goes too far. I agree that attributing desires to a being that is not capable of having any desires at all goes too far. After all, the informed desire account is an account of welfare. Those who do not have feelings and desires have no welfare at all, according to that account. However, attributing the desire to stay alive to a being that *is* capable of having desires does not go that far. It does not attribute desires to a being

³ Bradley (2009): 128-129.

⁴ Marquis (2009): 144.

⁵ Singer (2009): 156. Italics mine.

that has none. Even if an animal does not consciously or directly desire to stay alive because it has no concept of life and death, it might well desire other things. Staying alive might be instrumental for the fulfillment of those desires. Ascribing desires to go on living goes further than correcting existing desires for errors. It means taking into account a desire of which the being in question is not aware, provided that the fulfillment of that desire is instrumental for the fulfillment of the desires of which the being is aware. It seems to me that this is not obviously ruled-out by preference utilitarianism. After all, staying alive can be considered *a preference the being has*, even if it has this preference rather implicitly. As I said, staying alive is instrumental for the fulfillment of more explicit preferences.

A possible counter argument might be that some desires do not presuppose staying alive. Examples are the desire to stop hunger or the desire for the cessation of pain. If this is what an animal desires, death might serve this goal.⁶ It seems to me that by desiring food animals do not only desire the cessation of hunger. They do also desire the positive feeling that goes together with fulfilling their desire for food. Fulfilling a desire for food does not only cancel out a negative feeling, it can and usually does add something positive. This holds for other seemingly ‘negative’ experiential desires as well. Animals seem to positively enjoy what they necessarily have to do. This positive aspect of fulfillment is missed if the hunger problem is ‘solved’ by killing the animal. Furthermore, even if some desires do not presuppose staying alive, others might. At least for the fulfillment of those desires, staying alive would be instrumental.

Against the argument that even animals that have only short-term desires do (instrumentally) desire to stay alive, it can be said that when those animals are asleep, they might not have those short-term desires. Thus while asleep they do not retain desires. Unlike persons, they do not desire things for after they have woken up. So, killing them while asleep

⁶ DeGrazia (2003): 428.

would not frustrate any desires. Killing those animals painlessly while asleep would indeed not frustrate any desires if those animals do really not have any desires while asleep. Obviously, animals do not have conscious desires while they are sleeping (unless they have them in their dreams). However, desires that are not in the forefront of one's mind count as well. The animal might have such desires: it might desire to go on sleeping, to stay in the company of others (such as parents or offspring) to smell or feel certain things (such as the smell of their parents or offspring or the warmth of the sun). There always seem to be at least some dispositional desires. Staying alive is instrumental for the fulfillment of those desires.

So far, I have described a prominent account on the harm of death, based on the desire-fulfillment account of welfare. According to that account, animals are harmed less by death, because they lack the desire to go on living and have little or no other future-oriented desires. I have shown that it can be questioned whether the desire account must deny that animals desire to go on living. After all, the account does not focus on actual desires, but on 'ideal' desires. Thus, for instance, desires are corrected for errors of reasoning or errors about matters of fact. Furthermore, dispositional desires are taken into account along with actual desires. One might also take into account desires the fulfillment of which is instrumental for the fulfillment of other desires that the being has. I called those 'implicit desires'. Even animals that lack the conception of life and death but desire other things have an implicit desire to go on living that can be taken into account. As animals have at least some desires, including dispositional desires while being asleep, and – if one is willing to accept this – also an implicit desire to go on living, they are harmed by death. However, if the harm of death is defined in terms of the frustration of the desires that the being has when killed, it is true that those with less frustrated desires are harmed less by death. Therefore, according to that account, death is a lesser harm for animals (as well as for babies, infants, and some mentally ill people).

An alternative account of the harm of death does not focus on the frustration of what the being wants at the moment of its death, but on the loss of what would have been valuable for the being. That alternative account of the harm of death is compatible with the desire-satisfaction account of welfare, and also with every other account of welfare. Thus, unlike the account that has been presented above, the alternative account to be presented in the following section is not linked to a particular account of welfare.

3. Wants versus value

I will now present an alternative account of the harm of death, which claims that the harm of death is not frustration of wants but forbearance of value. This view is called the foreclosure or forbearance view, or the deprivation view concerning the harm of death.⁷ Instead of focusing on how much a being *wants* the future it would have had, this approach focuses on how much *value* that future would have had *for the being*. According to this view the harm of death consists in the foreclosure of the value that the subject's future life would have been for the subject. In other words, the harm of death for the subject can be determined by calculating the difference between the value of a being's actual life for her and the value of the being's counterfactual life for her, in case she had not died when she did.⁸ For instance, if a baby of three weeks old dies in an accident, and the baby would otherwise have had a happy childhood, studied philosophy and made a good career as a philosopher, had a happy life and died at the age of 80, then the baby is harmed a lot by death. Imagine that a philosophy student dies that has had an equally happy childhood as the baby would have had, and who would have had a good career as a philosopher and an equally happy life as the baby would have had. According to the foreclosure view, death is a lesser harm for

⁷ Kaldewaij (2008) defends this view on the harm of death.

⁸ Bradley (2009): 69-72. Bradley defends this view on the harm of death, which he calls the Difference Making Principle.

the student than for the baby because the student is deprived of less value.

The foreclosure view is compatible with different accounts of welfare. For instance, a hedonist account of welfare obviously does not put any weight on future-oriented desires or the desire to stay alive. On a hedonist account, death is harmful because it takes away all the enjoyment that the being otherwise would have experienced. The more pleasure is taken away, the greater the harm. Thus, all sentient beings that would otherwise have continued a pleasant life are harmed by death. It does not matter how future-directed those beings are. On a hedonist account of the harm of death, non-human animals are significantly harmed by death, insofar as death deprives them of the pleasurable experiences they would otherwise have enjoyed.

If one accepts the desire-satisfaction account of welfare, the forbearance view of the harm of death is a plausible option as well. According to the desire-satisfaction version of the forbearance view, death is harmful insofar as it precludes all future satisfaction of desire that a being would otherwise have had. This desire-satisfaction that is missed consists of both the satisfaction of the desires that the being already has, and the satisfaction of the desires that the being would come to have. In a way it is more straightforward to take into account all desire-satisfaction that is lost due to death, rather than counting only the lost satisfaction of desires that existed prior to death.

On other accounts of welfare it matters that the student has already invested more effort in his future. Furthermore, the student's life has already begun to take a shape or narrative, and death makes this a very bad narrative, while in the baby's life the narrative hasn't really begun. The foreclosure view can incorporate those concerns, "as long as it is intrinsically bad to have one's investments rendered futile or to have a life with a bad narrative structure. And if those things are not intrinsically bad,

why think they are relevant to the evil of death?”⁹ If those things were so bad that the student’s life would have been worse than the baby’s life, then, assuming that their total welfare when they had not died would be the same, the student loses more by dying than the baby. Hence, even if one accepts those views, this provides no reason to dismiss the foreclosure view.

In animal husbandry, animals are usually killed early in their lives. For instance, cows that are used for dairy production are usually killed after four years, while a cow could otherwise live for about twenty years. In egg production, male chickens are killed when they are a day old, because they cannot become laying hens and they are not selected for growing meat. Animals that are raised for meat are also killed in their early lives, for reasons of meat quality and efficiency. Those animals would have otherwise been able to continue the pleasant lives that they are assumed to have in animal-friendly animal husbandry. Killing those animals would be a significant welfare loss. It significantly harms those animals, according to the foreclosure view concerning the harm of death, because it precludes all future welfare they would otherwise have had.

According to the foreclosure view the harm of death is not determined by how much the being *wants* the future he would have had, but by the loss of the *value* that the future would have provided for the being. Hence, according to that view, animals can be significantly harmed by death even if they do not desire to go on living, and even if they have no or little future-oriented desires. In the following section, I will discuss a proposed modification of the foreclosure view that is relevant for the harm of death for animals.

⁹ Bradley (2009): 127.

4. Does psychological connectedness matter?

Although according to the foreclosure view the possession of future-oriented desires does not determine the harm of death, it has been suggested that death is a lesser harm for beings that have less psychological connectedness. If the harm of death for a being consists in the deprivation of that being's future welfare, it seems to matter in how far the being that is killed is connected to its future self. With other words, it seems to matter in how far what is taken away is really *her* future. This has to do with the concept of personal identity, in particular personal identity over time. This is relevant in order to determine the amount of personal harm. This is what we are after, when we determine the harm of death: it is the harm that death is for somebody. In principle different conceptions of personal identity over time are compatible with utilitarianism. With regard to the harm of death, it has been proposed that the underlying assumptions about personal identity over time matter:

In addition to asking how much good a person's future life would have contained, we must also ask [...] How close would the prudential unity relations have been between the individual as he was at the time of his death and himself as he would have been at those later times when the goods of his future life would have occurred? [...] The badness of the loss must be discounted for the absence of [psychological connections].¹⁰

So, according to this view, which McMahan calls the Time-Relative Interest Account, the loss of welfare is discounted in proportion to the lack of psychological connectedness between the being's present self and its future 'selves'. In order to determine the degree of psychological connectedness, it matters, for instance, how far a person's beliefs and

¹⁰ McMahan (2002): 183-4. Cited by Bradley (2009): 131.

desires remain the same, how rich the person's mental life is, how much the person remembers things from the past, and how much a person strives to satisfy desires that occurred in the past.¹¹

On this account of the harm of death, animals are harmed less by death than normal adult humans are. Although the degree of psychological connectedness varies per animal, animals generally live more in the moment, and have less far-reaching memories and desires. Even if this account of the harm of death was true and animals would indeed be harmed less by death than normal adult human beings, it is not clear that this would justify the routine killing of animals for minor purposes. After all, the harm of death for animals would not be zero, on this account. Generally, at least mammals and birds are known to have the capacity for rich mental lives, memories and desires for the (at least nearer) future. So, death might still be a significant harm for those beings.

This Time-Relative Interest Account, which is in fact a modification of the foreclosure view on the harm of death, has been criticized. In the remainder of this section I will present some controversial implications of the Time-Relative Interest Account.

One implication of the Time-Relative Interest Account is that killing a fetus does not harm it. Thus it is likely that this account can serve to justify abortion. According to the forbearance view, death is a significant harm for a sentient fetus, if it deprives the fetus of the pleasant experiences she would otherwise have had. As the time-relative interest account discounts the harm of death for lack of psychological connectedness, it implies that death is not very harmful for a fetus. After all, a fetus has no desires for the (distant) future and her later self will have no memories of her life as a fetus. Even though the foreclosure view considers death harmful for the fetus it needs not preclude abortion. Even if a utilitarian takes the significant harm to the fetus into account,

¹¹ McMahan (2002): 80. See also Bradley (2009): 130.

killing a fetus will be justified if there are countervailing benefits. A utilitarian will include the welfare of the mother and that of others that might once be in her situation in the calculation. The Time-Relative Interest Account does also imply that late embryos, demented and brain-injured humans are harmed by death only very little. Those implications have been brought forward to question the plausibility of the Time-Relative Interest Account on the harm of death.¹²

One influential argument in favor of the Time-Relative Interest Account is based on an imaginary case. A person has a disease that would soon cause her death. The person could be cured. The cure would cause a profound change in the person's values and desires. After the cure the person would live on for many happy years, but there would be very little psychological connectedness between the person before the cure and after the cure. The Time-Relative Interest Account can explain that it would be reasonable for the patient to reject the cure.¹³ This capacity to explain why it would be reasonable to reject the cure has been brought forward as an advantage of the Time-Relative Interest Account, as compared to the foreclosure view without that modification. Proponents of the foreclosure view without modification have considered rejecting the cure shortsighted and irrational, "like the decision of a child to ignore the consequences of his behavior on his adult self, since he does not currently care about the things his adult self will care about."¹⁴

If one explores why it might be reasonable to reject the cure, it becomes clear that those considerations are compatible with the unmodified foreclosure view. First, if the cure causes a change of identity that implies that the person has no future, even if she accepts the cure. Such a view on personal identity is compatible with the foreclosure view.

¹² Liao (2007). Feder Kittay (2005).

¹³ McMahan (2002): 77.

¹⁴ Bradley (2009): 117.

In that case the unmodified foreclosure view would not recommend taking the cure. Second, another reason for not taking the cure can be the fear that the cure causes a change of the person's moral values. For instance, the person might not be a good mother or vegetarian anymore, after having taken the cure. That reason for rejecting the cure is not based on prudential value. Any view about the harm of death does not rule out that a person cares for other things, except her own welfare. Third, another reason to reject the cure can be that the person's life with the cure might have less narrative unity. Attributing prudential value on the basis of narrative unity is compatible with the foreclosure view. I have already explained that the foreclosure view is compatible with any account of welfare. Fourth and finally, another reason for rejecting the cure can be that the person considers it unattractive that her values are changed in a kind of 'unnatural way', by a medical treatment, rather than, say, by education or maturing. Again, attributing disvalue in case of medical interventions in one's life is compatible with the foreclosure view, because the foreclosure view is compatible with any account of welfare. The point is that the reasons against taking the cure are compatible with the unmodified foreclosure view. They do not provide any arguments in favor of the modification as proposed by the Time-Relative Interest Account.¹⁵

What does the Time-Relative Interest Account imply when a person is harmed without knowing it by being deprived of a benefit? Here is an example:

Suppose Claire has an infant son, Charlie, who has a trust fund that he may use when he turns 25. For simplicity's sake suppose, sadly, that the trust fund is the only potential source of happiness in Charlie's life.

¹⁵ Bradley (2009): 120-121 has brought this argument forward.

Claire intends to drain the fund secretly, and prevent Charlie from ever finding out about its existence.¹⁶

The Time-Relative Interest Account implies that draining the fund as early as possible is least harmful for Charlie. As the psychological connectedness between the three-week-old Charlie and the 25-year-old Charlie are much weaker than those between Charlie at 23 and 25, the harm that is caused by draining the fund must be discounted. Another implication of the Time-Relative Interest Account seems to be that it cannot account for harms that are caused before a person exists. Suppose, for example, that the trust fund for Charlie existed and was drained by Claire before Charlie was born (or before he existed as a sentient being). In that case, the Time-Relative Interest Account would imply that Claire's draining the fund was not harmful for Charlie at all. This seems odd, because in all three cases Charlie would be equally deprived, unbeknownst to him, of all the happiness in his life.¹⁷

McMahan is aware of those problems and acknowledges that future interests may be relevant to the moral assessment of present acts:

If our concern is with individuals' time-relative interests, we must take account of all the time-relative interests affected by our action. The important consideration is whether one's action frustrates a time-relative interest; it does not matter whether the act is done before the time-relative interest exists.¹⁸

This would mean that even though Charlie was very young or even non-existent when Claire drained the trust fund, Charlie's later time-relative

¹⁶ Bradley (2009): 136.

¹⁷ Bradley (2009): 136-137.

¹⁸ McMahan (2002): 283.

interests count as well. If this is the case, the Time-Relative Interest Account does not have the above-mentioned implications in the case of draining the trust fund. It will not do to say that draining the fund when Charlie is very young matters less, because of his weak psychological connectedness with his future self at the time he would have received the fund. All time-relative interests that Charlie will have in the fund have to be taken into account, including those he has later, after the fund has been drained. Just as with the unmodified foreclosure view the timing of draining the trust fund makes no difference, if one understands the Time-Relative Interest Account in that way.

If the Time-Relative Interest Account is supposed to yield the same result as the unmodified foreclosure view in case of draining the trust fund, why does it yield different results in cases concerning the harm of death? The explanation is as follows: The difference is that if a being is killed, he or she will never again have any time-relative interests. In contrast, Charlie lives on. He will continue to have time-relative interests in receiving the trust fund, even after the fund has been drained.

Now, imagine that either an infant or a 20-year-old person had to die. If the 20-year-old dies, the infant lives on and will continue to have time-relative interests in continued life. As the 20-year-old dies, he will have no more time-relative interests in continued life. Thus, the time-relative interests in continued life of the infant will outweigh those of the 20-year-old, and we can afterwards conclude that it was better that the infant lived. If however, the infant dies and the 20-year-old lives on, we must conclude that this was the better outcome. Hence, on this definition of the Time-Relative Interest Account, what is the better outcome depends on what actually happens. Furthermore, contrary to what McMahan wishes to claim, this account does not imply that the death of a 29-year-old is worse than the death of an infant.¹⁹

¹⁹ Holtug (2010).

To sum up, the Time-Relative Interest Account proposes a modification of the foreclosure view. It discounts the harm of death (defined as the loss of what would have been valuable for the being) for lack of psychological connectedness between the being at the time of death and the being's future self at the time the value would have occurred to it. This modification of the foreclosure view is controversial. Even if one accepts the Time-Relative Interest Account, this does still imply that at least mammals and birds and at least some species of fish are harmed to a certain extent by death. Thus, it is unlikely that the routine killing of mammals and birds for minor purposes could be justified on that basis.

5. Conclusion

Competing views on the harm of death that are compatible with utilitarianism have different implications about how much animals are harmed by death. However, all views imply that pigs, cattle, chicken and some sorts of fish are harmed by death at least to some extent. This is because those animals can not only be deprived of welfare, they also have been shown to have memories and intentions and at least some future-directed desires.

Furthermore, let me note again here that quite apart from what the harm of death is for the animals, the deaths of animals that could otherwise have continued pleasant lives present a welfare loss. All else being equal, killing an animal that could have continued a pleasant life is a loss of welfare. For proponents of the Impersonal View on the evaluation of outcomes this explains why killing animals that would have had a pleasant future is problematic. For proponents of the Person-Affecting View on the evaluation of outcomes, killing an animal is morally

problematic in so far as this causes harm to sentient beings, whoever they are.

In the following chapter, we will therefore explore the Replaceability Argument that has been brought forward within utilitarianism. This argument is supposed to show a way of compensating for the welfare loss that is caused by killing an animal that would otherwise have had a pleasant future.

4. THE REPLACEABILITY ARGUMENT

1. Introduction

The Replaceability Argument proposes a way of compensating the welfare loss that is implied by the killing of animals that would otherwise have had a happy future. In this chapter, I will explore this argument. Peter Singer has explicitly discussed the Replaceability Argument in animal ethics. Apart from Singer's work, the argument is rarely discussed. If other authors discuss the Replaceability Argument it is always a reaction to Singer. Therefore, my exploration of the Replaceability Argument will be presented in dialogue with Singer's take on the issue.

The underlying idea of a compensation for lost welfare can best be grasped by looking at the numbers. Therefore, I will introduce the logic of the argument by algebraically depicting outcomes of alternative actions and compare those distributions of welfare (section 2). Once the underlying idea is clear, a definition of the Replaceability Argument can be given. The conditions will be specified under which the Replaceability Argument sanctions the killing of animals (section 3). The question arises whether those conditions can be met in the practice of animal husbandry. As we will see, certain forms of animal-friendly animal husbandry might

conform to those conditions (section 4). That far, it seems that utilitarianism needs the Replaceability Argument in order to sanction the routine killing of animals in the practice of animal-friendly animal husbandry.

Therefore, I will explore what it takes to embrace the Replaceability Argument. If one wants to claim that animals are replaceable as proposed by the Replaceability Argument, the question comes up whether this implies that humans are replaceable as well. I will critically discuss the only proposal so far (by Singer) for limiting the scope of the Replaceability Argument. I will explore what it takes to limit the scope of the Replaceability Argument as suggested by Singer, and whether this is possible at all (section 5). Furthermore, and more fundamentally, the Replaceability Argument is based on the Total View, which takes the welfare of all existing and all possible beings into account in the evaluation of outcomes. This view is controversial within utilitarianism (section 6).

2. The underlying idea of the Replaceability Argument

Within utilitarianism, as we have seen, actions are not categorically permitted or forbidden. Whether an action is right or wrong depends solely on its consequences in terms of welfare. Only the action that maximizes aggregative welfare is morally right. If animal-friendly animal husbandry is to be evaluated from this perspective, it must be admitted that the killing of animals that would otherwise have had a happy future seems, at first sight, to be impermissible, given that we have dismissed the average view as a method of aggregation. After all, it implies a loss of the welfare that the animals otherwise would have had. Let me illustrate the choice between killing and not killing the animals by depicting the numbers. I will show the distributions of welfare in both outcomes. Each distribution shows the condition of the animal at different times, throughout its life. When the animal exists, its condition is its welfare, which can be positive, negative or neutral. If the animal does not exist at a

time I put an Ω to distinguish nonexistence from neutral welfare.¹ One option is letting the animal live; the other option is killing it. For simplicity's sake, I assume a constant welfare level of 4 throughout the animal's life.

The animal is not killed at t_5 : $(\Omega_1, 4_2, 4_3, 4_4, 4_5, 4_6, 4_7, 4_8, 4_9, \Omega_{10}, \Omega_{11})$

The animal is killed at time t_5 : $(\Omega_1, 4_2, 4_3, 4_4, 4_5, \Omega_6, \Omega_7, \Omega_8, \Omega_9, \Omega_{10}, \Omega_{11})$

So, this shows the amount of welfare that is implied in each of two outcomes: when the animal is killed at t_5 as compared to when it is not killed at t_5 . From a utilitarian perspective, the consequences of two possible actions - namely killing or not killing the animal at t_5 - matter. Not killing the animal, *ceteris paribus*, results in more overall welfare and is therefore required.

Now, the idea of the Replaceability Argument is that the welfare loss due to the killing can be compensated by the bringing into existence of a new animal, which would not otherwise exist and which will be at least as happy as the killed one would have been. The outcome contains the welfare of both the killed animal until its death and the newly created animal after that time. The distribution for each animal is put in a separate row.

The animal is killed at and replaced by another animal at t_5 :

$(\Omega_1, 4_2, 4_3, 4_4, 4_5, \Omega_6, \Omega_7, \Omega_8, \Omega_9, \Omega_{10}, \Omega_{11})$

$\Omega_1, \Omega_2, \Omega_3, \Omega_4, \Omega_5, 4_6, 4_7, 4_8, 4_9, \Omega_{10}, \Omega_{11})$

¹ This method of notation is based on Broome (2004).

This sequence of killing and replacing animals could go on, as the new animal can be killed and replaced as well, and so on. The crux is that the total amount of welfare remains unaffected by the killing, provided that the killed animal is replaced, as illustrated above.

Utilitarianism in general holds that it is permissible to kill an animal provided that this action maximizes welfare. The Replaceability Argument takes into account the welfare that could be gained by adding a new animal that would not otherwise exist. With other words: Killing an animal with a pleasant life is permissible if and only if the lost welfare is compensated by bringing into existence another animal, which is at least as happy and would not otherwise have existed.

3. Defining the Replaceability Argument: specifying the relevant conditions

After having offered a first idea of what the Replaceability Argument is about, I will now strive for a more formal definition of the argument. I will build on definitions that have been proposed by Sapontzis and Pluhar, as those, to my knowledge, are the only formal definitions that have been brought forward in the literature. Sapontzis defines the Replaceability Argument as follows:

It is permissible, *ceteris paribus*, to use an animal and to kill it (for food or research or anything else), provided that the following conditions are met:

- (a) the life of the animal is on balance a life worth living,
- (b) the animal otherwise would have no life at all (would not exist),

and

(c) the animal will be replaced, at or after death, by another animal in the case of which conditions (a) and (b) hold.²

I will explain why I will deviate from that definition. First, Sapontzis suggests that the Replaceability Argument is about the permissibility of ‘using and killing’ an animal. Yet, I want to focus on the Replaceability Argument as an argument for justifying the killing of animals. By talking about using and killing, Sapontzis lumps together two arguments, which I prefer to distinguish. I would like to distinguish the Replaceability Argument, as an argument that specifies conditions that can compensate for the welfare loss caused by the killing of an animal from another argument, called the Logic of the Larder.

Very briefly, the Logic of the Larder claims that keeping and killing animals for consumption actually benefits the animals in question, because they would not exist at all if it were not for being consumed and, according to this argument, a short pleasant life is better for an animal than no life at all. Clearly, the Logic of the Larder is not merely about compensating the welfare loss that is caused by killing an animal that could otherwise have had a pleasant future. The Logic of the Larder argument does positively promote the keeping and killing of animals under certain conditions. This goes further and brings in an extra dimension. Actually, Singer seems to lump together what I take to be two different arguments as well. Yet, I think things get clearer if we separate them. I propose to concentrate on the replacement aspect of the Replaceability Argument and discuss the additional claims of the Logic of the Larder argument later on, in chapter 9. In order to justify the killing of an animal that could otherwise have continued a pleasant life, the replacement of the lost welfare would be sufficient. The aspect of replacement is what I wish to capture in my

² Sapontzis (1987): 177.

definition of the Replaceability Argument, without mingling it with any other and further claims.

Secondly, Sapontzis says that the life of the animal must be, on balance, 'worth living'. I think that taking a position on the controversial issue of when a life is worth living is not necessary here. This can be avoided by saying instead that the animal should have a positive score on lifetime welfare. More precisely, we are talking about animals that would otherwise have had a future life with positive welfare. Assuming a mental state account of welfare, this means that the animal's future life would have been, on balance, a pleasant one. Thus, the positive welfare that the animal would have had, if it were not killed is what has to be replaced, in order to compensate for the welfare loss due to the killing. The past welfare is irrelevant and therefore, strictly speaking, the lifetime welfare is not what matters.

A third aspect that I find problematic in Sapontzis' definition is condition (b), which says that the animal must otherwise not have existed. As I propose to understand the Replaceability Argument, this condition is only necessary for the animal that is meant to replace the previous one. It is not necessary for the animal that is to be replaced. Again, the reason that Sapontzis mentions this condition also for the animal that one is considering to kill shows that he lumps together what I take to be two different arguments. For the other argument, the Logic of the Larder, it is indeed relevant that the animal that one considers to kill would not have existed if it were not for the purpose of ultimately being killed.

Pluhar and Sapontzis take the *ceteris paribus* clause to imply that the killing causes no pain or fear for the animal nor any suffering for others, such as mothers or mates.³ Yet if the killing would have such negative consequences on welfare, this would not necessarily make the killing and replacement impermissible. The negative welfare consequences for the

³ Pluhar (1995): 185. Sapontzis (1987): 177.

animal or its relatives that would be caused by the killing could be outbalanced by the superior happiness of the animal that takes the place of the previous one. Therefore, I will rather claim that the *ceteris paribus* clause requires that the killing does not have any *unbalanced* negative (side-) effects. Instead of taking this into a *ceteris paribus* clause, I will add it as an extra condition.

These considerations lead me to the following definition of the Replaceability Argument:

It is permissible to kill an animal, provided that the following conditions are met:

- (a) The future welfare of the animal would have been positive,
- (b) The animal will be replaced, at or after death, by another animal, whose lifetime welfare is at least as positive as the future welfare of the killed animal would have been, and which would not otherwise exist.
- (c) The killing does not have any unbalanced negative side effects (such as fear or suffering for the animal or others).

This can be further specified according to any particular account of welfare. For instance, on a hedonist account, one would say that the future life of the animal should be pleasant overall. On a desire-satisfaction account one would say that the balance of the amount of preference-satisfaction minus the amount of preference frustration should be positive.

Now we know the conditions under which the Replaceability Argument allows killing, it is interesting to find out whether those conditions can be met in practice. I will turn to this question in the next section.

4. Practical implications for animal-friendly animal husbandry: can the conditions be met?

I will now discuss whether and in how far the Replaceability Argument would justify the practice of animal-friendly animal husbandry.

At first glance, the Replaceability Argument clearly supports animal-friendly animal husbandry. First of all, the animals in animal-friendly animal husbandry are indeed assumed to have *positive lifetime welfare*. Furthermore, animal husbandry implies that killed animals are ‘replaced’ by new ones. Those new animals would indeed not otherwise exist. (After all, it is not the purpose of animal husbandry to produce animals as such, but to produce animal products, i.e. to turn animals into products for human consumption.) Furthermore, it is assumed that the animals can be killed painlessly and more generally that the killing *does not produce any unbalanced negative (side-) effects*.

Does actual practice live up to the conditions of the Replaceability Argument? In order to answer this question, we must know whether ‘animal-friendly animal husbandry’ can fulfill the criteria that are specified in the Replaceability Argument. I have already mentioned that the purpose of animal husbandry is to turn animals into products. At a certain stage, animals leave the farm in order to be killed (or are killed on the farm), and new animals take their places. Either the animals are killed all at once and replaced by new ones (as I remember with geese on a German farm, which all suddenly disappeared before Christmas) or the replacement happens gradually, as when once in a while an older animal goes out and a younger animal comes in. No matter how it happens in practice, the principle is that animals are replaced and that the existence of the new ones might somehow compensate for the death of the killed ones. Above, I have expressed this idea more schematically, in a model. So, the second condition of the Replaceability Argument is fulfilled in practice. Therefore,

there are two conditions left, which deserve closer examination. First, would the animals in ‘animal-friendly animal husbandry’ indeed have pleasant lives? This is the first condition of the Replaceability Argument. Secondly, what about unbalanced negative (side-) effects of the killing? This concerns the third condition of the Replaceability Argument. Let me explore those questions in turn.

4.1 Do the animals have pleasant lives?

Do the animals that are kept in so-called animal-friendly husbandry systems actually have pleasant lives? In order to answer this question, we first need to know what ‘animal-friendly animal husbandry’ would mean in practice. Proponents of ‘animal-friendly animal husbandry’ often refer to *organic animal husbandry* when it comes to putting the idea into practice. At least, there seems to be a broadly shared image, according to which organic animal husbandry is animal-friendly.⁴ The organic sector re-enforces this image by depicting happy and healthy animals in a natural surrounding. The organic sector tells stories about farmers who had a bad conscience about their treatment of the animals in regular husbandry systems and who, out of love for the animals, turned to organic husbandry.⁵ The emotional and intuitive idea of many people is that intensive animal husbandry is not in line with how we morally ought to treat animals. Organic animal husbandry is viewed as a real alternative. It seems to be much better than intensive animal husbandry. But is it good enough, with respect to the ideal of animal-friendly animal husbandry?

Do the animals in organic animal husbandry have happy lives? Researchers have supported the idea that the lives of organic animals are happier in many respects. They have demonstrated many advantages of organic animal husbandry as compared to regular animal husbandry

⁴ Stassen (2006). Meijs (2007).

⁵ Hentzepeter (2008): 23.

systems, in terms of animal welfare.⁶ For instance, dairy cows feed themselves more naturally, because they spend at least some time outside on the meadow and they receive as much raw food as they want. There are fewer occurrences of metabolic diseases. Furthermore, the cowsheds are bigger. Therefore, there is less competition about feed and space, which in turn results in lower-ranked animals being less frightened. Organic pigs have the possibility to root and explore and they are more active, due to exposure to daylight. Sows can exhibit nest-building behavior. Some organic sows have the opportunity to bath in the mud. They receive raw feed and straw, which reduces stereotyped behavior and enhances peace in the group. There is less tail biting, less aggression and less fear among organic sows. Organic pigs have healthier paws and a more intact skin and a better metabolism.⁷

In some respects, however, it has also been suggested that organic animal husbandry scores worse in terms of animal welfare. For instance, in organic husbandry there are less possibilities of treating sick animals, due to stricter rules about medication. More animals ‘fall out’, i.e. die. Claw problems are the major welfare problem for cows. Cows on organic farms are not better off in this respect, because also in organic farms most cowsheds have concrete floors, which causes claw problems. Organic cows, as compared to regular cows, are more frequently outside. However, most of the time, notably in the winter, they are inside as well. Organic cows do also suffer from mastitis and other infections. As noted before, due to restrictions on medication, they cannot always be effectively treated when they are sick. Elsbeth Stassen, professor in Animals and Society at Wageningen University goes as far as stating that she would rather be a cow in intensive husbandry systems than a cow in organic farming.⁸

⁶ Ruis & Pinxterhuis (2007).

⁷ Ruis & Pinxterhuis (2007).

⁸ Stassen (2006).

It is true that farmed animals suffer from various causes, *also in organic animal husbandry*. For instance, piglets are usually castrated. Even if this is done with anesthetics, it is still a painful and stressful treatment. (Even the anesthesia causes considerable pain.) Organic cows (not Demeter cows, though) are usually de-horned.⁹ This is done with anesthetics. Still, the treatment is at least stressful. Furthermore, the fact that dehorning is necessary in order to avoid injuries among the cows (and save space) might indicate problems with their housing or treatment more generally. Dairy cows are made pregnant every year in order to be milked. The calves are taken away from the cows after a short period of time, while calves would naturally suckle for 2-6 months. The separation of cow and calf is highly distressing, due to the strong maternal bond and the strong nursing need of the calf.¹⁰ In general, even organic animals are selectively bred and live in more or less unnatural conditions. Expressions of species-specific behavior are restricted, and all this is likely to influence their health and welfare in a negative way. Farmed animals are often sick, as the need for antibiotics and other medications shows, also in organic farming.¹¹ So, it is open for discussion, whether any farmed animal, even in organic animal husbandry, has an overall pleasant life, after all.

However this might be, the current systems can be ameliorated with respect to animal welfare, for instance by stopping the castration of pigs and the de-horning of cows and by keeping the calves together with the cows. Some infringements on animal welfare seem to be linked to animal production as such. However, there might be systems of animal production that are compatible with each animal's lifetime welfare being positive. After all, the animals do not need to be extremely well off in order for the Replaceability Argument to be applicable. Their welfare just

⁹ Demeter is a label for organic products, with additional requirements based on a certain worldview.

¹⁰ Ruis & Pinxterhuis (2007).

¹¹ Ruis & Pinxterhuis (2007). Hentzepeter (2008): 24.

needs to be better than neutral. Therefore, let us assume for the remainder of this thesis that forms of animal-friendly animal husbandry can meet the first condition of the Replaceability Argument.

4.2 What about unbalanced negative (side-) effects?

What about the third condition, which requires that the killing does not have any *unbalanced* negative (side-) effects? It can be argued that the killing has negative side effects. For instance, the killing might cause fear, discomfort and pain for the animal that is killed due to the situation in the slaughterhouse and the foregoing transport. The killing might cause suffering for others e.g. the animal's mother as in the case of the male calf that is taken away and killed. It seems to be unlikely that by other husbandry systems, negative (side-) effects can completely be avoided. Negative effects can in principle be compensated by the *greater* happiness of the animal that replaces the first one. Yet, in the case of animal husbandry, this is very unlikely to happen. It is unlikely that the animals that replace the killed ones will have increasingly positive lifetime welfare.

Even if an animal would not be killed, it is likely that the animal would experience some discomfort in old age before dying naturally. Thus, in any real life examples, unlike my example in section 2, the welfare level of the animal would possibly be less at the end of its life. Even if this is the case, the assumption is that the animal would have had an overall pleasant future, even taken this likely decrease in welfare into account. Therefore, killing the animal still reduces the overall quantity of welfare, even if it has the result that the welfare decrease due to old age will be avoided. If the killing causes additional welfare loss, because it causes suffering or desire-frustration, this must be taken into account over and above the welfare loss that is caused by the loss of the welfare that the animal would have experienced in the future. So, in order for the practice of animal-friendly animal husbandry to meet the third condition of the Replaceability

Argument, the killing should take place without any unbalanced negative side effects, such as pain or fear for the animal or others. This is far from being realized with most current methods of killing animals.¹² However, more welfare-friendly methods of killing might become available. Hence, let us assume for the remainder of this thesis, that the practice of animal-friendly animal husbandry in principal is able to meet also the third condition of the Replaceability Argument.

To sum up, it is questionable whether current practices of animal husbandry could satisfy the conditions of the Replaceability Argument. As far as we can tell, there *might be* practices of animal husbandry that do grant animals a sufficiently pleasant life and that avoid any unbalanced negative side effects of the killing. Probably, those systems would need to score better in terms of animal welfare than current systems of organic agriculture. With the aid of the Replaceability Argument, utilitarianism would admit those practices as morally permissible. In that sense, utilitarianism would support animal-friendly animal husbandry.

In the remainder of this thesis, I will explore the assumptions and implications of the view that animals are replaceable. I will also explore the implications and assumptions of an alternative utilitarian view, according to which animals are not replaceable. In the following section, I start by exploring the implication, that – once the Replaceability Argument is accepted – not only nonhuman animals would be replaceable, but also human beings would be.

5. The scope of the Replaceability Argument

The idea that animals are replaceable as proposed by the Replaceability Argument might suit well those who wish to support animal-friendly

¹² Foer (2009): 228-233, 255-256.

animal husbandry. In contrast, the idea that animals are replaceable has also been criticized for being incompatible with taking animals seriously.¹³ Even Singer, who accepts the Replaceability Argument admits that it “feels odd” to compensate the killing of one being with the creation of another being. Singer admits: “The proposition that the creation of one being should somehow compensate for the death of another does have an air of peculiarity.”¹⁴

If beings are replaceable, this holds not only for non-human animals, but also for human beings, unless a morally relevant criterion can be found for making distinctions. The distinction between who is replaceable and who is not should not, according to utilitarianism, be made in an arbitrary way. Hence, the question comes up whether all sentient beings are replaceable. What would proponents of the Replaceability Argument think of the idea that not only non-human animals but also humans are replaceable? The idea that the Replaceability Argument would also hold for human beings has led to consternation.¹⁵ It has been brought forward as a reason for rejecting utilitarianism altogether.¹⁶ Even those who are not troubled by replaceability of non-human animals tend to think differently if humans are concerned. So, a major challenge for those who want to adhere to the Replaceability Argument and want to avoid its application to humans is to exclude at least some animals from its scope.

Peter Singer is to my knowledge the only one who has brought forward an argument for restricting the scope of the Replaceability Argument. In this section, I will explore what it takes to reduce the scope of the Replaceability Argument, and whether this is possible at all. Singer’s argument for restricting the scope of the Replaceability Argument is based on a particular account of welfare, which is the desire-satisfaction account,

¹³ Regan (2004): 208, Pluhar (1995): 190.

¹⁴ Singer (1995): 229.

¹⁵ Singer (1993): 122.

¹⁶ Pluhar (1995): 190.

in combination with a ‘moral ledger model’ (section 5.1). I will explore several implications of Singer’s account. First, Singer’s account of welfare is pessimistic (section 5.2). Secondly, Singer’s account of welfare is linked to a particular conception of the harm of death (section 5.3). Third, Singer’s account of welfare is ad-hoc in several respects (section 5.4). Fourth, Singer’s account of welfare is still rather inclusive (section 5.5). Very recently, Singer has defended a somewhat revised position, which will also be explored (section 5.6).

5.1 Singer’s account of welfare: desire-satisfaction account plus moral ledger model

As I have explained in chapter 2, utilitarians have to choose an account of welfare. Singer accepts a desire-satisfaction account of welfare. This means that Singer defines welfare in terms of desire-satisfaction. At first glance, the Replaceability Argument works the same for a desire-satisfaction account of welfare as for a hedonist or objective list account. Thus, how can Singer restrict the scope of the Replaceability Argument?

The desire-satisfaction account as such is not sufficient for restricting the scope of the Replaceability Argument. Suppose, for instance, that a human person was killed and replaced by another human person. One could say that the killed person, had he not been killed, would have had a certain amount of welfare, in terms of desire-satisfaction. The being would have had many more desires, most of which would probably have been fulfilled, resulting in positive welfare. Provided that the new person has at least as much welfare as the killed one would have had, the replacement seems to compensate for the welfare loss and hence make the killing morally acceptable.

In order to preclude this possibility, Singer makes the following crucial move: He declares that every unsatisfied desire counts negative on the welfare scale. With other words, every unsatisfied preference creates

negative welfare. As soon as the desire is satisfied, the negative effect on welfare is cancelled out and becomes neutral. Singer suggests that we should “think of the creation of an unsatisfied preference as putting a debit in a kind of moral ledger that is merely cancelled out by the satisfaction of the preference”.¹⁷ A ‘ledger’ is a book (or what nowadays might be a computer program) in which for instance a shopkeeper memorizes debits until they are paid back. The shopkeeper writes down all the debits, and once a debit is paid back, the negative entry in the book is cancelled out. With this ‘moral ledger model’ in place, replaceability seems to be ruled-out.

How can the moral ledger model rule out replaceability? A simple calculation can show how this works: Killing an adult human being normally means that some of her preferences - and usually very strong preferences, such as wanting to go on living or finishing some projects - are not fulfilled. Thus killing a human adult being has the consequence that certain important preferences cannot be satisfied. Those negative entries cannot be canceled out in this being’s life any more. The welfare score of the being remains negative. It would be of no use to replace this killed being by a new one. After all, the best that the new being can achieve is to satisfy all his or her preferences and end up with a neutral welfare level. The new being’s welfare level can never be positive. This is because of the moral ledger model, according to which the fulfillment of desires does only cancel out negative welfare and does never produce positive welfare. Therefore, a compensation of the negative effect on welfare caused by the killing of the human adult is not possible. There can be no positive welfare with which the negative welfare due to the killing can be compensated. So, in order to limit the scope of the Replaceability Argument, Singer recurs to a desire-satisfaction account of welfare, in combination with the moral ledger model. In the remainder of this discussion, I will refer to this combination as “Singer’s account of welfare”.

¹⁷ Singer (1993): 128. See also Singer (forthcoming 2011): 114.

Who can be excluded from replaceability in that way? For Singer's account of welfare it is crucial that unsatisfied preferences count as negative welfare. The more unsatisfied preferences a being has at the moment of its death, the worse it is in terms of effects on welfare that this being dies. Now, which beings are likely to leave the most unsatisfied preferences when they suddenly die? It seems obvious: the more future-oriented desires a being has, the more likely it is that death frustrates the fulfillment of those desires. If, in contrast, a being does not have any desires for the future, death will not frustrate those desires. Killing such a being would not have any negative effect on welfare. Thus, generally, the more future-oriented a being is, the less it is replaceable. Those who have more and stronger future-oriented desires are likely to leave greater 'minuses' when killed. Singer brings forward the desire to go on living as an important criterion. According to Singer, those who are future-oriented in their desires and desire to go on living are not replaceable. As this characteristic varies in degrees, so varies the degree to which an individual is replaceable. For instance, a being that though conscious has no desires for the future is completely replaceable. That means it is replaceable without leaving a 'minus'. In the second edition of *Practical Ethics*, Singer mentioned fish in that regard, but after new insights into the capacities of fish, Singer does not consider fish completely replaceable any more. Now, Singer simply states that "if such animals exist" they are fully replaceable. Indeed, it is unclear whether sentient beings without any desires exist at all. Beings that have desires, but whose desires concern only the near future, are likely to leave only small minuses, and are thus more replaceable, according to Singer.

Those familiar with the second edition of *Practical Ethics* might note that this gradual view of the possession of the relevant characteristics and thus of replaceability is not made explicit there. To the contrary, in his second edition of *Practical Ethics*, Singer frames the distinction as one between persons and non-persons. Persons are those beings, human or non-human, who possess certain characteristics, such as self-awareness and

rationality. Persons possess a sense of their own existence over time and future-directed desires.¹⁸ Those characteristics are deemed relevant with regard to replaceability. Singer's framing of the relevant distinction as one between 'persons' - and 'non-persons' and Singer's talk (in the introduction to his second edition of *Practical Ethics*) about 'drawing clear lines' between who is replaceable and who is not - suggest a dichotomy between those who are replaceable (non-persons) and those who are not (persons). In any case, Singer currently explicitly assumes a gradual account of the possession of those characteristics, and thus also a gradual account of personhood. As personhood, as Singer now makes clear, comes in degrees, so does replaceability.¹⁹

So, now we know how Singer proposes to limit the scope of the Replaceability Argument. He does this by recurring to the desire-satisfaction account of welfare, in combination with the moral ledger model. We also have at least a rough idea of who would be replaceable and who would not. The division is not drawn between humans and non-humans, as this would be likely to be unjustified discrimination, and thus what Singer calls 'speciesist'. Instead, criteria are used that follow directly from Singer's account of welfare: what matters is how future-oriented a being is.

Before I will critically discuss Singer's account, I need to mention one important revision that Singer has made to the moral ledger model. As people have usually at least some unsatisfied preferences when they die, the moral ledger model implies that all of us score negative on lifetime welfare and thus that all of us had better not lived. In order to remedy for this unwelcome implication, Singer has proposed the following modification of the moral ledger model:

¹⁸ Concerning the use of the term 'person' in philosophy and ethics, see Sturma (2006).

¹⁹ Singer (forthcoming 2011): 119-122.

Thus the moral ledger model of creating and satisfying a preference will not do. It might be saved by attaching to it a stipulation that sets a given level of preference satisfaction, below complete satisfaction, as a minimum for overcoming the negative entry opened by the creation of a being with unsatisfied preferences. This might be the level at which we consider a life ceases to be worth living, from the perspective of the person leading that life. Such a solution seems a little ad hoc, but it may be possible to incorporate it into a plausible version of preference utilitarianism.²⁰

The idea is that a being's welfare level is negative as soon as a being has unsatisfied preferences. It is immediately restored to a zero level if the preference is satisfied. The welfare level remains negative, as long as one or more preferences remain unsatisfied. In order to avoid the implication that most of us had better not lived, Singer claims at this point that not all preferences need to be fulfilled in order to consider a life worthwhile. So, a life in which some preferences are not fulfilled can still come out of the negative score. The 'zero' is moved down the scale, in order to transform minuses into zero. It seems that Singer wants to say that the 'minus' will be erased and one will score 'zero', even if not all preferences are fulfilled. Thus, the neutral situation does not require the satisfaction of all preferences. This is probably in line with people's experience of their welfare. Furthermore, it avoids the implication that all of us had better not lived. Most people will achieve the zero-level of welfare, if one needs not to have all preferences fulfilled in order to come out of the negatives.

While it is obvious that according to the moral ledger model, welfare could only be negative, or at best neutral, it is less clear how this would be according to the modified moral ledger model. Would there be positive welfare, or would the neutral zone be wider? Unfortunately, Singer's description of the modification is restricted to the above-

²⁰ Singer (2005): 129.

mentioned citation. It is interesting, though, that Singer takes counter-intuitive implications such as the implication that all of us had better not lived, or that all of us are replaceable, as a reason for modifying his moral theory. Since Singer has introduced the Replaceability Argument, he has revised his position on this issue several times, and he always admitted to remain unsatisfied with his proposal. This is mainly due to counter-intuitive implications of each proposal. Here, I will focus on Singer's most recent strategies.²¹ Appeals to counter-intuitive implications in general are often made in the evaluation of moral theories. In chapter 8, I will discuss this issue, and in particular utilitarianism's stance on it.

Now Singer's account of welfare is more or less in place, I will turn to explore its implications.

5.2 Singer's account of welfare is pessimistic

Singer's moral ledger model is not in line with how preference-satisfaction is usually experienced. In some cases of preference satisfaction, the moral ledger model seems adequate. For instance, when we are thirsty, drinking some water takes away the thirst and feels like the wiping out of a 'debit'. However, as Singer admits elsewhere:

We are glad that we are hungry if delicious food is on the table before us, and strong sexual desires are fine when we are able to satisfy them.²²

In those cases, the fulfillment of a desire is experienced as creating a credit, something of positive value. People do actually sometimes refuse to ease their hunger in time when it comes, in order to get more satisfaction from a delicious meal that is awaiting them. Thus, it seems that the fulfillment of desires can do more than wiping out a debit. Elsewhere, Singer argues along those lines. Singer claims that the value of the creation and

²¹ See Pluhar (1995, p. 190-217) for a discussion of Singer's earlier strategies.

²² Singer (1998).

satisfaction of a desire depends on whether “the experience as a whole will be desirable or undesirable.”²³ Thus, it seems that Singer’s moral ledger model does not match with people’s experiences of their welfare. This is a disadvantage for a definition of welfare. A welfare scale with positive, negative and neutral welfare matches much better. Of course, there are examples of people who have defended such a pessimistic account of welfare. However, the conclusions those people draw from those accounts are not the conclusions that Singer wishes to draw, as we will see later on in this section.

5.3 Singer’s account of welfare is linked to a particular account of the harm of death

Singer’s argument for the limitation of the scope of the Replaceability Argument forces him to accept a particular account of the harm of death. Singer’s account of welfare implies that the harm of death consists in the fact that present desires cannot be fulfilled anymore. So, for instance, a person is harmed by death because her present desire to go on living and her present desire to finish certain projects are frustrated. Those who do not find this account of the harm of death plausible, have a reason to reject Singer’s account of welfare, because it implies this account of the harm of death.

As I have argued in chapter 3, Singer defines the harm of death in terms of the frustration of ‘hanging’ desires, i.e. of desires that were already present before death and cannot be fulfilled due to death. Instead, utilitarians can define the harm of death as a deprivation of future desire-satisfaction. If that is the case, not only ‘hanging’ desires, but all future desires that the being would have had are taken into account. This is the forbearance view on the harm of death. The forbearance view is the

²³ Singer (1998): 389.

natural view on the harm of death for hedonists, and it is also compatible with other accounts of welfare, such as objective list accounts and desire-satisfaction accounts.

Hence, in order to restrict the scope of the Replaceability Argument, Singer needs his particular account of welfare, which implies a particular account of the harm of death, focused on the lack of fulfillment of the future-directed desires that the being had before its death.

5.4 Singer's account of welfare is ad-hoc in several respects

Another critique against Singer's account of welfare is that it is ad-hoc in several respects. It will become clear what I mean by 'ad hoc' in what follows.

First, as Singer himself admits, it is ad-hoc *where exactly in the modified moral ledger model one should locate the sufficiency level*. The sufficiency level is supposed to indicate the amount of unsatisfied preferences that is declared to count as neutral value. So, the neutral zone is shifted down the scale, in order to avoid the implication that most of us had better not lived. It is arbitrary where exactly in the originally negative area the 'zero' should be put.

Secondly, Singer's account of welfare is ad-hoc in the sense that it is not linked to the rest of his work. *On other occasions, Singer does not indicate in any way that he disagrees with the standard assumption that there is positive, negative and neutral welfare*. To the contrary, elsewhere Singer speaks about welfare as if it could be positive, negative and neutral in a way that seems incompatible with the moral ledger model. For instance, Singer thinks that it is a bad thing knowingly to bring a miserable being into the world, and it is a good thing knowingly to bring a happy being into the world. If the best one could hope for was reaching a zero level in case all one's desires were satisfied, why would it be a good thing to bring additional beings into the

world?²⁴ It seems inconsistent with Singer's account of welfare to claim that it could be a good thing to create happy children. Furthermore, one of the conditions of the Replaceability Argument is that the animals in question have pleasant lives and would have had a life with positive welfare. However, the idea that anyone can score positive on the welfare scale seems inconsistent with Singer's account of welfare.

It is a very uncommon assumption to hold that welfare can only be negative, to different degrees, or neutral. One author who has (after Singer) defended such a view is Christoph Fehige. Fehige's view is called 'anti-frustrationism' (also 'frustrationism'). Fehige claims that whether one has a satisfied desire or no desire at all is equally good. Having an unsatisfied desire, by contrast, is bad.²⁵ The implication of Fehige's model, and a conclusion that Fehige accepts, is that it would be better not to create people. Singer, however, refuses to draw that conclusion: "It is absurd to hold, of even the most fulfilling lives that because some desires could not be satisfied, it would have been better if they had never been lived."²⁶ Singer takes this implication, of which he apparently has not been aware when proposing his model, as a reason for dismissing the original moral ledger model. However, if Singer really believes that having unsatisfied desires is prudentially bad, and having satisfied desires is not better than neutral, the implication that life is not such a good thing simply follows. *Given Singer's account of welfare, the whole idea of shifting the sufficiency line in order to avoid the implication that most of us were better not born seems inconsistent.* That is the third element that is ad-hoc about Singer's account of welfare.

A further reason why I find it ad-hoc is related to Singer's adaptation of a more gradual account of personhood. According to Singer, as personhood comes in degrees, so does (non-) replaceability. The idea would be that those who leave a lot of unsatisfied desires are not

²⁴ Singer (1993): 123.

²⁵ Fehige (1998): 518.

²⁶ Singer (1998): 390-391.

replaceable, while those who leave only a few unsatisfied (short-term) desires are replaceable, to a certain degree. It seems to me though, that when the bringing into existence of a new being cannot compensate for the welfare loss of the killed being, the killed being is *not* replaceable. It is not replaceable, because the welfare loss cannot be replaced as proposed by the Replaceability Argument. If other things than the welfare of the newly created being can be used to compensate for the remaining welfare loss, then this might in principle hold also for the welfare loss that is caused by killing a person. It is unclear where the limit should be. Alternatively *if a (little) remaining welfare loss is deemed acceptable, it is unclear how much loss is supposed to make a being irreplaceable*. It seems that those lines have to be drawn somewhere, and that can only be done in an arbitrary way.

5.5 Singer's account of welfare fails to restrict replaceability in the desired way

The purpose of Singer's argument is to limit the scope of the Replaceability Argument. Therefore it is interesting to point out who would still be replaceable, if one accepted Singer's argument. Even if Singer's account of welfare was accepted, the Replaceability Argument would still be rather inclusive.

On Singer's account, only beings that are highly future-oriented are completely excluded from replaceability. That leaves many animals, human and non-human, replaceable. Human babies and young children are usually not highly future-oriented, neither are many non-human animals. Even great apes, which Singer would otherwise want to count as 'persons'²⁷, are not highly future-oriented. Also other mammals, such as pigs and cows, which are used in animal husbandry and which Singer wants to give the

²⁷ Singer (1993): 132.

'benefit of the doubt' concerning personhood²⁸, live very much by the day. So, the criterion of being highly future-oriented might well be very strict. All those who are not highly future-oriented are supposed to be replaceable to a certain degree. So far, at least normal human adults seem to be excluded.

However, Singer cannot even insist on that. Singer must grab one of two horns of a dilemma: either a being is replaceable, or it had better not lived. Remember that in order to avoid the implication that most of us had better not lived the sufficiency level must be put low enough, well beneath the original 'zero'. However, if it is set low enough for that purpose many people might reach that level during their lives. So, they would reach a level where they still have important unsatisfied preferences, but where a 'sufficient' amount of preference satisfaction can be enjoyed. People who would reach that level during their lives would be out of the negatives, and thus be replaceable. Would there be any sufficiency level that avoided both, the implication of being replaceable and the implication of better not having lived? There would be no such level. To the extent that one implication is avoided, the other is present. Singer suggests the level "at which we consider a life ceases to be worth living, from the perspective of the person leading that life." Singer seems to be after the level of preference satisfaction at which a person judges his life not worth continuing. If all people above this unfortunate level are replaceable, then most of us are.²⁹

Those who would still score negative on Singer's amended welfare scale would not be replaceable. But that would hardly be an advantage. Remember that one of the conditions for the Replaceability Argument to be applicable is that the beings in question have positive welfare. If the life

²⁸ Singer (1993): 132.

²⁹ Alternatively, Singer might refer to the level at which we judge a life not worth starting. Arguably, this level would be higher, closer to zero. Wherever the level is set: either one is replaceable, or one had better not lived.

of a being is miserable, i.e. if it scores negative on the welfare scale, there is no need to replace it. Simply killing that being would be the best one could do (*ceteris paribus*) with respect to overall welfare. So, those with negative future welfare do not even need to be replaced.

Furthermore, if a zero-level were established at a point where a sufficient amount of preferences were satisfied, a being with significantly more satisfied preferences would seem to be on the positive side. This being would have positive welfare, and therefore, its life would be a good thing. Thus, as Singer now admits, bringing into existence such a being could compensate for the killing of another being. Thus, it would no longer be the case that persons would not be replaceable.³⁰ Hence, Singer's account of welfare cannot restrict the Replaceability Argument.

5.6 Singer's most recent strategy

Singer's most recent strategy for restricting the Replaceability Argument to non-persons includes the moral ledger model. This model, as we have seen, implies that all of us had better not lived:

There is, however, one serious objection to this account of preferences: if the creation of each preference is a debit that is cancelled only when the desire is satisfied, it would follow that it is wrong, other things being equal, to bring into existence a child who will on the whole be very happy and will be able to satisfy nearly all, but not quite all, of her preferences. Because everyone has some unsatisfied desires, even the best life anyone can realistically hope to lead is going to leave a small debit in the ledger. The conclusion to be drawn is that it would have been better if none of us had been born!³¹

Thus, according to the moral ledger model, an unsatisfied desire count negatively on the welfare scale, and the fulfillment of a desire can only

³⁰ Singer (1998): 391.

³¹ Singer (forthcoming 2011): 114.

cancel out the negative score. Welfare can never be positive. In order to avoid the implication that all of us had better not lived, Singer has earlier proposed to modify the moral ledger model, as discussed above. This modification did not have the desired results. Singer does not defend this modification any longer.

In order to avoid the unwelcome implication of the moral ledger model, Singer now proposes to accept preference-independent value. Singer now accepts “a notion of value that goes beyond the minimalist basis for preference utilitarianism”. Singer needs to accept that there are two kinds of value: “preference-dependent value, which depends on the existence of beings with preferences and is tied to the preferences of those specific beings, and value that is independent of preferences”.³² The underlying idea seems to be that if desire satisfaction cannot create positive welfare (as the moral ledger model implies), and yet if lives in which nearly all of our preferences are fulfilled are to be considered valuable overall, then the value has to come from somewhere else. Thus, other things, besides preference satisfaction must have value. If that were the case, it would be possible to claim that our lives have been valuable, even if not all of our desires have been fulfilled. After all, desire-satisfaction would not be the only value.

Is the acceptance of desire-independent value new for utilitarians? Obviously, the acceptance of desire-independent value goes beyond the desire-satisfaction account of value, which says that desire-satisfaction is the sole ultimate value. Hedonists, in contrast, accept an account of value that is desire-independent. Hedonists accept that pleasure or happiness is the sole ultimate value. If Singer now wishes to accept a combination of a desire-satisfaction account with a hedonist account of value, an obvious challenge would be to point out how such a combination could work. After all, both accounts might have different implications as to what

³² Singer (forthcoming 2011): 117.

counts as welfare and of why something counts as welfare. Furthermore, Singer would need to point out what the acceptance of this partly hedonist account of welfare would imply for the rest of his moral theory. For instance, hedonism goes together with the foreclosure view on the harm of death.

Accepting a preference-independent account of welfare would have consequences for Singer's account of the harm of death. As we have seen, according to the foreclosure view, every being that is deprived by death of a pleasant future is thereby harmed, no matter whether the being has any conception of its own future, or any future-directed desires. Thus, it seems that even those beings would be harmed by death and therefore would not be replaceable without welfare loss. Singer could, of course, accept the Time-Relative Interest Account, along with the foreclosure view on the harm of death and thus discount the harm caused by the deprivation of a pleasant future for lack of psychological connectedness. In any case, the implications of Singer's new account of welfare would need to be shown.

However, it would not do for Singer, to accept a hedonist account of desire-independent value. After all, the reason Singer accepted the moral ledger model in the first place was to show that persons are not replaceable. They are not replaceable, because killing them leaves a minus on the welfare scale and that minus cannot be compensated by bringing about a new person. This is because, according to the moral ledger model, that new person cannot score positive on the welfare scale either: even if all her preferences are fulfilled the score will be no more than zero. This whole story changes when welfare is considered to consist of desire fulfillment along with pleasure. If pleasure is taken into account as well, a being can score positive on the welfare scale. If it scores positive on the welfare scale, the implication that it better had not lived is avoided. However, in that case it does not hold any longer that this person is irreplaceable. Hence, again, it seems that Singer cannot have it both ways: claiming that we are not replaceable because our welfare is always negative,

and avoiding the implication that all of us had better not lived by accepting that our welfare can be positive.

Now, what else could be a preference-independent value? According to Singer: “We could hold a pluralist view of value and consider that love, friendship, knowledge, and the appreciation of beauty, as well as pleasure or happiness, are all of value.” This can be understood as Singer suggesting an alternative account of welfare, maybe an objective-list account with those items on the list. If this is the case, my above-mentioned comments are applicable here as well: If this implies that welfare can be positive, persons are no longer irreplaceable. Another way of understanding this is that Singer is suggesting an account of value that is not only desire-independent, but also welfare independent. Then, the items Singer mentions would be considered valuable, independent of their contribution to welfare. Again, such an account of value would need to be justified and to be related to the rest of Singer’s theory. The resulting theory would no longer be a welfarist moral theory, and therefore, depending on how one decides on terminological issues, it would arguably not be a utilitarian moral theory any more. Furthermore, as on this view the lives of persons could yield value, however defined, this value could be positive and hence persons would no longer be irreplaceable.

Singer’s acceptance of different kinds of values necessitates the weighing of those values against each other. This would result in a very different form of consequentialism, and it is remarkable that Singer would embrace it just in order to save the moral-ledger model and to reduce the scope of the Replaceability Argument. The view proposed by Singer implies that there is at least something to be said in favor of procreation, even if having children is not what a couple prefers. On a large scale, the view rules out a common decision of the world population to stop reproducing. Indeed, this is an implication that Singer seeks. The moral ledger model would imply that there is nothing wrong with a general voluntary stop on reproduction. Singer makes it explicit that his

stipulation of desire-independent value is intended to morally criticize such a general voluntary stop of reproduction.

At the same time, Singer admits that instead of accepting preference-independent value, a preference utilitarian might just bite the bullet and accept the implication that all of us had better not lived.³³ “Is this too absurd to take seriously?” Singer asks.³⁴ Apparently it seems quite absurd to Singer, because as we have just seen he is prepared to radically revise his moral theory in order to avoid that implication. That lifetime welfare can only be negative or at best neutral is an implication of the moral ledger model, which - as I have already mentioned - resonates nowhere else in Singer’s work. It is remarkable how Singer struggles with some unwelcome implications of his theory. It seems that Singer, as probably most of us, finds it implausible that all of us are in principle replaceable, or that all of us had better not lived. How problematic is it that a moral theory has counter-intuitive implications? This issue will be discussed in chapter 8, in particular in relation to utilitarianism.

Thus Singer’s most recent strategy for limiting the scope of the Replaceability Argument consists of either accepting the unmodified moral ledger model or of accepting it along with desire-independent (and maybe even welfare-independent) value. This strategy seems ad-hoc, as either way the implications do not resonate with the rest of Singer’s work. If the acceptance of desire-independent value implies that a person’s life can have positive value, it does also imply that persons are replaceable. So this move, as theoretically remarkable as it is, does not help. On the other hand, accepting the original moral ledger model implies that the lifetime welfare of all of us can never be positive. This seems to have far-reaching practical implications for utilitarianism, comparable to those mentioned in David Benatar’s *Better Never To Have Been: The Harm of Coming into Existence*. According to Benatar, the world population should

³³ Singer (forthcoming 2011): 118.

³⁴ Singer (forthcoming 2011): 114.

be reduced to zero by ceasing to have children, because life is not worth starting.³⁵

5.7 Summing up

I have explored some implications of Singer's recent efforts to restrict the scope of the Replaceability Argument. The unmodified moral ledger model implies that all of us had better not lived. It rules out the possibility of positive welfare and positive value. Modifying this moral ledger model by shifting down the neutral zone in order to avoid the implication that all of us had better not lived requires the choice between two horns of a dilemma: either we have neutral or even positive welfare and are thus replaceable, or we have negative welfare and our death involves no direct welfare loss at all. In that latter case, we would not even need to be replaced, simply killing us would do. Acceptance of the unmodified moral ledger model in conjunction with acceptance of desire-independent value also confronts us with two horns of a dilemma: either our lives can have positive value or not. If our lives can have positive value, we can be killed and replaced by others whose lives have as much positive value as our future lives would have had. If our lives cannot have positive value, we cannot be replaced without causing greater welfare loss, but then simply killing us would do. Killing us would, *ceteris paribus*, maximize overall value.

All those strategies are ad-hoc: they do not fit well with the rest of Singer's theory. The modified moral-ledger model necessitates drawing the line for neutral welfare somewhere below the point of full preference satisfaction, and that can only be done in an arbitrary way. Furthermore, if personhood and thus replaceability comes in degrees, it seems that beings are more or less easily replaceable, rather than either replaceable or not. If after replacement a remaining welfare loss is accepted, or can be

³⁵ Benatar (2006).

compensated by things besides the newly created being, it is unclear where the line should be drawn between who is and who is not replaceable in that way.

The moral ledger model is based on an account of welfare that does not fit with how people usually experience welfare. Usually people experience welfare as ranging between positive and negative, rather than as only negative or neutral. In addition, Singer's account of welfare is linked to a particular definition of welfare, i.e. the desire-satisfaction account. It is also linked to a particular view on the harm of death, which is focused on hanging desires. It is not compatible with all other possible definitions of welfare or views on the harm of death. Finally, only the unmodified moral ledger model in combination with Singer's view of welfare succeeds in restricting the scope of the Replaceability Argument. This account has the implication that all of us had better not lived.

The conclusion is not that the scope of the Replaceability Argument cannot be restricted. I have rather shown what the implications of such a restriction would be. Proponents of the Replaceability Argument must either accept the implications of one of those strategies, or else accept that all animals, human as well as non-human, are in principle replaceable.

What else does it take to accept the Replaceability Argument? The Replaceability Argument is based on a certain view that is controversial within utilitarianism. An exploration of that controversy will be central to this thesis. In the following section, I will provide a first sketch of the theoretical presuppositions of the Replaceability Argument.

6. The Replaceability Argument is based on a controversial view

The Replaceability Argument rests on an assumption that is disputed within utilitarianism. It presupposes that the possible welfare of possible

individuals must be taken into account in the evaluation of outcomes. This entails not just future individuals, but also *contingent beings*, i.e. *beings that might or might not come to exist*. Let me mention an example in order to provide a first idea of what the relevant issue is. Imagine a couple considering whether or not to have a child. Imagine that the couple would be slightly less happy when they would have a child, as compared to their lives without any child. Assume that no others would be affected in their welfare. Should the couple have a child? This is the choice I have depicted in chapter 2, section 4.2 in a table. If the potential welfare of the contingent child is allowed to enter into the calculation, the couple will be *required* to have a child, if the child can be expected to lead a happy life. This is because all the positive welfare of the child that is caused by the choice to have the child would outweigh the slight loss of welfare for the couple. In contrast, if the potential welfare of the contingent child were not allowed to enter into the calculation, than the woman would be required on utilitarian grounds to remain childless. This is because the couple would be less happy after having a child.

The Replaceability Argument rests on the view that the potential welfare of contingent beings should count in the aggregation of welfare. It takes the welfare of the possible new animal into account when considering whether to kill the existing one. Killing an animal that would have a pleasant future is a welfare loss. The outcome in which the animal is killed and replaced yields as much welfare as the outcome in which the animal is not killed if and only if the welfare of the contingent second animal is taken into account. In the following chapter, I will introduce the controversy about that view. The main part of this thesis will be concerned with exploring the alternative views on this issue and finding out what they entail.

7. Conclusion

The Replaceability Argument proposes a way of compensating the welfare loss that is implied by the killing of animals that would otherwise have had a happy future. The idea of the Replaceability Argument is that the welfare loss due to the killing can be compensated by bringing into existence a new animal, which would not otherwise exist and which will be at least as happy as the killed one would have been. I have proposed the following definition of the Replaceability Argument:

Replaceability Argument: It is permissible to kill an animal, provided that the following conditions are met:

- (a) The future welfare of the animal would have been positive,
- (b) The animal will be replaced, at or after death, by another animal, whose lifetime welfare is at least as positive as the future welfare of the killed animal would have been, and which would not otherwise exist.
- (d) The killing does not have any unbalanced negative side effects (such as fear or suffering for the animal or others).

As far as we can tell, there *might be* practices of animal husbandry which do grant animals sufficiently pleasant lives and which avoid any unbalanced negative side effects of the killing. Probably, those systems would need to score better in terms of animal welfare than current systems of animal-friendly animal husbandry. With the aid of the Replaceability Argument, utilitarianism would be able to justify those practices. Therefore, utilitarianism's support of those practices of 'animal-friendly animal husbandry' depends on the Replaceability Argument.

The Replaceability Argument seems to imply that all animals, including humans, are replaceable. In order to avoid that implication,

Singer has proposed a way to limit the scope of the Replaceability Argument to non-persons. For that purpose, Singer has adopted a certain account of welfare, consisting of a desire-satisfaction account in conjunction with the ‘moral ledger’ model. Only the unmodified moral ledger model implies that persons are not replaceable. It also implies, among other things, that all of us had better not lived and that welfare can never be positive. Those who want to adhere to the Replaceability Argument might accept this implication, or else accept that all animals, including humans, are in principle replaceable.

The Replaceability Argument is based on the Total View, which claims that the potential welfare of contingent beings should count in moral considerations. Contingent beings are beings that do not yet exist and whose existence depends on the moral choice that is being contemplated. It is disputed whether the possible welfare of contingent beings should be taken into account. Exploring the implications of the alternative views on this issue will be the main issue of this thesis. The relevant controversy will be further explored in the following chapter.

5. TOTAL VIEW VERSUS PRIOR EXISTENCE VIEW

1. Introduction

As has briefly been mentioned at the end of the previous chapter, the Replaceability Argument works only if one takes the welfare of the animal into account that will be brought into existence if and only if the existing animal is killed. That animal does not exist at the time that the choice about killing the existing animal is made. It might possibly exist in the future, namely in case the first animal will be killed. As its existence depends on the killing of the other animal, it is a so-called *contingent* being: whether it will exist at all is contingent upon the moral choice under consideration, i.e. upon whether or not the other animal will be killed.¹ That is the definition of ‘contingent being’: it does not exist yet, and

¹ Again, this is an abstract description of what happens in the relevant cases of animal use that I have mentioned in the first section of chapter 1.

whether it will exist at all depends on the moral question under consideration. The animal, whose killing is considered, is therefore *not* a contingent being. It is an actual existing being. The fact, if it is one, that this animal exists only because it is part of a practice in which animals are brought into existence in order to be used and ultimately killed is irrelevant here. The existing animal whose killing is contemplated is not a contingent being, because it is not a possible, but an actually existing being. It already exists while the moral choice about whether or not to kill it is contemplated, therefore its existence does not depend on this choice.² As explained, the Replaceability Argument depends on counting the potential welfare of contingent beings - namely the possible newly created animal - in the aggregation of welfare. Should the possible welfare of contingent beings count? This question is controversial within utilitarianism.

In the present chapter I will introduce the theoretical controversy. The controversy is about whether or not to take into account the possible welfare of contingent beings in the aggregation of welfare (section 2). The two main opposing views are the Total View that counts the possible welfare of contingent beings, and the Prior Existence View that does not count the possible welfare of contingent beings (section 3). As impartiality is an important aspect of utilitarianism, it seems that excluding the welfare of contingent beings is in need of a justification. How can it be justified not to count the welfare of contingent beings, given the utilitarian impartiality? With other words: is the Prior Existence View compatible with utilitarianism at all? The Prior Existence View depends on two assumptions for being a consistent *utilitarian* view. It depends on a particular view about what matters in the evaluation of

² The argument that focuses on the fact that this animal which we are considering killing would not exist if it were not for being used and killed is another argument. This argument, called the Logic of the Larder, will be discussed in chapter 9.

outcomes and on a particular view about whether coming into existence can benefit or harm a being (section4).

2. Whose welfare should count?

As I have already pointed out, the Replaceability Argument only works if the welfare of the not yet existing animal is taken into account. This second animal does not yet exist, but it might possibly come to exist. Therefore, it can be called a 'possible' being: a being that might possibly exist. As is common in this debate, I will also refer to possible beings as potential beings: those who do not yet exist. Whether the welfare of possible beings should count in the aggregation of welfare is controversial within utilitarianism. There are the following three positions:

- a) The welfare of possible beings *should never count*.
- b) The welfare of *some* possible beings should count.
- c) The welfare of possible beings *should always count*.³

Arrhenius introduces the controversy about whether or not to count the welfare of possible beings as follows. He explains that different theoretical approaches count the welfare of sentient beings differently depending on the temporal location or the modal features of their lives.⁴ The criterion of 'temporal location' is easy to grasp. One might argue, for instance, that the welfare of future beings should not

³ Those positions might also be defended in soft versions. It can, for instance, be claimed that the welfare of possible beings or the welfare of a certain category of possible beings should count for *less*, rather than not at all. I will not consider the soft versions here, as they have basically the same problems as the strict ones.

⁴ Arrhenius (2000): 152.

count. Thus, one might hold that only the welfare of presently existing beings counts. The 'present' is a temporal location. So, one can argue that whether a being's welfare counts depends on whether it exists on this temporal location, i.e. whether it presently exists. This is exactly what *presentists* claim. They claim that only the welfare of presently existing beings counts. Hence they opt for (a) above. What about the other criterion that Arrhenius mentions? What does it mean to decide about whether or not to take a being's welfare into account on the basis of *the modal features of a being's life*? That criterion is about the relation of a being's life to reality. It is about whether a being does in fact live, will live or has lived, as compared to might live or surely will never live. Thus, whether the possible welfare of a possible being counts might depend on how likely it is that the potential being will come into existence. That kind of (modal) distinction will be further discussed below.

As explained in chapter 1, for a utilitarian having moral status means that a being is morally considerable. Effects of an action on this being's welfare have to be included in moral considerations. The general utilitarian answer to the question of who has moral status is simple and straightforward: As far as an action (including omission) affects anyone's welfare those effects have to be taken into account. That is why all sentient beings are accorded moral status and plants and stones are not. That is also why it is irrelevant whether a sentient being is black or white, male or female, human or nonhuman. On this basis, it can be argued that future beings should count morally. This means that in our moral decision-making we should take into account the effects on the welfare of beings that live in the future, for instance in choices that affect the environment. Presentists deny this. They can only account for the welfare of future individuals via the interests of (some) present beings. In general, utilitarians agree that the welfare of all sentient beings should equally be taken into account, no matter which individual the being is, where it lives, and when it lives. For a utilitarian, the question whose

welfare to count in the aggregation of welfare is the question who should have moral status.

It is important to note that not all potential beings will once be actual beings. A being that is a possible being at a certain time, might in fact never come to live. For instance, a couple's possible 'next child' might never be conceived. It might never be an actual sentient being. Therefore, potential beings do not always live at a different time. They might not come to exist at all. Often it will be uncertain whether they will come to exist. As I already mentioned, beings that might or might not come to exist are called contingent beings. They are distinguished from 'necessary beings', i.e. beings that live, have lived, or will for sure live. The distinction between a contingent being and a necessary being is one of the modal features of their lives. It is about how 'real' they are or will be. That distinction is arguably different in kind from the distinction on the basis of gender, species, location, or temporal location. A crucial question is whether the distinction is morally relevant.

In what follows, I will introduce two rival views about that issue. One of those views, the Prior Existence View, decides whether or not to count a being's welfare on the basis of the modal features of that being's life, i.e. whether the being is a necessary or contingent being. The other view, the Total View, does not make such a distinction.⁵

⁵ In addition to the Prior Existence View, there are other views that make a distinction on the basis of the modal features of a being's life. The most prominent of those is 'actualism', which counts those beings that do exist, have existed or will exist, and does not count 'merely possible beings', i.e. those who might but will in fact never be actual. I will, however, discuss only the Prior Existence View, because this one is most relevant for the evaluation of the Replaceability Argument, which is the purpose of this thesis. Actualism would imply that when the animal one considers killing is not killed, the possible welfare of the never actually existing replacement animal needs not be taken into account. However, if the animal one considers killing is in fact killed and replaced, the possible welfare of the replacement animal, which is than an 'actual' (it will once actually exist) being, must be taken into account. Thus, when the animal is killed and replaced, this action turns out to be morally right, and when the animal is not killed, this action turns out to be morally right as well. The problem is that what one actually does

3. Prior Existence View versus Total View

In the field of animal ethics, the Replaceability Argument evokes the question of whether to take into account the potential welfare of potential beings in the aggregation of welfare. In animal ethics, two opposing positions on this issue are discussed: the Total View is opposed to the Prior Existence View. The Total View represents category (c) above. It is the (only) view that always and fully counts the possible welfare of possible beings in the aggregation of welfare. In chapter 2 I have mentioned the example of a couple contemplating whether or not to have a child. The child is a possible being: it does not yet exist, but it might possibly exist. Thus, the Total View would take the possible welfare of that possible child into account in the aggregation of welfare. Assuming, as in the example, that the parents would be slightly less happy after having the child and that the child would be very happy, on the Total View the welfare of the child can compensate for the welfare loss to the parents. Thus, on the Total View, having the child would be the better outcome, even if the parents would be less happy with a child. In general, the Total View has the following implication: On the *Total View*,

[...] we aim to increase the total amount of pleasure (and reduce the total amount of pain) and are indifferent whether this is done by increasing the pleasure of existing beings, or increasing the number of beings who exist.⁶

The Total View counts the welfare of possible beings. Hence, when considering whether to kill and replace an animal, the Total View

determines who will or will not be actual. The discussion in animal ethics, and the discussion concerning the Replaceability Argument in general, is limited to the Total View and the Prior Existence View, both to be discussed in this thesis.

⁶ Singer (2005): 103.

would count the possible welfare of the possible next animal that would take the killed animal's place. Therefore, the Total View, as has already been mentioned, would allow the Replaceability Argument.

The rival view to be explored in this thesis is the Prior Existence View.⁷ On the *Prior Existence View*, we

[...] count only beings who already exist, prior to the decision we are taking, or at least will exist independently of that decision.⁸

Two versions of the Prior Existence View have been distinguished. The *restricted Prior Existence View* holds that the potential welfare of as yet non-existent beings should not count in the aggregation of welfare. Only the welfare of beings that do actually exist should count.⁹ This view is equivalent to the above-mentioned presentism. The *extended Prior Existence View* counts not only the welfare of actually existing beings in moral deliberations. Next to the welfare of actually existing beings, it counts the welfare of those who will exist anyway. More precisely, it counts the welfare of those who do already exist and of those who will exist, independent of the decision that is at issue.¹⁰ Thus, the Prior Existence View

distinguishes between beings who will exist *regardless* of the particular action we are now contemplating [...] and those beings whom we are

⁷ Arrhenius (2000) calls this view Necessitarianism, because it counts only necessary beings, i.e. those who exist or will exist independently of the moral choice that is contemplated.

⁸ Singer (2005): 103.

⁹ Pluhar (1995: 190-191) ascribes this view to Salt (1892) and the early Singer (1975).

¹⁰ Singer (2005): 87.

now considering creating.¹¹

Those who exist regardless of the particular action we are now contemplating are accorded moral status. Those whose existence depends on the action that is contemplated are not accorded moral status, i.e. their welfare does not count in the aggregation of welfare.

This framework of discussion can be sketched as follows:

Do potential beings count in moral considerations?

Restricted Prior Existence View: No, only actually existing beings count morally.

or

Extended Prior Existence View: Only existing beings count, as well as those *who will exist, regardless of the choice we are now contemplating*. Potential beings whose coming into existence depends on the choice we are now making do not count.

versus

Total View: All existing beings and all potential beings count morally (including contingent beings).¹²

¹¹ Pluhar (1995): 193. See also Sapontzis (1987): 194-195.

¹² Note that 'Total' here refers to the maximal scope of moral objects. All actual beings along with all potential beings count as moral objects. Their (potential) welfare has to be taken into account in moral considerations. This is not the 'total' of the total view as a method of aggregation. The 'total' of the 'total view' as a method of aggregation refers to the fact that all effects on welfare should be summed-up. This holds independently of how the scope of moral objects has been determined. So, while the 'total view' is

The Restricted Prior Existence View will not concern us in this study. It is usually considered an implausible theory because it implies that we have no moral obligations whatsoever towards future generations. Furthermore, and more importantly, the Restricted Prior Existence View does not play any prominent role in the debate that I will be concerned with. Henceforth, when I talk about the Prior Existence View, I refer to the Extended Prior Existence View. Thus, the opposing views I am concerned with differ on whether the possible welfare of contingent beings should count in the evaluation of outcomes.

4. Are both views coherent utilitarian views?

In chapter 2, I have introduced the basic elements of utilitarianism. I now want to explore whether the Total View and the Prior Existence View are compatible with those basic elements of utilitarianism, and thus whether they can be coherent utilitarian views. Considering the basic elements of utilitarianism, namely welfarism, maximization, impartiality and universality, it seems that the Prior Existence View is in need of justifications for what concerns impartiality. After all, impartiality means that *each counts for one*. It means that the welfare of every affected being has to be taken equally into account. The Prior Existence View implies that the welfare of contingent beings should not be included in the evaluation of an outcome. This seems to be at odds with the requirement of impartiality.

about *how* to aggregate, the ‘Total View’ is about *across whom* to aggregate. I distinguish those views by using small, respectively capital letters.

An example will make this clear. Consider, again, the case of the couple contemplating whether or not to have a child. Here are the distributions of welfare, one for each outcome:

Outcome A: remaining childless		Outcome B: having a child	
	Welfare level		Welfare level
Man	10	Man	9
Woman	10	Woman	9
Child	Ω	Child	9

The child in this case is a contingent being: whether it will exist at all depends on how the moral choice in question is decided.

The Total View would count the welfare of the couple, as well as the possible welfare of the contingent child, and therefore would require choosing outcome B. The Prior Existence View would not count the possible welfare of the contingent child. It would require choosing outcome A, because outcome A maximizes welfare across necessary beings. How can it be justified within utilitarianism to take only necessary beings into account? It seems that this is at odds with the utilitarian duty to impartially maximize welfare. How can a utilitarian neglect the welfare consequences for the child, in this example? After all, if the couple chooses to have a child, that child will exist. It is in need of justification not to include the child's welfare in the aggregation of welfare for this outcome.

A utilitarian can offer such a justification. Utilitarian proponents of the Prior Existence View need two assumptions. First, they need to assume the Person-Affecting Restriction. This particular view about what matters in the evaluation of outcomes has been introduced in chapter 2, section 4.4. It holds that not the quantity of welfare as such matters in

the evaluation of an outcome. Instead, as we have seen, the Person-Affecting Restriction focuses on how people are *affected*. More specifically, it focuses on whether and in how far people are made better off or worse off in an outcome. With other words, an outcome, according to the Person-Affecting Restriction, must be evaluated in terms of harms and benefits to sentient beings.

The other assumption that a utilitarian needs in order to defend the choice for outcome A is that causing a being to exist cannot benefit or harm that being. Consider, again, the case of the couple contemplating whether or not to have a child. Assume that coming into existence does not benefit that child. Granted, the child has a certain positive welfare level during its live. However, according to the assumption we are currently contemplating, coming into existence does not make the child better off, because the counterfactual situation would be nonexistence, and comparisons in terms of welfare between an existent child and the child's non-existence cannot be made.

Those two assumptions together would justify not counting the possible welfare of contingent beings in the aggregation of welfare. Consider the above distribution of welfare for the outcome in which the couple has a child as compared to the outcome in which the couple does not have a child. If the outcomes have to be evaluated in terms of net benefits and if coming into existence does not benefit the child, it follows that the welfare of the child should not be taken into account. After all, the assumption is that an outcome should be evaluated in terms of the net benefits, rather than in terms of the quantity of welfare that it contains. Furthermore, the second assumption is that coming into existence cannot harm or benefit anybody. It follows that the welfare of contingent beings can be left out of consideration. Hence, the Prior Existence View is consistent with utilitarianism, if and only if those assumptions are accepted. Indeed, the Person-Affecting Restriction and

the assumption that causing a being to exist cannot benefit that being together imply the Prior Existence View.¹³ Each of those assumptions will indeed be defended in the following chapters: the assumption that causing a being to exist cannot harm or benefit that being will be defended in chapter 6, and the Person-Affecting Restriction will be defended in chapter 7. They will be defended as possible views, without the pretension to provide any knockdown arguments in favor of those assumptions.

5. Conclusion

The Replaceability Argument only works if the possible welfare of the possible animal that might replace the killed one is taken into account in the aggregation of welfare. This possible animal is a contingent being: it does not yet exist and whether it will exist at all depends on whether the existing animal is killed. It is controversial whether the welfare of contingent beings should be taken into account in the aggregation of welfare. The two main views in this debate are the Total View and the Prior Existence View. The Total View claims that all existing beings and also all possible beings including contingent beings count. The Prior Existence View counts all existing beings as moral objects, as well as those possible beings whose existence does not depend on the moral choice that is contemplated. These are the beings that already exist or will exist, independently of the moral choice that is contemplated: so-called necessary beings. Contingent beings, on the other hand, whose existence depends on the moral choice that is contemplated, do not count morally. So, the Prior Existence View makes a distinction between those who do or will exist, independently of the moral choice that is

¹³ Arrhenius (forthcoming) has denied that those assumptions together imply the Prior Existence View. Arrhenius' argument will be discussed and dismissed in appendix E.

contemplated and those whose existence depends on the moral choice that is contemplated. The first do count morally, the second do not.

At first glance, not counting the welfare of contingent beings in the evaluation of outcomes seems to be at odds with the utilitarian requirement of impartiality. The Prior Existence View needs two assumptions in order to be compatible with utilitarian impartiality. It needs the assumption that the evaluation of outcomes should be concerned with benefits and harms to sentient beings. In addition it needs the assumption that causing a being to exist cannot benefit or harm that being. After all, if what counts is maximizing the net benefit of sentient beings, and if bringing a being into existence cannot benefit or harm that being, then, and only then, it makes sense to neglect the possible welfare of contingent beings in the evaluation of outcomes. Both assumptions will be explored and defended in the following two chapters.

6. CAN EXISTENCE BE BETTER FOR A BEING THAN NON-EXISTENCE?

1. Introduction

As we have seen, the Prior Existence View needs the assumption that causing a being to exist cannot harm or benefit that being. In the present chapter I will defend that assumption.

If we talk about making someone ‘better off’, we implicitly make a comparison about the goodness of two or more different states of affairs for a being. For instance, if I say that accepting a job offer was good for me, I make a comparison between two states of affairs in terms of how good they are for me. In case of bringing someone into existence, if that being was not brought into existence, then what? If the being, say the animal, was not brought into existence, it would not live at all. So, we need to compare a situation in which the animal has a good, bad or neutral life with a situation in which the animal does not exist at all. We need to make this comparison in terms of what would be better for the animal. Can we say that one situation is better for a being than another

situation, when in one of those situations the being does not exist at all? Comparisons between existence and non-existence are special. Can those comparisons be made, and, if so, how? This is what we need to find out, in order to determine whether bringing a being into existence can benefit or harm that being.

How to address the question whether existence can be better for a being than non-existence? In order to give a positive answer to this question one must compare the animal's welfare in an outcome in which the animal exists with its welfare in an outcome in which it does not exist. It is controversial whether, and if so how, existence and non-existence can be compared in terms of welfare. I will first introduce the view that coming into existence can make a being better or worse off (section 2). Then, I will explore the most important challenge to that view which claims that existence and non-existence are incommensurable (section 3). One way of comparing existence to non-existence is to put the welfare level in case of non-existence at the zero level. I critically discuss that strategy (section 4). Those who want to claim that nonexistence counts as zero welfare, rather than no welfare, have to point out that such a comparison is possible and what it entails. They also have to point out for whom non-existence has zero welfare. This cannot be done without violating some conceptual assumptions about what 'better off' and 'worse off' mean (section 5). I will conclude that existence cannot make a being better or worse off than non-existence (section 6).

2. The view that coming into existence can make a being better or worse off

As I said, if we talk about being made better off, we usually compare two situations. Consider the question whether accepting a certain job offer made me better off. In order to answer this question, one might compare my welfare increase that is due to accepting the offer with my welfare

before I did so. Alternatively, one might compare my welfare after having accepted the offer with what my welfare would otherwise have been, had I not accepted the offer. How does this example translate to the question whether coming into existence can benefit a being? The state of affairs in which the being exists would need to be compared with the state of affairs before the being existed. Alternatively, the state of affairs in which the being exists would need to be compared with the counterfactual state of affairs in which the being does not exist. More precisely, we need to compare those states of affairs in terms of what is better *for the animal*. We need to compare those states in terms of prudential value.

A complication with the comparison of a certain state of affairs with what would otherwise have happened, i.e. the so-called counterfactual comparison, is that it might be hard to guess what would otherwise have happened. For instance, it might be hard to guess what had happened, had I not accepted the job offer. When we compare existence with non-existence, this problem does not arise. If we want to compare the animal's coming into existence with the counterfactual situation, we recognize that had the animal not come into existence, it would not exist at all. So, in this particular case, there is really no doubt of what the counter-factual situation would have been. That is an advantage. However, comparisons that involve non-existence raise particular problems, as we will see. The major problem is that we have to compare two states of affairs in terms of prudential value, i.e. in terms of their goodness for the being. Such a comparison is necessary in order to determine whether existence is better for the being. However, in case of non-existence, there is no being at all.

Nils Holtug argues that existence can be better or worse for a being than non-existence. This does not mean that existence as such is conceived to be valuable and always better than non-existence. Rather, existence is viewed as an extrinsic good, which makes the having of intrinsic goods, such as having a life in which the good outweighs the

bad, possible. Holtug's argument is based on a comparative judgment. According to Holtug, if someone's life has an overall positive value to him (in terms of mental states, preference satisfaction or items on an objective list), this can be compared to a state of non-existence. If this person would not exist, no value at all would accrue to him. According to Holtug, if positive value accrues to the person that is better for him or her than if no value at all accrues to him or her. The major objection against this view, to be discussed in the remainder of this chapter, is that the absence of welfare is not commensurable with any score on the welfare scale. Before discussing this objection, which is called the 'incommensurability objection', I will briefly mention two other objections that have been brought forward against the view that coming into existence can benefit a being.

One objection claims that future people cannot be identified and therefore cannot be harmed (or benefited). Future people, it is said, cannot be identified by 'rigid designators', that means by terms that denote one and the same individual in all possible worlds in which it has denotation. For instance, when a couple refers to their future child as 'our future child', this is no rigid designator, because it refers to different possible children, depending on which sperm fertilizes which egg. Holtug counters, however, that rigid designators could in principle be provided, although they would have to be complex, depending on the situation and the criteria for being one and the same person that one assumes. In the case of the couple, one could provide rigid designators for all the children they could have within nine months by itemizing possible combinations of gametes and perhaps specifying certain further conditions.

Still, it could be objected that this only tells us about whom their possible future children are, but we need to identify the child they actually *will* have, if any. That, however, is not necessary for the view that

coming into existence can benefit a being. This view only claims that *if* this person comes into existence, it might thereby benefit or be harmed.¹ Anyway, even the identification of possible future people is not necessary for Holtug's Value of Existence View to hold. "For the view can be taken to mean that, if one brings a person into existence, then whoever she is she may thereby be benefited (or harmed)." For instance we might refer to a couple's possible next child as 'their next child' and claim that "if, in the future, anything satisfies it, this individual might be benefited (or harmed) by coming into existence."² I agree with Holtug that the identifiability problem is not damaging for the Value of Existence View.

Another objection against the view that coming into existence can benefit a being claims that this view is logically incomprehensible. After all, if existence is claimed to be better for a person than non-existence, than non-existence must be worse for that person than existence. This can be considered a conceptual or even logical requirement. But a non-existent person cannot have any properties, so non-existence cannot be worse for her.³

Holtug tries to dismiss this so-called No Properties of the Non-Existent Objection. Holtug defends the following proposition:

P: Non-existence is worse for Jeremy than existence.⁴

Holtug argues as follows. The assumption is that Jeremy's life includes a surplus of positive values and that his non-existence involves no such

¹ Holtug, p. 369.

² Holtug, p. 369.

³ Parfit (1984) and Broome (1993) adhere to this argument.

⁴ Holtug (2001): 372.

values. Both claims are compatible with the No-Properties of the Non-Existent Principle. Holtug goes on:

[...] it seems to be better to have a surplus of positive value than to have no value. Contrariwise, it seems to be worse to have no value than it is to have a surplus of value. This judgment relies only on the nature of positive value and no value.⁵

Now, what can we make of Holtug's claim if we assume that value must always be the value *for someone*? Someone who does not exist cannot 'have' any properties, even not 'zero value' on the welfare scale. After all, the welfare scale is about prudential value. That value is always subject-related. We cannot talk about subject-relative value without reference to any subject. The comparison only holds if value as such counts, but the comparison cannot be about subject-relative value, because in case of non-existence there is no subject, to which value accrues. The comparative judgment cannot be made without ascribing properties to the non-existent.

One can say that the view that coming into existence can benefit a being needs not attach any properties to the being in the state of non-existence, only the absence of properties. But ascribing to the non-existent only the absence of properties is not sufficient for making a comparison between two states of affairs in terms of prudential value. If a comparison is to be made in terms of welfare, what we need are two levels of welfare. However, it has been argued that the absence of value in case of non-existence cannot be measured with the welfare scale: it cannot be indicated on that scale. The most natural point on the scale would be to count the absence of value as zero value. However, it has been objected that this is impossible. This objection has been called the

⁵ Holtug (2001): 373.

Incommensurability Objection. The Incommensurability Objection claims that non-existence and existence are incommensurable in terms of welfare. The welfare scale measures prudential value, but nothing can be good, bad, or neutral for a non-existent being. I will discuss this objection in the following section.

3. Are existence and non-existence commensurable?

It has been argued that it is nonsensical to assume that being caused to exist can benefit someone. Narveson puts this critique as follows:

The issue here is this. Can we say that the creation of a given life – its conception, say – is the conferring of a benefit on the person who results? Suppose that the resulting person is happy; and suppose too that some of his happiness is due to a fortunate genetic make-up. Since conception is not something which happens to someone who is already around, being instead the creation or production of the person in question, we cannot say that the resulting person is better off than he or she would have been if that event had not taken place. The best one could do is set the utility-state of that possible person at zero if not conceived, and then say that his conception has benefited him to, say, the extent of the contribution to his life-long utility attributable to his genetic make-up as opposed to other contributing factors. But that is to equate the condition of someone who is neither happy nor unhappy with that of one who is dead, and that seems wrong; worse yet, it is to equate it with ‘one’ who was never born in the first place, and thus has no identity at all. And that seems very strange indeed.⁶

So, Narveson argues that in order to compare non-existence with existence in terms of utility, it has to get a utility score. Utility is defined

⁶ Narveson, (1978): 48. See also Warren (1996).

in terms of prudential value: welfare. Can non-existence get a welfare score? One might argue that non-existence should score a zero in terms of welfare. After all, in case of non-existence, there is neither positive nor negative welfare. However, zero is the score that is attributed to a person that is neither happy nor unhappy. According to Narveson, it seems odd to equate the situation of somebody who is neither happy nor unhappy with ‘somebody’ who was never born. It seems that Narveson wants to suggest that non-existence cannot be measured on the scale that measures welfare: neither can it be located at the zero point of that scale nor anywhere else. What Narveson depicts as ‘very strange’ is not the use of a welfare scale on which the subjectively experienced happiness of a being in a certain period can be expressed, at least in theory. What Narveson dismisses is the use of the welfare scale for measuring something of, according to him, a completely different dimension, namely death and non-existence.

Heyd has brought forward the same objection. He claims that

[...] there is no way to compare the amount of suffering of states of actual people and the state of non-existence of these people. We should resist the temptation to assign a zero-value to non-existence, thus making it quantitatively commensurable with either the positive or the negative net value of the life of actual people.⁷

So, describing non-existence in terms of the absence of intrinsic value is one thing, but to accord it a place on a scale of subjective value is another thing. Yet, this last thing is necessary for the purpose of comparing which of two states of affairs is more valuable for a person.

⁷ Heyd (1992): 113.

The criticisms of Heyd and Narveson suggest that non-existence cannot be valued in terms of welfare, not even zero welfare. This means that existence and non-existence are considered incommensurable in terms of welfare. This is the Incommensurability Objection. I will illustrate and strengthen the Incommensurability Objection with two examples, the first by Narveson, and the second by myself. Narveson supports his argument with an example from semantics:

The situation here is rather like the one about the treatment of sentences whose referring terms lack reference. Shall we say that “John’s children are ill”, when John has no children, implicitly asserts that he has some and is thus false? Or shall we say that its having a truth-value at all presupposes that he has children, and hence that it is to be reckoned as lacking truth-value at all? In making utilitarian calculations, do we reckon the utility-level of the unborn at zero, and thus reckon conception and birth as promoting utility on all fours with its promotion in beings who exist independently? Or shall we say that utility presupposes subjects of utility, so that the utility of the unconceived may properly be disregarded?⁸

We can read Narveson’s example as an enforcement of his earlier point: death and non-existence are beyond measurement with the welfare scale. Yet, the force of this example depends on one’s position on how sentences whose referring terms lack reference should be treated. It is unclear whether one of the answers is more plausible than the other.

Here is my own example. Imagine that respondents are asked in a questionnaire to indicate their child’s reaction to their parental requests. The child’s reaction should be indicated on a scale ranging from minus 10 (very subversive) to plus 10 (very cooperative). Parents might be able

⁸ Narveson (1978): 48.

to indicate their child's general reaction on the scale: is the child mainly subversive or mainly cooperative, is it, for instance slightly more cooperative or is it nearly always subversive, or do the cooperative and subversive reactions hold the balance? Now, Holtug's argumentation, transferred to this example would be: "If you don't have a child, than you have neither a cooperative nor a subversive child. So, you can safely put the mark at the zero-point." According to Holtug, a life in which positive and negative experiences hold the balance should count for the same as non-existence: a mark at the zero-point. Back to the example, I would say that a respondent without children should not put a mark at the zero-point. He should put his mark at another place, namely 'not applicable'. Maybe the questionnaire is not intended for people without children at all, or else provisions have to be made for answering 'not applicable', or else what we get is great confusion. This suggests that non-existence cannot and should not be measured on the welfare scale.

Recently, new contributions have been made to this dispute. Furthermore, in the field of physics there is a comparable dispute about whether the absence of value is the same as zero value. In the following section I will point out the argumentative strategies that are used in those debates, and draw some conclusions.

4. How to decide whether the absence of any value is the same as zero value?

To recapture, the challenge for those who want to claim that coming into existence can make a being better or worse off is to point out how a state of non-existence can be compared to a state of existence in terms of welfare. After all, such a comparison is necessary in order to compare both states in terms of welfare. There is agreement about the assumption that the non-existent lack properties and thus do not have any positive or negative welfare. The dispute is about whether the absence of welfare

in case of non-existence is the same as having zero welfare in case of existence. After all, if it is the same, both states are comparable. How can it be decided whether the absence of any positive and negative value for a person in case of his or her non-existence is the same as the person having zero welfare in case of his or her existence? Some have tried to argue for one position or the other by way of analogies. For instance, some have drawn an analogy with temperature. The scale with which we measure temperature, just like the welfare scale, has ranges from positive to negative scores, and includes a zero score. Now, it would be nonsense to claim that the color green, or the letter R, is neither hot nor cold, thus it must score zero on the thermometer. However, the claim that the absence of welfare is the same as zero welfare cannot be refuted so easily. The temperature example might not be a good analogy, because the zero-point on the thermometer does not indicate the absence of temperature.

Interestingly, there is a parallel dispute in the field of quantum physics. There, it is disputed whether for instance the claim that the photon has zero mass is different from the claim that the photon has no mass. There are parallel disputes about zero spin and zero charge: is it the absence of spin and charge, or is it zero spin and charge, and what is the difference? For instance, Balashov has argued that the claim that the photon has zero mass is different from the claim that the photon has no mass. According to Balashov the photon has actually zero mass.⁹ Here, of course, our interest is not in quantum physics. However, we are interested in how to decide upon those questions.

In order to decide whether something is an instance of zero value or rather of the absence of value, or both, it makes sense to employ some general principles of empirical inquiry. This is actually what happens in those disputes, in the field of physics, as well as in the field of

⁹ Balashov, J. (1999).

philosophy. The idea is, quite simply, to assemble known instances of a certain property, P, and finding among them putative instances of $P=0$ and, then, opposing such instances to those in which P is absent. I will sketch four general argumentative strategies along those lines. I will show how those argumentative strategies have all been used in the recent dispute about whether the absence of welfare in case of non-existence should count as zero welfare.

1. *The Argument from Composition.* This argument claims that a combination of two or more P-hoods (positive and negative) cannot amount to complete P-lessness. With respect to welfare, this means that we can speak of zero welfare if instances of positive and negative welfare (in terms of enjoyment and suffering, desire satisfaction and frustration or items on an objective list) hold the balance. If positive and negative welfare hold the balance, that is different from not having any welfare. That point can easily be granted. Nobody claims that existing people who score zero on the welfare scale lack welfare. They are considered as having zero welfare. The dispute is about whether there are other instances of zero welfare, besides positive and negative welfare holding the balance. The following arguments are relevant for that dispute.

2. *The Argument from Parity.* This argument holds that if anything else, except the above-mentioned states of combinations of positive and negative P-hoods, claims to have zero P-hood, it must show relevant similarities with the known instances of zero P-hood. Thus, when one claims that the nonexistent have zero welfare, one must point to relevant similarities between the non-existent and those in which positive and negative welfare hold the balance. In the recent debate Bradley has brought forward an argument that is supposed to show relevant similarities between non-existence and situations of positive and negative welfare holding the balance.

Here is Bradley's argument: Bradley claims that someone, let us say Kris, will be indifferent, as far as his welfare is concerned, as to

whether he dies immediately, or lives through a period of zero welfare before he dies. This is supposed to show that putative instances of zero P-hood (non-existence) and known instances of zero P-hood (neutral welfare of existing beings) are relevantly similar. It allegedly does not matter for Kris as far as his prudential value is concerned whether he first lives through a period of zero welfare or dies immediately.¹⁰

What to think about Bradley's argument? Now, Bradley argues that as Kris would value immediate death or death after a period of zero welfare equally, it follows that both future scenarios, are comparable in terms of welfare, and thus that Kris must have zero welfare while dead. However, Luper has replied that from the fact that Kris might be indifferent, as far as his welfare is concerned, between a period of zero welfare, followed by death, or immediate death, it does not follow that Kris has zero welfare while dead. In this sense, one can say that a shorter happy life is less valuable for a person than a longer happy life without assuming that persons have a welfare level when they are dead. The length of the person's life just determines, in this example, how much he gets of what is intrinsically valuable for him, namely welfare.¹¹

To recapture, Bradley argues that Kris is indifferent, as far as his welfare is concerned, between dying immediately and dying after a period of zero welfare. This is supposed to indicate that Kris has zero welfare in both scenarios. Luper shows that the absence of intrinsic value for Kris in case of death could be explained without reference to zero welfare. Thus, the absence of intrinsic value in both cases is not a *relevant* similarity between a state of zero welfare and a state of death. It does not show that Kris must have zero welfare when dead.

3. *The Argument from Unification.* This is the claim that instances of zero P-hood and instances of positive or negative P-hood must show

¹⁰ Bradley (2009): 98 ff.

¹¹ Luper (2009).

similarities, related to the possession of P-hood. After all, they are all instances of P-hood. Furthermore, having zero P should be shown to be different from P-lessness.

Luper has brought forward an argument in order to show that those who have zero welfare have relevant similarities with those who have positive or negative welfare. Luper argues that all those are responsive. According to Luper, only those who are responsive can have welfare. Luper argues that this distinguishes them from those who do not have welfare, such as the non-existent or the dead: the non-existent or dead are not responsive.

Here is Luper's argument: As just said, Luper claims that in order to have a welfare level at a time, a being must be responsive at that time. Responsiveness is defined as follows:

I say that a creature is 'responsive' at t if and only if its well-being may be affected at t – rising if certain conditions are met, and falling if certain other conditions are met.¹²

Luper refers to responsiveness in order to draw a relevant distinction between those who have welfare, including zero welfare, and the non-existent who have no welfare, not even zero welfare.

Bradley challenges this argument by offering a counter-example. According to Bradley, it is metaphysically possible that a dead person is revived. Currently, reviving dead people is not among our capacities, but it might be in the future, according to Bradley. Maybe only some of those who have recently died might be revived, but reviving people, according to Bradley, seems not metaphysically impossible. Bradley's counter-argument, which adheres to the metaphysical possibility of reviving the dead, is problematic. After all, death is defined as the permanent end of all functions of life in an organism. So, if the loss appears to be reversible, the person in question cannot have been dead.

¹² Luper (2007): 244. Cited in Bradley (2009): 102.

So, it seems that Bradley's challenge does not dismiss Luper's claim that in order to have a welfare level at a time, a creature must be responsive at that time.

Bradley thinks that it is metaphysically possible to affect the welfare level of a dead person. Therefore, Bradley concludes that Luper must have in mind a more restricted notion of possibility when he says that in order to be responsive it must be possible that one's welfare level can change. For instance, Luper might mean that the welfare level can change without external help. Bradley, however, dismisses this interpretation as too restrictive. After all, "a person who has been knocked unconscious or who is comatose" has no welfare level according to that interpretation. Yet, according to Bradley, "there seems to be no reason to say that any living person lacks a well-being level." Contrary to Bradley, one might argue that there are reasons to say that living beings at times lack welfare levels. One might say that while unconscious, they have no welfare. For instance, it makes sense to say that a comatose person lacks welfare while being comatose, because at that time she is unconscious and she cannot be affected in her prudential value: nothing can be good, bad, or neutral for her. Whether this is acceptable at all, though, depends on which definition of welfare one accepts. It might be acceptable for hedonists, such as Bradley.

Bradley brings forward another argument against Luper's claim that in order to have a welfare level, a being must be responsive. Bradley suggests that it is possible that someone lives with a permanent welfare level of zero and that this welfare level "could not rise or fall given the person's inherent capacities. Bradley makes up the following example, assuming hedonism:

Imagine Marsha is born without the capacity to feel pleasure or pain, and never develops that capacity; imagine Greg is born with that capacity, but due to his circumstances, he never actually feels any

pleasure or pain. Given Luper's account of responsiveness, Marsha is relevantly like a shoe – she has no well-being level at all – while Greg has a well-being level of zero. This just seems wrong.¹³

Again, assuming hedonism, it makes sense to say that a living person without any capacity of subjective feeling lacks a welfare level. Thus, Marsha might indeed be like a shoe in this respect: she lacks a welfare level. Whether Greg is also relevantly like a shoe depends on the 'circumstances' under which he lives. If he can feel pleasure and pain, but simply happens to have only neutral feelings, he still has a welfare level. Having only neutral feelings is very unlikely to last long in responsive creatures. If Greg's capacity to feel pleasure or pain has been knocked out, then he is relevantly like a shoe as well. The example of Marsha and Greg does not dismiss Luper's claim that having a welfare level is linked to responsiveness.

Bradley proposes a much weaker notion of responsiveness:

R: Person S is responsive at time t only if there is some world w and some time t_n such that S has a positive or negative well-being level at $\langle w, t_n \rangle$.¹⁴

I think that Bradley's notion of responsiveness is too inclusive. It implies that dead people are responsive. Of course, Bradley welcomes this implication, because it allows him to attribute a welfare level of zero to dead people without forcing him to do the same for shoes. However, Bradley has offered no convincing argument for this move. As Luper has already pointed out in reply to Bradley: "Just because something (is the

¹³ Bradley (2009): 103.

¹⁴ Bradley (2009): 104.

sort of thing that) is such that its attaining goods or evils at some time is not impossible, it does not follow that it has a welfare level at some given time, or at every time.”¹⁵ Neither does it follow from the fact that it is metaphysically possible that someone has a welfare level at a time, that this person actually does have a welfare level at that time.

4. *The Argument from Disparity*. This argument claims that if an alleged instance of P=O and an instance that has nothing to do with P-hood differ in a trait known to relate to P-hood, this strongly supports the claim that the first instance really has zero P.

Luper compares a shoe that is not befallen by any positive or negative values at time T and a person that is not befallen by any positive or negative values at time T. According to Luper, if the person has the capacity to experience positive or negative values (even though those capacities might be temporarily blocked), the person still is responsive in Luper’s sense and thus can have welfare. Thus, Luper brings forward a relevant difference between an instance of zero P-hood and a known instance of P-lessness (the shoe). That difference (responsiveness) can be explained in terms of P-hood (i.e. welfare): one is only responsive if one’s welfare level can rise or fall.

As we have seen, responsiveness is a vague concept. It is unclear what it means that one’s welfare level *can* go up or down. For instance, does it mean that it can rise without external help or that the rising is metaphysically conceivable, and why should this be relevant? I suggest that reference to consciousness instead of responsiveness would be clearer: Only conscious beings can have a welfare level, and they can only have it while conscious.

¹⁵ Luper (2009).

Can a subject S have welfare at a time T, when S has no subjective experiences whatsoever at T? I am inclined to say 'no'. That means that not only dead people and the non-existent lack welfare, but also at many periods during their lives might people lack welfare. For instance, when people are temporarily unconscious, due to an accident, they lack welfare at that time. Maybe even during dreamless sleep people lack welfare. I do not think that this is a strange idea. When calculating the welfare of certain periods or even whole lives, those blank periods can simply be left out. So, I think that zero welfare can only exist if positive and negative experiences hold the balance. Just like positive and negative welfare, zero welfare can only exist in the presence of welfare. It never indicates the complete absence of welfare. Thus, zero welfare is never the mere absence or privation of welfare. The two are not the same.

5. For whom would non-existence have neutral value?

So far, as we have seen, the arguments are not on the side of those who claim that non-existence can be accorded zero value on the welfare scale. Let us now move beyond this issue. If non-existence could be accorded zero value on the welfare scale, a further question would come to the fore. Those who assume that non-existence has neutral value must point out for whom it has neutral value. Neutral value is about welfare. Welfare is a prudential value. It is about what is good for someone. For whom would it be good in case of non-existence? It seems that there is nobody in case of non-existence whom this 'value' could befall.

It seems that in case of non-existence the non-existent must have neutral welfare. However, the non-existent can have no properties, not even neutral welfare. If one wants to argue that non-existence can be neutral for a person, then one speaks about a genuine relation, holding between a state of affairs (non-existence) and a subject (the person). For such a claim to be true, both the state of affairs and the subject must

exist in the same world. But the person does not exist in case of her non-existence. Therefore, it cannot be true that her non-existence has zero value for a person. The non-existent is not standing in any relations. So, nothing can be neutral for her.¹⁶

Neither does it work to claim that non-existence has zero-value for the *existent*. After all, something (A) can only be better for somebody than something else (B) if (B) is worse for that being than (A). Hence, if existence is said to be better for the person, this implies that non-existence must be worse for her. However, non-existence cannot be worse for her, as we have just seen. If a person does not exist, she does not stand in any relations. Furthermore, if something (A) is better for a person than (B), this must be true, no matter whether the person actually receives (A) or (B). It must be true, no matter whether (A) or (B) actually obtain. So, if it is claimed that non-existence can be better or worse for the existent, this must be true even if non-existence obtained. However, as we have seen, non-existence cannot be better or worse for the non-existent, because nothing can have any value for the non-existent.

Ascribing zero value or ‘neutral value’ to non-existence is ruled out by these arguments. However, ascribing zero welfare to non-existence is necessary in order to claim that it can be better or worse than existence in terms of the being’s welfare. So, we cannot claim that existence can be better or worse for a being than non-existence.

One might argue that in comparisons with non-existence one should relax the above-mentioned logical requirements. After all, non-existence is a special case. For instance, Parfit has argued that normally saying that someone is made better off requires a comparison between two states of affairs for that person. However, that requirement might be relaxed if one makes comparisons between existence and non-existence:

¹⁶ Bykvist (2007): 343.

We can admit that, in every other kind of case, we benefit someone only if we do what will be better for him. In the case of giving someone life, we understand the special reason why the alternative would not have been worse for him. We might claim that, in this special case, the Requirement need not be met.¹⁷

I think that we should not accept any exceptions to those requirements. Contrary to what has been suggested by proponents of relaxing the requirements, doing so is not necessary for accounting for the harm of death. We can account for the harm of death without making such exceptions. We benefit a person if we help him experience more rather than less welfare in his life. That is why saving someone's life might benefit that person. We harm a person by depriving her of welfare. In order to make those claims, no problematic comparisons between existence and non-existence are needed. What must be compared are rather two lives: a longer one and a shorter one. So, relaxing the requirements would only be necessary for being able to compare non-existence to existence and would be ad hoc.

6. Conclusion

I have argued that bringing an animal into existence cannot harm or benefit that animal. Although someone can of course have a good, bad or neutral life, it is impossible to compare the welfare of a being's life with the welfare that this being would have had, if he or she had not come into existence. The absence of value in case of non-existence is not the same as neutral welfare. Therefore, existence and non-existence are incommensurable. Furthermore, even if non-existence had neutral prudential value, there would be no one for whom it would have that

¹⁷ Parfit (1984): 489.

value. Thus, it cannot be the case that non-existence is better or worse for a person. Therefore, it can also logically not be the case that existence is better or worse for a person than non-existence. Hence, it is not the case that coming into existence can make a person better or worse off than he or she would otherwise have been. Coming into existence cannot be a genuine comparative benefit or harm.

It has been argued in chapter 5 that the Prior Existence View needs to assume that coming into existence cannot benefit or harm a being in order to justify not to count the welfare of contingent beings in the aggregation of welfare. In this chapter I have defended this assumption. Even if one thinks that I do not have provided any knockdown argument in favor of the view that coming into existence cannot harm or benefit a being, at least I have shown that this view can reasonably be upheld. That is all, as far as this issue is concerned, that is needed for the Prior Existence View to be a coherent utilitarian view on the aggregation of welfare.

As we have seen in the previous chapter, the Prior Existence View does also need another assumption in order to be a coherent utilitarian view on the aggregation of welfare. It needs to assume that outcomes should be evaluated in terms of harms and benefits to sentient beings. This assumption will be defended in the following chapter.

7. PERSON-AFFECTING RESTRICTION AND NON-IDENTITY PROBLEM

1. Introduction

I will now turn to the second assumption that the Prior Existence View needs in order to be a coherent utilitarian view on the question across which entities to aggregate. This is the assumption that what matters in the evaluation of outcomes are harms and benefits to sentient beings rather than the quantity of welfare as such. This particular view about what matters in the evaluation of outcomes is known as the Person-Affecting Restriction. It has been introduced in chapter 2, section 4.4. As explained, unlike the Impersonal View, which evaluates outcomes solely on the basis of their intrinsic aspects, i.e. the quantity of welfare that they contain, the Person-Affecting Restriction evaluates outcomes in a comparative way. The Person-Affecting Restriction evaluates outcomes in terms of the harms and benefits that they entail. In order to determine the harms and benefits that an outcome contains, the outcome must be compared to one or more other possible outcomes. In order to

determine which outcome yields most net benefits it matters, for instance, which people exist in each outcome and whether they would have existed in the other outcome as well.

The major challenge to the Person-Affecting Restriction is the fact that in some cases the outcome that contains most welfare and that at first glance seems to be the most beneficial does in fact not involve a benefit to any particular being. It is a challenge for the Person-Affecting Restriction how to deal with those cases. If what one does has implications for who will exist, different beings exist in different outcomes. In those cases, when no particular person exists in both outcomes, it seems impossible to say that the better outcome, i.e. the outcome that contains more welfare, is better for any particular person. After all, the persons that exist in one outcome do not exist in the other outcome. This challenge to the Person-Affecting Restriction has been labeled the Non-Identity Problem, because the beings that exist in both outcomes are not identical. No particular individual that exists in one outcome could possibly exist in the other outcome. In those cases, the outcome that seems to be preferable in terms of welfare does not benefit anybody (section 2).

Are those cases a reason for dismissing the Person-Affecting Restriction? No, they are not, or so I will argue. A wide interpretation of the Person-Affecting Restriction avoids the Non-Identity Problem. According to the Wide Person-Affecting Restriction what matters in the evaluation of outcomes are indeed *harms and benefits to sentient beings*. However the view does not focus on harms and benefits to *particular individuals*, but rather on harms and benefits to sentient beings *whoever they are* (section 3). The Wide Person-Affecting Restriction can be defended against possible criticism (section 4). It has advantages above the Narrow Person-Affecting Restriction (section 5). The Wide Person-Affecting Restriction is a possible view about what matters in the evaluation of outcomes.

Thus, the Wide Person-Affecting Restriction can avoid the Non-Identity Problem. Hence, the Non-Identity Problem is no reason for dismissing the view that outcomes should be evaluated in terms of harms and benefits to sentient beings rather than in terms of the quantity of welfare as such. Therefore, the second assumption that the Prior Existence View needs in order to be a coherent view within utilitarianism, namely the assumption that outcomes should be evaluated in terms of harms and benefits for sentient beings, can be defended. Given this conclusion and the conclusion of the previous chapter, both necessary assumptions for Prior Existence Utilitarianism can be defended. Hence, the Prior Existence View is a coherent utilitarian view about whose welfare to count in the aggregation of welfare (section 6).

2. The Non-Identity Problem

As we have seen, the Person-Affecting Restriction evaluates actions in terms of the difference they make for the welfare of sentient beings. What matters, according to the Person-Affecting Restriction, is in how far sentient beings are made better or worse off. I will now introduce what counts as the major challenge to that view on the evaluation of outcomes.

Derek Parfit has famously explored what he called the Non-Identity Problem. This problem refers to the fact that some of our choices, which affect the welfare of people, also affect their identity. With other words, the compared outcomes differ not only as to how well off the existing people are, but also as to who exists. Parfit points this out in the example of the 14-year-old girl:

The 14-year-old Girl: A 14-year-old girl has a baby. The baby has a difficult start in life and therefore also some problems later on in its life. So her child has a somewhat diminished quality of life overall, but a life

that is still overall pleasant. Had the 14-year-old girl waited several years before becoming pregnant, she would have been able to give her baby a better start in life and her child would have a better life overall.¹

Now, if one assumes that one's personal identity depends at least in part on the genetic material of which one is made, and given the fact that sperm (and eggs) have short lives, it follows that the time of conception is relevant for who exactly will exist. In the example, the girl would have a different child, growing from different sperm and egg, if she delayed conception. Therefore, it seems impossible to condemn the 14-year-old girl for having the child with reference to the interests or the welfare of the child. After all, the child has a pleasant life, and if the girl had waited, this child would not exist at all.

Here is another case that illustrates the Non-Identity Problem, let us call it the case of the Two Women:

The case of the Two Women: Two women plan to become pregnant. Woman A is already pregnant when she receives good and bad news. The bad news is that her child will be born with a defect that will lower its quality of life, although it will still be able to have a worthwhile life. The good news is that the woman can take a pill to cure this defect. Woman B gets good and bad news when she is about to stop with contraception. The bad news is that she has a medical condition that will result in a child with a defect that will lower its quality of life, just as serious as that, which woman A's child would have if she would not take

¹ See Parfit (1984): 358.

the pill. The good news is that woman B will have a healthy baby if she waits three months before becoming pregnant.²

Now, it seems that in order to do the best thing for her child, woman A should take the pill in order to have a healthy baby. Likewise, woman B should wait three months in order to have a healthy baby. However, the difference is that by taking the pill woman A would make her child (not yet born, but already existing in her womb) better off. Woman B, per contrast, would not benefit any baby. By delaying conception, she would have a different baby than she would otherwise have had. The Person-Affecting Restriction is focused on harms and benefits to sentient beings. It seems that in case of woman B delaying conception would not benefit anybody. After all, there is no particular being made better off by having a different baby, or so it seems. Hence, it seems that the Person-Affecting Restriction would require that woman A takes the pill, but not that woman B delays conception.

The Non-Identity Problem can also show-up on a larger scale. Major social policy decisions that affect the welfare of future generations also affect the identity of the people who are going to exist. Consequently the future populations that are at stake when deciding whether to implement a major social policy consist of different people. Imagine that the current population adopts a policy that causes disadvantages for the future generation, for instance in terms of depleted resources or pollution. Assuming that those future people will have lives that are pleasant overall and that they would not have existed had that policy not been adopted, it seems that adopting that policy does not harm them.³

² This case is based on Parfit (1984): 367. See also Singer (1993): 123.

³ Although cases like this one are commonly brought forward in order to illustrate the Non-Identity Problem, even by Parfit, they are less clear than both before-mentioned

Thus, the “Non-Identity Problem” refers to the observation that in different people choices the outcome that is intuitively considered better (such as delaying conception in case of woman B or the 14-year-old girl) does not make any particular person worse off. With other words: it seems impossible to condemn the action that produces less welfare or to recommend the action that produces more welfare in person-affecting terms. For instance, the 14-year-old girl would not benefit any particular child by delaying conception. If she delayed conception, she would have a different child. Note that the Impersonal View would not lead to the Non-Identity Problem. The Impersonal View would clearly evaluate the outcome that contains more welfare as better in terms of welfare. As delaying conception would result in more welfare, utilitarianism in conjunction with the Impersonal View would require just this. The Non-Identity Problem is a challenge for the Person-Affecting View on the evaluation of outcomes.

Let me recapture my argumentative strategy for this chapter, and its relation to previous chapters. A central question of this thesis is whether utilitarianism implies the Replaceability Argument. I have argued that the Total View on whose welfare to take into account in the aggregation of welfare implies the Replaceability Argument, while the Prior Existence View does not. The Prior Existence View depends on acceptance of the Person-Affecting Restriction on the evaluation of outcomes in order to be a coherent utilitarian view. The Non-Identity Problem counts as the major challenge to the Person-Affecting Restriction. Therefore, I will now explore whether and how the Person-Affecting Restriction can avoid the Non-Identity Problem. I will argue that the Non-Identity Problem can be avoided if the Person-Affecting Restriction is interpreted in a wider way.

cases. It is less clear that it holds for each individual that he or she could not have existed in the other outcome and thus be personally harmed or benefited by the policy choice. Roberts (2007) has made this point.

3. The Wide Person-Affecting Restriction

As explained above, the Non-Identity Problem has been brought forward as the major challenge to the view that outcomes should be evaluated in terms of benefits and harms to sentient beings. It seems that proponents of the Person-Affecting Restriction have nothing to say in favor of the outcome that maximizes welfare in different people choices. After all, it seems that the choice of this outcome does not benefit anybody, as different beings exist in both outcomes.

But are the beings that exist in both outcomes – such as the baby that the 14-year-old girl could have now or the baby that she could have a couple of years later - really different, in a morally relevant sense? It is true that the babies would be different individuals, genetically speaking. But is this the sense of identity that is relevant here? In what follows, I will introduce a version of the Person-Affecting Restriction that implies a looser notion of identity (3.1). I will focus on the concept of identity and argue that different notions of identity are relevant in different contexts. When questions about moral responsibility are asked, as in non-identity cases, the notion of identity should capture what is morally salient. I will argue that what matters in non-identity cases is not genetic make-up. A much looser notion of identity is adequate (3.2). As the Wide Person-Affecting Restriction evaluates outcomes in terms of benefits and harms for sentient beings without being focused on genetically distinguished individuals, it can accept that in case of the 14-year-old girl, having a child later would be better *for that child*. Thus, utilitarianism in conjunction with the Wide Person-Affecting Restriction would require delaying conception and thus choosing the outcome that maximizes welfare. It would avoid or circumvent the Non-Identity Problem.

3.1 *De dicto* betterness

As we have seen, according to the Person-Affecting Restriction what matters in the evaluation of outcomes is not the quantity of welfare, but rather harms and benefits to sentient beings. Two interpretations of the Person-Affecting Restriction can be distinguished. In one interpretation the focus is on harms and benefits for *particular individuals*. In the other interpretation the focus is on harms and benefits to sentient beings, *whoever they are*. On the basis of this distinction two possible interpretations of the claim that an outcome can only be better (worse) if it is better (worse) for sentient beings can be distinguished. On the first, narrow, interpretation what is good must be good for someone and what is bad must be bad for someone, where ‘someone’ is understood as one or more *particular persons*.⁴ Thus, we can define the Narrow Person-Affecting Restriction as follows:

The Narrow Person-Affecting Restriction: An outcome is better (worse) if and only if it is better (worse) for at least one *particular person*.

On the wide interpretation what is worse for people should be understood in a wider sense, without reference to particular people that must be harmed or benefited. Parfit has defined this wide interpretation of the Person-Affecting Restriction as follows:

The Wide Person-Affecting Principle: Call outcome X ‘worse for people’ in the wide sense if the occurrence of X would be less good for the X-people than the occurrence of Y would be for the Y-people.⁵

⁴ Roberts (1998) accepts this view.

⁵ Parfit (1984): 396.

Thus, we can define the wide Person-Affecting Restriction as follows:

The Wide Person-Affecting Restriction: An outcome is better (worse) if and only if it is better (worse) for people, *whoever they are*.

The Wide Person-Affecting Restriction differs from the Impersonal View, because it is person-affecting: it holds that in order to be morally better, an outcome must be better for people. It differs also from the Narrow Person-Affecting Restriction, because it is not concerned with how particular individuals as such are affected. It seems that individual identity does not matter, according to the Wide Person-Affecting Restriction. More precisely, as we will see in what follows, identity matters less, or matters in a different way. My defense of the Wide Person-Affecting Restriction focuses on the concept of identity. I will propose, as it were, to employ a looser notion of identity in the relevant contexts.

The Wide Person-Affecting Restriction is in line with an interpretation of utilitarianism as requiring the maximization of welfare *for sentient beings whoever they are*. Usually, utilitarianism is not preoccupied with how particular individuals *as particular individuals* fare in the outcomes; what matters is rather the overall welfare of sentient beings. Thus, the Wide Person-Affecting Restriction fits particularly well with utilitarianism, resulting in Wide Person-Affecting Utilitarianism. How would Wide Person-Affecting Utilitarianism deal with non-identity cases? Consider the case of the 14-year-old girl. Wide Person-Affecting Utilitarianism would require postponing conception, because this would be better for *her next child, whoever it will be*.

The different concerns of the Wide and the Narrow Person-Affecting Restriction with regard to non-identity cases have a parallel in the distinction that philosophers of language make between ‘*de dicto*’ and ‘*de re*’. For instance, one can distinguish *de re* betterness from *de dicto* betterness:

De re betterness: An outcome O1 is *de re* better for someone S, than an alternative outcome O2, if and only if *the actual S* is better off in O1 than in O2.

De dicto betterness: Comparing two outcomes O1 and O2, O1 is *de dicto* better for S if and only if the thing that is S in O1 is better off in O1 than the thing that is S in O2 is in O2.⁶

In case of *de dicto* betterness, the S can be filled in by a descriptive referring term, such as ‘her next child’ or ‘the 35th president of the United States’. Thus, in the case of the 14-year-old girl, one can say that delaying conception is better than conceiving immediately, because her next child will be better off when she delays conception.

Obviously, very often, making things *de dicto* better is not morally praiseworthy or relevant in any way. One can make things *de dicto* better for someone, without making it *de re* better. For instance, if I exchange my sick dog for a healthy one, by disposing of the sick one and buying a new one, I have made things *de dicto* better for my dog: the being that is ‘my dog’ after those actions is better off than the being that was ‘my dog’

⁶ Hare (2007): 514.

before. ‘My dog’ in the *de dicto* sense is better off after the exchange.⁷ Arguably, it is not a good thing that I made my dog better in that sense. However:

[...] it does not follow from the fact that *de dicto* betterness is not *always* morally significant that it is *never* so. Sometimes it is appropriate to expect people to care about making things *de dicto* better in some ways.⁸

So, the question is whether it can be shown that it is indeed *de dicto* betterness (or worseness) that matters in non-identity cases.⁹

3.2 A morally relevant notion of identity

The Non-Identity Problem assumes a particular concept of identity. It assumes the Kripke/Parfit Origin View, which is a broadly accepted view about personal identity. It is an essentialist position, claiming that in any possible world in which an individual exists, it has the same genetic origins. One implication of this view is that an individual cannot be harmed by pre-conception genetic choices. After all, an individual is harmed, if and only if it is caused to be worse off than it would otherwise have been. In case of pre-conception genetic choices the individual in question would not have existed, according to this essentialist view, if other choices had been made. The reason is that if any other choice had been made, the individual that would exist would be a different individual with different genes.

⁷ Hare (2007) brings forward a more funny example about Zsa Zsa Gabor having found a method of keeping her husband young.

⁸ Hare (2007): 516.

⁹ Hare (2007) defends that claim by pointing to an analogy with other cases in which *de dicto* betterness is supposed to matter. However, it has been shown that this analogy fails (Wasserman 2009 A).

This concept of identity is controversial. It is controversial whether identity depends on exact genetic origin. For instance, genetically identical twins per definition possess the same genetic make-up. However, it is not obvious that they are ‘identical’ in all relevant senses. This example suggests that ‘identity’ is not a clear-cut concept. What do we mean when we talk about identity? Obviously, nobody would want to suggest that genetically identical twins are the same individuals. Not only are they numerically different, in the sense that one can speak about several identical white plates. But also it seems that identical twins have substantially different identities, even though they share some features, such as their genetic make-up and their home and their family, which are more or less determinant for who each of them is.

For instance, Wrigley argues that it is possible to allow “identity claims for an individual despite changes in genetic origins”.¹⁰ Wrigley argues that the counterpart of a subject ‘S’ in another possible world, even if it might have a different genetic origin, can still be S. Wrigley’s view on personal identity is based on Lewis’s (1986) Genuine Modal Realism. This view identifies a person’s counterparts across possible worlds by using a similarity relation, rather than the strict identity of genetic origin. According to Wrigley, identity is not tied to exact genetic origin. So, the same individual could have existed with a slightly different genetic make-up. That makes it possible to claim that an individual has been harmed (or benefited) by pre-conception genetic choices. Had those changes in its genetic material not occurred, the same individual might have lived, and might have been better or worse off:

¹⁰ Wrigley (2006): 507.

[...] [A]n individual can lay claim to having been harmed [or benefited, TV] by choices regarding its genetic characteristics prior to conception by appeal to that person's counterparts across possible worlds.¹¹

Relevant counterparts of the individual, on that account, are only individuals that have the same identity, for only then can we compare the possible worlds in order to assess *personal harm*. The crucial question then is: which counterparts in other possible worlds count as relevantly similar? Which counterparts can be said to have the same identity? This, according to Wrigley, is determined by context. The context determines the range of possible worlds and the selection of the particular counterpart relation(s) that are relevant.¹²

I think that this is an interesting account, according to which identity is not tied to exact genetic origins. For instance, when I think about my daughter, I have a certain idea about what makes her the very particular, unique person that she is. Genetic identity plays a role. It matters for her identity that she is the child of her parents, and that she looks the way she looks. However, does this mean that it is crucial for who she is that she has the exact genetic make-up that she actually has? If her genetic make-up would be slightly different, even unnoticeably from the outside, would that mean that it would not be her and that it would instead be a different child? I think it is conceivable that exact genetic makeup is not crucial for who she is. In particular, differences in genetic identity seem to be less important when they are only on the level of genes (genotypic). Differences seem less identity determining when they are not expressed (phenotypic) in an individual's looks or behavior. But even if my daughter looked slightly different, for instance if her ears

¹¹ Wrigley (2006): 507.

¹² Wrigley (2006): 512.

where one centimeter bigger, it is not obvious that she would have been a different individual.

A looser account of identity is conceivable. According to such an account, the same individual might have existed with a slightly different set of genes. A problem with this account is that it is not clear how much variation of genetic material is tolerable before a being would not be an individual's counterpart anymore but rather a different individual. It is not clear what it means that it is determined by context whether a being that would have existed in a different world would have been an individual's counterpart or rather a different individual.

As we have seen, an alternative to the genetic essentialist view on which the Non-Identity Problem is based, is the view that identity is vague. Identity can be conceived as dependent on a certain characteristic or (more likely) on a set of characteristics. Sufficient changes in this set of characteristics will change a person's identity. Very minor changes, such as nearly undetectable changes of genetic make-up, do not change a person's identity. In case of somewhat bigger changes, it might be unclear whether a person's identity has changed:

Where P's identity is determined by C, and C is a property or set of properties that can vary by degrees, P's identity will be vague.¹³

Another way of putting this is to say that identity is 'not perfectly fragile': It won't break down, as it were, after a minor change in the relevant property or set of properties. Along that route, both Wrigley, and, as we will see, also Wolf have engaged with the Non-Identity Problem. I consider this a promising route.

¹³ Wolf (2009): 103.

If identity is vague, how to determine which change of properties constitutes a change of identity? As an answer to this question, it is possible to just stipulate where the limits are. Thus it is possible to *stipulate* that any change of genetic material or of time and circumstances of conception leads to the existence of a different child. So, it is possible to make agreements about, or offer definitions about the relevant boundaries. An interesting idea of Wolf is that identity is not a univocal concept and that for different purposes different definitions of identity are adequate:

[...] it is worth noting that identity is not a univocal concept. [...] When we employ the term, it will be crucial to identify what is relevant about identity in the particular context so that we will use the right concept of identity, and not the wrong one. There is no single concept of “Identity” that we can appropriately employ in all circumstances.¹⁴

Thus, just like Wrigley, Wolf appeals to the context. Wrigley appeals to context for determining which counterparts of a being in other possible worlds would still count as the same individual. So, the context determines the limits of the range of differences that are within a particular personal identity. Wolf appeals to context for determining which criteria are relevant for personal identity. It is in fact common that different notions of identity are used and thought to be relevant in different contexts. For instance, when we want to know who can rightly inherit property under the terms of law, a description of the person as “John and Mitzi’s fourth child” might be sufficient.¹⁵ If someone wants to know towards whom she has the duty to leave the toilet tidy, a suitable answer might be “towards those who are going to use the toilet

¹⁴ Wolf (2009): 104.

¹⁵ Wolf (2009): 104.

after you". It seems that those answers appeal to particular practices in which a particular understanding of identity seems relevant. In every case, it must be determined which definition of identity is appropriate.

The question that interests us here is which definition of identity should be applied in non-identity cases. Non-identity cases confront us with moral questions. It's exactly the moral implications that are drawn from those cases that are under scrutiny. Therefore, according to Wolf, the relevant conception of identity for those cases should be determined by moral theory. According to Wolf, in *Non-Identity Cases* genetic identity is irrelevant from a moral point of view.¹⁶ I interpret this as meaning that whatever criterion one appeals to, it has to be made plausible that it is morally relevant. It can be disputed that genetic essentialism is the morally relevant understanding of identity in non-identity cases.

Saying that different definitions of identity are possible and that it depends on the context which of those concepts is relevant is one thing. Justifying that a particular concept of identity is to be chosen in non-identity cases is another thing. How can one argue for a loose understanding of identity being appropriate in non-identity cases? How can one argue against it? How to determine which identity-determining features are relevant from a moral point of view? One answer to this question is that the intuition behind the Non-Identity Problem is that strict genetic identity is not what matters in those cases. The intuition that strict genetic identity is not what matters leads to the judgment that the 14-year-old girl should delay conception. But then, what is the moral relevance of such intuitions?

Which identity-determining feature is considered morally relevant depends on one's understanding of morality and on the moral theory that one adopts. From a utilitarian perspective, a looser understanding of

¹⁶ Wolf (2009): 107.

identity in non-identity cases is at least possible. That is all that I need to have established here. A utilitarian can adopt a minimal conception of identity as adequate in those cases. This would imply that genetic make-up, personalities or other characteristics of individuals are irrelevant when we consider whether we, members of the present generation, should needlessly destroy the natural resources that the members of future generations will need. Similarly, when a couple contemplates their duties concerning the welfare of their possible child, the child's exact genetic make-up can be deemed irrelevant. Also in more conventional cases, possible victims and claimants are identified in vague and minimalist ways. For instance, the duty not to roll boulders from the top of a mountain is owed towards those who might climb the path behind you. The fact that we can affect 'their' welfare can be considered relevant from a utilitarian perspective. Our impact on their welfare is what matters, or so a utilitarian can claim. According to that wide interpretation of the Person-Affecting Restriction it does not matter who exactly 'they' are.

Note that proponents of this approach do not present an alternative definition of identity. Instead, they claim that identity can be defined in various ways. Depending on the context, different definitions are adequate.

A utilitarian can claim that the Non-Identity Problem is problematic precisely because the different genetic identities of future people should not matter from the moral point of view.¹⁷ According to that view, genetic definitions of identity are not relevant in non-identity cases. As Wolf explains:

Where the question involves responsibility, our theory concerning the identity-determining characteristics of future persons should capture

¹⁷ Wolf (2009): 98.

what is significant about them from the *moral* point of view, and this might be quite different from other senses of identity that we use to individuate persons for other reasons, or within other theories or projects.¹⁸

So, utilitarians are in a position to claim that a wide notion of identity is adequate in non-identity cases. With other words, they can appeal to a de-dicto notion of betterness in non-identity cases. For instance, they can claim that in the case of the 14-year-old girl, what matters is how her future child will be affected, whoever this will be. In that way, by appealing to a Wide Person-Affecting Restriction, the Non-Identity Problem can be circumvented. There is a reason to choose the better outcome in terms of welfare, without abandoning the idea that outcomes should be evaluated in terms of their effects on sentient beings.

To sum up, the Narrow Person-Affecting Restriction does indeed run into the Non-Identity Problem. It holds that an outcome is better if and only if it is better for a particular person. This causes problems in different people choices, where the apparently optimal outcome in terms of welfare cannot be said to be better for any particular person. The Wide Person-Affecting Restriction holds that the better outcome must be better for persons, whoever they are. The Wide Person-Affecting Restriction does not run into the Non-Identity Problem, because the optimal outcome in different people choices is indeed ‘better for people’ in this wide sense.¹⁹

¹⁸ Wolf (2009): 105.

¹⁹ Caspar Hare has brought forward a rather weird and ingenious argument, which builds on the idea that identities are vague. Hare’s argument supports the conclusion that choosing the optimal outcome in non-identity cases is required and hence lends additional support to the Wide Person-Affecting Restriction. A summary and discussion of Hare’s argument can be found in Appendix F.

4. Possible criticisms

I will now discuss two possible criticisms that might be brought forward against the Wide Person-Affecting Restriction. First, it might be argued that individuals in the wide, *de dicto* sense cannot have interests nor can they be harmed and benefited (section 4.1). Secondly, it might be argued that the Wide Person-Affecting Restriction opens the door to the Replaceability Argument and thus brings back in what the Prior Existence View was supposed to avoid (section 4.2).

4.1 Individuals in the wide sense cannot have interests

It can be argued against the Wide Person-Affecting Restriction that a wide notion of identity is not helpful. After all, individuals in the wide, *de dicto* sense, such as ‘the 14-year-old girl’s next child’ cannot have interests, nor can they be better or worse off. Only concrete particular (present or future) individuals can. One can have interests in virtue of being a certain type or description, but the interests relate to a particular person, *de re*. Therefore, it can be argued that recourse to the *de-dicto* sense makes no sense.²⁰ It is of course true that only particular individuals can have interests or be better or worse off. Even if we say in the case of the 14-year-old girl that ‘her next child’ will be better off if she waits a couple of years before having it, it is not the case that a particular individual (*de re*) of flesh and blood would be better or worse off depending on what the 14-year-old girl does. Proponents of the Wide Person-Affecting Restriction are of course aware of that fact.

In the same vein, the Wide Person-Affecting Restriction can be criticized because it does not explain why a child born to the 14-year old

²⁰ Heyd (2009): 13.

girl would have any reason to complain. Utilitarianism in conjunction with the Wide Person-Affecting restriction would require that the girl should do the best for 'her next child', whoever it will be. However, utilitarianism in conjunction with the Wide Person-Affecting Restriction would not establish that any child born to the 14-year-old mother had any reason to complain. That particular child would not have been better off if the girl had waited before having a child. Rather, it would then not have existed at all and the girl would have had another child. Even though, arguably, personal identity is not relevant for the moral duty of the mother concerning the time of conception, it seems still to be relevant with respect to the child's reasons to complain about her mother not having waited a couple of years. Thus, there is still a sense in which the non-identity problem remains intact.

I do not think that this is very disturbing from a utilitarian perspective. It can be established with the aid of the Wide Person-Affecting Restriction that the mother ought to strive for the optimal outcome in terms of welfare. Utilitarianism is not concerned with particular individuals *as particular individuals*. What counts is the overall maximization of welfare (possibly defined as aggregate net benefit) for sentient beings. It can be granted that the child in question cannot complain with reference to its particular interests. Still, the mother should have acted differently, morally speaking.²¹

So, even if the optimal outcome in terms of welfare is chosen in non-identity cases, it is not the case that a particular individual (*de re*) is made better off. Likewise, if the worse outcome is chosen, it is not the

²¹ Weinberg (2002) proposes an approach that is compatible with the Wide Person-Affecting Restriction and that is also able to tackle this aspect of the Non-Identity Problem. Thus, Weinberg claims to be able to explain why the worse-off child in question can rightfully complain to her mother. Weinberg's approach is interesting. However, I will not discuss it here. After all, the pressing problem, from a utilitarian perspective, is that in non-identity cases the choice of the optimal outcome in terms of welfare should be morally required. The Wide Person-Affecting View can show that indeed the optimal outcome should be chosen and can explain and justify why this is so.

case that a particular individual (*de re*) is harmed. Why should this be problematic? Claiming that this is problematic seems to presuppose the Narrow Person-Affecting Restriction, i.e. the view that outcomes should be evaluated in terms of their harms and benefits for particular individuals (*de re*). However, as I have argued, utilitarians need not agree with that. In section 5 of this chapter I will argue that actually the Wide Person-Affecting Restriction seems to be more in line with utilitarianism than the Narrow Person-Affecting Restriction.

4.2 The Wide Person-Affecting Restriction invites the Replaceability Argument

Does the Wide Person-Affecting Restriction open the door to the Replaceability Argument? It might seem that it does, because it is not concerned with particular individuals (*de re*) but rather with ‘roles’, so to say (*de dicto*). This might bring to mind the Replaceability Argument, which can be framed, in line with the de-dicto/ *de-re* distinction, as follows: “Farmer Johnson’s pigs (whoever they are) are well off. He does regularly kill some (*de re*) and replace them by others (*de re*). His pigs (*de dicto*) are always doing fine.”

If the Wide Person-Affecting Restriction brought back the Replaceability Argument, that would be interesting news. After all, so far I have claimed that the Prior Existence View does not imply the Replaceability Argument. The whole purpose of introducing the Wide Person-Affecting Restriction was to show that the Prior Existence View could be a coherent utilitarian view on whose welfare to take into account. We have seen that the Prior Existence View needs to assume the Person-Affecting Restriction in order to be a coherent utilitarian view. The Wide Person-Affecting Restriction has been introduced in order to avoid the Non-Identity Problem, which is the major challenge to the Narrow Person-Affecting Restriction. Now, if the Wide Person-Affecting Restriction sanctioned the Replaceability Argument it would

seem that even the Prior Existence View would not be able to avoid it (because the Prior Existence View needs the Wide Person Affecting Restriction in order to be a coherent utilitarian view and to avoid the Non-Identity Problem).

Therefore it is crucial to point out that the Wide Person-Affecting Restriction needs not sanction replacing animals. After all, the Wide Person-Affecting Restriction is compatible with the Prior Existence View. The Prior Existence View implies that the welfare of contingent beings does not count in the aggregation of welfare, and therefore killing a being that could have had a pleasant future is a welfare loss, even if a new being would be put in the killed being's place. Hence, the Wide Person-Affecting Restriction in conjunction with the Prior Existence View on whose welfare to count would rule out replaceability. What is more, the link between the (Wide) Person-Affecting Restriction and the Prior Existence View is quite strong. As has been indicated in chapter 5.4, the (Wide) Person-Affecting Restriction together with the view that causing a being to exist cannot harm or benefit that being, implies the Prior Existence View. The Prior Existence View, in turn, rules out replaceability. That means that the Replaceability Argument is not compatible with the Wide Person-Affecting Restriction.

Not only is the Wide Person-Affecting Restriction a possible view that utilitarians can adopt in order to avoid the Non-Identity Problem. The Wide Person-Affecting Restriction actually has advantages above the Narrow Person-Affecting Restriction. I will point out some of those advantages in the following section.

5. Advantages of the Wide as opposed to the Narrow Person-Affecting Restriction

The Wide Person-Affecting Restriction has the advantage as opposed to the Narrow Person-Affecting Restriction that it can avoid the Non-

Identity Problem. I will now point out some differences of both views in particular cases. Subsequently, I will discuss the relevance of those differences for utilitarianism.

Consider, as a first example, two possible outcomes of our choices, outcome A and outcome B, represented as two possible future states of the world. Outcome A consists of population OK with people whose welfare is only slightly above the neutral level. Outcome B consists of population GREAT that is twice as large and consists of very well off people. The outcomes can be illustrated as follows, where the height of the graphs represents the welfare level of the population and the width of the graph represents the number of people. Bringing about each outcome would be equally right, according to the Narrow Person-Affecting Restriction, if both populations consisted of different people.

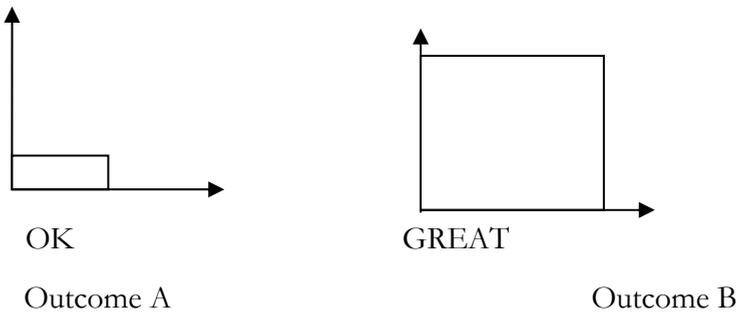


Figure 3: OK versus Great

After all, bringing about B rather than A would be better for no one. The Wide Person-Affecting Restriction would prefer bringing about population B, even if both outcomes contained different people and no single being would exist in both outcomes. The reason is that outcome B

would be better for the B-people than outcome A would be for the A-people.²²

The second example is a choice situation in which we have the same options A and B, but one person, Fred, would exist in both outcomes. In outcome B, Fred would be as well off as the other very well off people in B. In outcome A, Fred would be even happier than that and much happier than the OK people. As A would be better for Fred and worse for no particular person, the Narrow Person-Affecting Principle would require to bring about outcome A. The Wide Person-Affecting Restriction would not have this implication.²³ It would hold that B is better for the B-people than A is for the A-people. Therefore, the Wide Person-Affecting Restriction would evaluate outcome B as better.

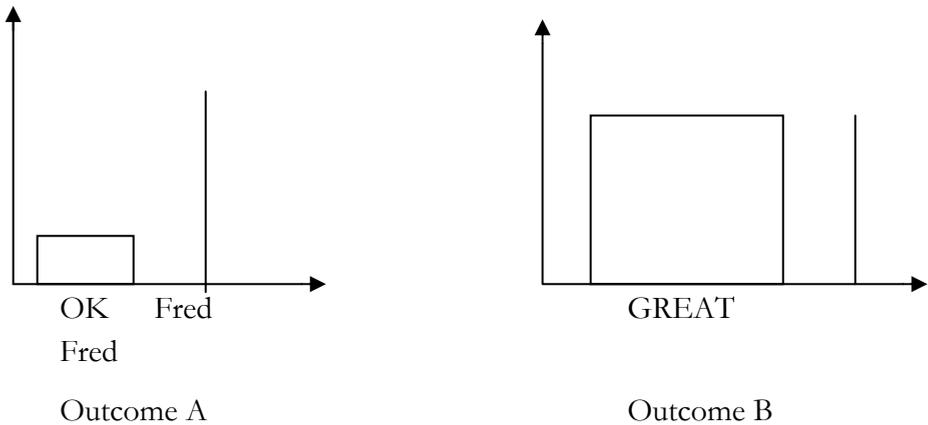


Figure 4: Two populations and Fred

²² Note that in this example, not only the genetic identities of the beings in both populations differ, as well as their welfare level. The number of beings that exist in each population differs as well. In Appendix A I explain how the Wide Person-Affecting Restriction can deal with different number choices. In order to avoid that complication one could just as well imagine this example (and the next) with populations of equal size.

²³ Norcross (1999) has brought forward those examples as reasons to dismiss the (Narrow) Person-Affecting View. Norcross has not considered the Wide Person-Affecting View.

So far, I have only described the different verdicts about those cases that the Wide and Narrow Person-Affecting Restriction would yield. Now, I want to suggest that the implications of the Wide Person-Affecting Restriction are more in line with utilitarianism. Consider again, the second above-mentioned example, featuring Fred. It seems that Fred is in a privileged position here, according to the Narrow Person-Affecting Restriction. After all, Fred's interests alone determine how to choose between both populations. Fred's interests alone determine whether the population that he will be part of will consist of people that are barely well off or very well off. Such a privileged position for one individual seems to be at odds with the utilitarian duty to neutrally maximize welfare.

This is no knockdown argument. After all, utilitarian proponents of the Narrow Person-Affecting Restriction could reply that the choice for outcome A above outcome B in the second above-mentioned example is in line with the duty to neutrally maximize welfare, interpreted as the duty to maximize net benefits to particular individuals. That is granted. Yet still this focus on maximizing net benefits of particular individuals leads us far astray from original utilitarian concerns with maximizing welfare and impartiality.

Consider now the first of the above-mentioned examples. A utilitarian proponent of the Narrow Person-Affecting Restriction would be indifferent as to which outcome to choose if different people would exist in both outcomes. After all, no particular individual would be benefited or harmed by either choice. Again, this indifference between those outcomes seems to be far removed from the original utilitarian concern with maximizing welfare.

Now, couldn't utilitarian proponents of the Impersonal View on the evaluation of outcomes bring forward similar concerns against the Wide Person-Affecting Restriction? Remember that the Impersonal View evaluates outcomes solely on the basis of the quantity of welfare

that they contain. If anything is straightforwardly in line with the utilitarian requirement to neutrally maximize welfare, it seems to be the Impersonal View. It is granted that the Impersonal View is indeed most straightforwardly focused on the maximization of welfare. However, one might wonder whether utilitarianism as a moral theory should be concerned with the maximization of welfare as an abstract quantity, even if it harms and benefits nobody. Or, alternatively, should utilitarianism be concerned with sentient beings and how they are affected in their welfare? One famous way of presenting this choice is whether utilitarianism should be concerned with making happy people or with making people happy. One infamous implication of the first, impersonal, interpretation of utilitarianism is known as the Repugnant Conclusion. It will be discussed in the following chapter.

For now, it should be clear that utilitarians could accept the Person-Affecting Restriction on the evaluation of outcomes. As compared to the Narrow Person-Affecting Restriction, the Wide Person-Affecting Restriction is not focused on harms and benefits for particular individuals, but rather on effects on sentient beings, whoever they are. In that respect, the Wide as compared to the Narrow Person-Affecting Restriction seems to be somewhat closer to the Impersonal View. However, the Wide Person-Affecting Restriction is still crucially different from the Impersonal View. It is not an intrinsic aspect view that evaluates outcomes only in terms of the quantity of welfare that they contain. The Wide Person-Affecting Restriction is concerned with aggregate net benefits. If the Person-Affecting Restriction is interpreted in a wide way, it is not vulnerable to the major challenge that has been brought forward against this view on the evaluation of outcomes, namely the Non-Identity Problem.

6. Conclusion

The Prior Existence View needs a particular assumption about what matters in the evaluation of outcomes in order to be a coherent utilitarian view about the question across whom to aggregate. It needs to assume that what matters in the evaluation of outcomes are harms and benefits to sentient beings. The Person-Affecting Restriction implies that an outcome can only be better (worse) if it is better (worse) for sentient beings. In its narrow interpretation, this means that the outcome must be better (worse) for any particular being. This Narrow Person-Affecting Restriction runs into the Non-Identity Problem. The Non-Identity Problem refers to the observation that in different people choices, the optimal outcome in terms of welfare cannot be said to be better for any particular being.

A wide definition of the Person-Affecting Restriction can avoid the Non-Identity Problem.²⁴ The Wide Person-Affecting Restriction claims that outcomes must be evaluated in terms of their aggregate net benefits. The Wide Person-Affecting Restriction implies that an outcome is better (worse) if and only if it is better (worse) for sentient beings, *whoever they are*. This way of circumventing the Non-Identity Problem is based on the idea that different definitions of identity are relevant in different contexts. The idea is that in non-identity cases, which concern moral responsibility, our notion of identity should capture what is relevant from a moral point of view. In those cases, we are not interested in individuating persons. Thus the genetic notion of identity is not necessarily the relevant one. People might, instead, be identified in terms

²⁴ Thus, the Wide Person-Affecting Restriction can deal with different people choices. As mentioned, in Appendix A I explore how the Wide Person-Affecting Restriction can deal with cases in which not only different people live in different outcomes, but also different numbers of people. This is interesting with respect to the Wide Person-Affecting Restriction as such, but not directly relevant for the issue of this thesis, i.e. the exploration of the implications and assumptions of Prior Existence Utilitarianism and Total Utilitarianism.

that are morally relevant, such as ‘our next child’ (whoever it will be) or ‘those who will be living after us’ (whoever they will be). This, however, builds on a notion of morality and on a position concerning what is morally relevant. It seems that as far as utilitarianism is concerned, the Wide Person-Affecting Restriction is a coherent view about the evaluation of outcomes. The Wide Person-Affecting Restriction differs crucially from the Narrow Person-Affecting Restriction, which is focused on effects on particular individuals. It differs also from the Impersonal View, which is focused on the quantity of welfare as such, rather than on effects on sentient beings.

So far, both assumptions that the Prior Existence View needs in order to be a coherent utilitarian view about which entities to count in the aggregation of welfare have been defended. I have argued in the previous chapter that causing a being to exist cannot harm or benefit that being. In this chapter I have argued that the view that outcomes should be evaluated in terms of their effects on sentient beings is a coherent utilitarian view on the evaluation of outcomes. Therefore, we can conclude that the Prior Existence View is a coherent utilitarian view: If indeed outcomes should be evaluated in terms of harms and benefits for sentient beings, and if bringing a being into existence cannot harm or benefit that being, then it makes sense not to take into account the possible welfare of contingent beings in the evaluation of outcomes. This is what Prior Existence Utilitarianism implies.

In the following two chapters I will explore the implications of both Total Utilitarianism and Prior Existence Utilitarianism.

8. REPUGNANT CONCLUSION AND EXPECTED MISERY ARGUMENT

1. Introduction

As we have seen, both the Prior Existence View and the Total View are possible utilitarian views on the moral status of contingent beings. Total Utilitarianism implies the Replaceability Argument and is therefore compatible with the aim of animal-friendly animal husbandry, while Prior Existence Utilitarianism does not imply the Replaceability Argument and is not compatible with the aim of animal-friendly animal husbandry. In the foregoing three chapters I have explored the assumptions of both views. In the remainder of this thesis I will explore some implications of both versions of utilitarianism.¹

¹ One implication, concerning transitivity, will be briefly discussed in Appendix C.

In the literature, implications of both versions of utilitarianism are compared in terms of their fit with commonsense judgments or what can be called ‘intuitive appeal’. This raises the question about the role of intuitions within utilitarianism (section 2). The most broadly criticized implication of Total Utilitarianism is that it might require bringing into existence additional people in order to raise the overall quantity of welfare. This might hold, even if adding people to the population significantly reduces the welfare of already existing beings, and even if the added people themselves are barely happy. If the number of people whose lives contain welfare is large enough, their lives together might contain more welfare than the lives of a smaller number of very happy people. This implication has been called the Repugnant Conclusion (section 3). Prior Existence Utilitarianism can avoid the Repugnant Conclusion, because it does not take into account the possible welfare of contingent beings. However, exactly the omission of taking into account the possible welfare of contingent beings seems to lead to an implication that many consider no less repugnant. Prior Existence Utilitarianism seems to be unable to account for the intuition that bringing into existence a being that would surely have a miserable life is morally blameworthy. This implication can be called the Expected Misery Argument (section 4). Contrary to what has been claimed, I will argue that Prior Existence Utilitarianism can deal with the Expected Misery Argument and account for the intuition that wantonly bringing into existence a miserable being is morally blameworthy (section 5). Thus, Total Utilitarianism implies the Repugnant Conclusion, while Prior Existence Utilitarianism might not be as vulnerable as has been suggested to the Expected Misery Argument (section 6).

2. The role of intuitions within utilitarianism

We have seen in previous chapters, and we will also see in this chapter, that counter-intuitive implications are brought forward in the evaluation

of Total Utilitarianism and Prior Existence Utilitarianism. For instance, I have discussed the scope of the Replaceability Argument in chapter 4. The idea that normal human persons would in principle be replaceable has been found intuitively objectionable. We have seen that Singer takes much effort to avoid this implication. In order to reduce the scope of the Replaceability Argument, Singer proposes the moral ledger model. The moral ledger model implies that welfare can only be negative or neutral and that all of us who have some unsatisfied desires when we die had better not lived. Singer, again, takes efforts to avoid this implication. It seems that even Singer, who has been very critical about the role of intuitions in the evaluation of moral theories, is concerned with the intuitive appeal of the theory that he defends.²

In this chapter, I will explore implications of the Total View and the Prior Existence View that have played a prominent role in the evaluation of both views. The reason why those implications have received so much attention is that they evoke a feeling in many people that something is wrong about those implications. The implications do not fit with many people's judgments about what is right or wrong. Those judgments are spontaneous judgments and they are broadly shared. Intuitions are "things that strike us as true without us knowing entirely why they do".³ An intuition is usually experienced as a strong and persistent feeling that something is right or wrong. So, the intuitions I am talking about are tokens or types of psychological phenomena that are often triggered in the evaluation of moral theories. The idea is that intuitions are about something and they can be true or false: they are considered truth-valued propositional attitudes.⁴ Given the prominent role that those judgments play in the debate about the Total View and

² For an illuminating discussion of Singer's position concerning intuitions see Singer (2005) and Sandberg & Juth (2010).

³ Cohnitz and Häggqvist (2009): 3.

⁴ Cohnitz and Häggqvist (2009).

the Prior Existence View, it is interesting to inquire whether and in how far it matters that a particular moral view has counter-intuitive implications. In particular, we are interested here in the utilitarian stance on that issue.

So, what is the status of those attitudes, marked by “directness, (apparent) non-inferentiality, (subjective) compellingness or ‘glow’”?⁵ In how far should intuitions be taken seriously in the evaluation of moral theories? It is disputed what the status of such intuitions should be. At first glance, there seem to be two ‘extreme’ views, which claim either that moral intuitions are completely irrelevant, or that moral intuitions are indisputable data that a moral theory should accommodate. I am not sure whether anybody even pretends to adhere to any of those extreme positions, but it will be useful to distinguish them for analytical purposes. Besides those extremes, there is a broad spectrum of possible views in between. That a view is ‘extreme’ in that sense, does not establish that it is wrong; nor is the number of adherents to a view an indicator of its truth. So, what else can be said about that issue?

It is a possible position to deny the relevance of counter-intuitive implications in the evaluation of a moral theory. According to that position, it is irrelevant whether a moral theory fits well with existing intuitions. After all, a moral theory is not intended to be the most comprehensive summary of our existing moral judgments. Instead, a moral theory should be able to critically assess common moral practices and judgments. Moral theory should provide a critical point of view vis-à-vis our convictions. A moral theory should offer a rational justification of moral judgments, and unmask practices and judgments that are morally wrong. A moral theory should provide an answer to the question “How should I act?” For a utilitarian, a reasonable answer to this question can be: “Ignore all our ordinary moral judgments, and do what

⁵ Cohnitz and Häggqvist (2009): 4.

will produce the best consequences.”⁶ If there is a convincing story of why and how a moral theory arrives at its judgments, a need to match with intuitive judgments can be deemed unnecessary. According to that view, if some of our beliefs are not in line with what the theory requires, that is so much the worse for our beliefs. It shows that we are on the wrong track, and we ought to adjust our beliefs, or at least, not let them influence our actions.

Evidence suggests that intuitive moral judgments are influenced by many factors. Evolutionary contingencies, cultural upbringing, subconscious cues, and situational circumstances have all been found to influence people’s spontaneous moral judgments. For instance, in nearly all of the 200.000 years of human history, hurting others was only possible in a direct, face-to-face way. It might have been adaptive for social beings that depend on each other, not to hurt each other, at least within one’s own group. This might be the reason why we tend to feel that it is wrong to hurt others in face-to-face situations. Hurting others in more indirect ways across distances of time and space is only possible relatively recently, evolutionary speaking. It has been shown that hurting others across distances of space and time, without having to face them, so to say, does not evoke the same deeply ingrained negative feelings and judgments. This evolutionary explanation can explain that we tend to arrive at different moral judgments about different forms of violence. It can explain why we tend to feel uncomfortable about violence in face-to-face situations and are less disturbed with violence in impersonal situations. If the difference were claimed to be morally relevant, an additional explanation would be needed. It has also been shown in experiments that invoking disgust by, for instance, presenting a case in a

⁶ Singer (2005): 346.

disgusting room affects people's moral judgment on a presented case. The feeling of disgust seems to play a role in negative moral judgments.⁷

The fact that certain judgments are persistent and widespread does not necessarily show that those judgments are morally right. If indeed evolutionary contingencies or the cleanliness of a room can influence intuitive moral judgments, it needs at least a further argument to conclude that those intuitive judgments are morally relevant and even morally right. The above-mentioned findings have evoked debates about the trustworthiness of spontaneous judgments as indicators of what is morally right or wrong. Those who do not wish to bother very much about a moral theory's fit with commonsense judgments have used those findings about possible explanations for moral intuitions in order to support their cause. Their idea is that if intuitive judgments are influenced by such contingent factors, it can be doubted that they are reliable guides to what one morally ought to do.⁸ Let me only note here that an assessment of that criticism, and indeed of any claim about what the role of intuitive judgments in moral theorizing should be, depends on what one considers the purpose of moral theorizing to be. It depends on one's understanding of morality. With other words: "Whether intuitions are evidence depends, of course, on what they are supposed to be evidence *for*."⁹

Instead of denying the importance of a fit between moral theories and widespread moral intuitions, another possible position is that intuitions should be taken very serious. It seems that the more commonsense judgments are taken to be indisputable and the more a moral theory is modeled in order to fit with them, the less critical

⁷ Schnall et al (2008).

⁸ Haidt et al (1997), Greene (2003), Singer (2005).

⁹ Cohnitz and Häggqvist (2009: 9) distinguish "three different conceptions of what philosophical analysis is concerned with that correspond to three different evaluations of whether intuitions should be considered evidence in philosophy".

potential the resulting theory has. Those who take moral intuitions in the sense of spontaneous moral judgments about particular cases seriously, are not to be confused with moral intuitionists who accept moral intuitionism as a meta-ethical position about the justification of moral claims. Moral intuitionists hold that certain moral judgments cannot be based on more fundamental judgments. Those moral judgments are called ‘intuitions’. However, the intuitions that the meta-ethical intuitionists talk about are not spontaneous, unreflected, emotional responses.¹⁰ So, intuitionists are not representatives of the view that intuitions, in the sense of spontaneous beliefs, should be taken at face value. Rather, it must be people who are prepared to take counter-intuitive implications seriously as indicators that a moral theory might be wrong.

While a moral theory should retain a critical potential to unmask and correct prejudices and false judgments, even if they are persistent and broadly shared, it seems also hard to completely neglect intuitive judgments about what is right and wrong in the evaluation of a moral theory. If counter-intuitive implications are not considered beforehand either knockdown arguments against a moral theory, or completely irrelevant, what can be an appropriate reaction to appeals to counter-intuitive implications? Those who accept a certain theory that has a counter-intuitive implication can bring forward reasons for accepting the theory in spite of its counter-intuitive implications. Maybe the theory is built on some highly attractive basis that one is not prepared to give up. Maybe it can be shown that rival theories score even worse in terms of intuitive appeal. Alternatively, one can attempt to discredit the intuition and argue that it has to be abandoned as morally irrelevant or prejudiced after further reflection.

¹⁰ Roeser (2002): 11.

Traditionally, the fiercest opponents to the idea that counter-intuitive implications are morally relevant can be found among utilitarians. Utilitarianism has a history of being used to challenge what counts as accepted wisdom. Utilitarianism has been brought forward against broadly accepted ideas and religious prescriptions that certain sorts of actions are plainly wrong or that some sentient beings count for less than others. As we have seen, according to utilitarianism, no class of action is wrong as such; it always depends on the consequences. Furthermore, each counts for one, and no one for more than one. Utilitarians have always defended positions that “challenge prevailing orthodoxy”.¹¹ Utilitarians have always criticized many broadly shared beliefs as untenable prejudices that ought to be abandoned. Utilitarians have defended moral requirements and implications that indeed many people have found repugnant.

The founding father of utilitarianism, Jeremy Bentham (1748) and his famous scholar John Steward Mill (1806) were influential social reformers. Bentham was a political radical. He defended the separation of church and state, equal rights for women, the right to divorce, and the decriminalization of homosexual acts. He opposed slavery and he spelled out the basis of a utilitarian view on animal rights. Mill was arguing for women’s rights, such as their right to vote. Mill was aware that he attacked “an almost universal opinion” at his time.¹²

More recently, utilitarians have become infamous for their denial of the moral difference between active and passive euthanasia, the claim that we ought to give significant amounts of our possessions to needy strangers, and the argument that the alleviation of a big enough number of headaches can justify the killing of an innocent person.¹³ One might get the impression that utilitarians are not impressed by counter-intuitive

¹¹ Rachels (1999).

¹² John Steward Mill, *The Subjection of Women*.

¹³ See Rachels (1999), Singer (2009), and Norcross (1997).

implications of their moral theory, or of any moral theory, for that matter. As reformers, utilitarians can expect that they have to argue against the prevailing opinion.

Although utilitarians do not shy away from making claims that strike many as utterly implausible, they tend to be prepared to defend those claims. In those defenses appeals to counter-intuitive implications are taken seriously, in one of the many above-mentioned ways. For instance, intuitive judgments are unmasked as being caused by evolutionary contingencies that are morally irrelevant;¹⁴ attempts to avoid certain counter-intuitive implications are shown to be unconvincing;¹⁵ and judgments that strike many as counter-intuitive are presented, from a different perspective, as normal and reasonable, after all.¹⁶ Some leading utilitarians even defend a coherentist method of arguing what we have reason to believe. According to that method, “in defending a moral theory, we must see how well that theory fits in with a wide variety of judgments that we are inclined to make about many different matters.”¹⁷ Also in that process, judgments are submitted to critical (re-) consideration and are confronted with other judgments about cases, moral theory or the purpose of morality.

As those reactions reveal, even utilitarians take appeals to counter-intuitive implications seriously as an invitation for further discussion. Therefore, it should bear no surprise that appeals to counter-intuitive implications figure also in the debate between Total Utilitarianism and Prior Existence Utilitarianism. My aim in what follows is not to evaluate the implications of both views in terms of their intuitive appeal or plausibility. I will leave this to the reader. After all, if

¹⁴ Singer (2005): 348.

¹⁵ Norcross (2008).

¹⁶ Norcross (1997): 158-167. Sicora, Sumner, Singer (2009).

¹⁷ Kagan (1998): 15. See also Tännsö (1998): 11, and Hooker (1996): 531.

intuitions are taken seriously in moral evaluations, this holds much more for one's own intuitions than for those of others. Furthermore, for my purpose, it is not necessary to take a position on whether and how intuitions should be taken seriously, and which view is intuitively more appealing. My aim, after all, is not to defend any view, but rather to explore the implications and assumptions of two coherent utilitarian views on whose welfare to count in the aggregation of welfare. Hence, in what follows, I will point out what the implications of Total Utilitarianism and Prior Existence Utilitarianism are.

3. The Repugnant Conclusion

In the aggregation of welfare, the Total View takes into account the welfare of all existing beings and the possible welfare of all possible beings. This seems to be attractively straightforward: all welfare consequences of the options are considered. The aim is neutral welfare maximization. The Total View accepts two ways of maximizing welfare as equally valid: One can either make existing beings happier or bring into existence more beings that are happy. This implies that moral agents would be required to add more people to the population, if this was the best way of maximizing overall welfare. Even stronger, extra people must be added, even if this decreases the welfare of the already existing people, as long as the total sum of welfare after the addition is greater than before. If a lot of people are added, the total amount of welfare can increase, even if the average welfare per person becomes very low. Consider the following simplified example: Two persons, A and B, each have a welfare level of 10. So, the total would be 20. If a person would be added to the scene with a welfare level of 3 and if this addition would reduce the welfare of A and B from 10 to 9, this would result in a total utility of 21. The Total View would therefore require the addition if no better way of maximizing welfare were available. This holds if the Total

View on who counts morally is accepted in conjunction with the total view as a method of aggregation.¹⁸

Parfit has famously stretched this implication of what he calls the 'Impersonal Total Principle' further. As Parfit exemplifies: "The greatest mass of milk might be found in a heap of bottles, each containing only a single drop."¹⁹ Back to the population case, this would lead to what Parfit called the Repugnant Conclusion:

The Repugnant Conclusion: 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.²⁰

So, Parfit describes a scenario in which adding people to the population maximizes welfare. If the welfare of the added people *together* is sufficiently high, it can compensate even huge losses in the welfare of

¹⁸ The Total View on who counts morally in conjunction with the *average view* as a method of aggregation would not yield this result. It implies that any addition to a population is forbidden if it lowers the average. Thus any child with a level of welfare beneath average, even if very happy, should not be born. Killing individuals with a welfare level below average does also raise the average. And where should the killing end? There will always be individuals beneath average, until only one individual or some equally happy individuals remain. Furthermore, the average view would require the addition of a very unhappy individual to a population, if the rest of the existing individuals fare even worse. (This would benefit nobody, of course, but it would raise the average welfare level.) For those reasons the average view has been dismissed. (This shows, again, that even utilitarians care about counter-intuitive implications of moral theories.)

¹⁹ Parfit (1984): 388.

²⁰ Parfit (1984): 388.

the previously existing people. The number of 'at least ten billion' is arbitrary.²¹ The welfare level of the population that Parfit depicts is slightly above zero. Parfit calls this 'barely worth living', assuming that the zero-level is a point where a life turns from being worth living to not being worth living. I will leave this interpretation for what it is at this point, and prefer to stick to describing the welfare level as 'slightly above zero', without making assumptions about when a life is worth living. Anyway, Parfit finds this implication of the Total View counter-intuitive and calls it the Repugnant Conclusion.

So, the Repugnant Conclusion has been presented as a problem for the Total View. Let us see whether and in how far the Total View can deal with this alleged problem. Can we somehow explain that it is not unpalatable? Must we accept it, fully aware of its unintended consequences, by lack of any convincing alternative? Can we embrace it as far as value theory is concerned, but try to keep it out of the deontic domain? Should we focus on the deontic domain and try to find alternatives for the Impersonal View? In what follows, I will elaborate a bit on those possible reactions to the Repugnant Conclusion, in order to get a more vivid understanding of what it entails. My question here is not whether any intuitive judgment about this implication of the Total View should be endorsed. My concern is rather to figure out what the implication that has been described as the Repugnant Conclusion actually is.

Several arguments have been brought forward for showing that the Total View can avoid the Repugnant Conclusion, or that the Repugnant Conclusion is not as repugnant as it might seem at first glance. One of those arguments is directed against the criticism that the

²¹ I guess that Parfit assumes a big number, because if only a few people were assumed to live, adding extra people - even at the expense of the previously existing people's welfare- might intuitively be a less unattractive thing to do. Ideas about the diminishing value of extra lives in bigger populations might play a role here. Hurka (1983) has explicitly defended that idea in his 'variable value view'.

Total View requires the expansion of the population. Against this criticism it is claimed that an expansion of the population is only required if there is no more efficient way to maximize welfare. As long as there are equally good options of raising the total amount of welfare, the Total View does not require the addition of happy people. If other ways of raising overall welfare are more efficient, proponents of the Total View will favor those other ways.²²

Another argument of defenders of the Total View is directed against the criticism that the idea of raising overall happiness by making more happy people is repugnant, because there are more urgent ways of raising overall happiness. Raising overall happiness can also be achieved by making existing people happier. In the present world, there are many people that struggle with a very low quality of life. According to the critics of the Total View achieving a higher quality of life for those people is more important than raising overall welfare by making more happy people. The same holds for animals. Also against this criticism defenders of the Total View have brought forward that the Total View does not require (and does not even allow) creating more people if there are more efficient ways to maximize welfare.

Defenders of the Total View have argued that the Total View would not lead to the Repugnant Conclusion. Given the present situation of the world, it seems unlikely that adding more people would be the best means of realizing more happiness, or so they have argued. According to some defenders of the Total View, it seems much more likely to achieve more happiness by focusing our efforts on ameliorating the situation of countless people that are now struggling with a very low standard of living. Especially if the welfare of non-human animals is taken into account as well, an ongoing expansion of the human and animal population is not likely to lead to the most overall welfare. A

²² Sumner (1978): 104.

redistribution of resources among a stable society might be a more efficient way of increasing overall happiness. Whether this is right is an empirical question. If it were right, the Total View would require just this: a redistribution of resources, rather than (or in addition to) the bringing into existence of more happy beings. In that case, the Total View would not require that we strive for a much larger population with barely happy people. If the idea that we would be required to do just that is what makes the Repugnant Conclusion so repugnant, and if the Total View would not require this in the current situation, the Repugnant Conclusion would indeed be less repugnant than it seemed to be.

Defenders of the Total View have also claimed that even if the Repugnant Conclusion follows from the Total View, it is not as repugnant as it might seem. The Repugnant Conclusion might appear repugnant, because a large population consisting of people whose happiness level is barely above zero seems to be a very unattractive thing to strive for. Just in itself, a world populated with many people whose happiness level is barely above zero might be considered a bad thing. The underlying idea might be that that outcome cannot be good, because everyone's welfare level is deplorably low. Defenders of the Total View have pointed out, however, that the population in question consists of people whose lifetime welfare is positive. They might not have much pleasure, but then they do not have much pain either. This does not seem to be an unacceptable sort of existence. What would probably be more realistic to assume is that those people have happy as well as unhappy episodes in their lives, but have lives that are more happy than unhappy.²³ Granted, 'barely above zero' means that their lives are surely not very happy.²⁴ But still, their welfare is positive. Thus, according to some defenders of the Total View, the implication called the Repugnant

²³ Sikora (1978): 117.

²⁴ Those with a more demanding definition of happiness would not consider them 'happy' at all (Haybron, 2008: 4).

Conclusion is not very repugnant after all.

Another idea of what is repugnant about the Repugnant Conclusion has been that people should be autonomous in their reproductive decisions. People should not be morally required to bring happy children into the world, nor to devote their energy to other projects that maximize happiness more efficiently. The underlying idea is, then, that in some domains autonomy is more important than results in terms of welfare. Defenders of the Total View have pointed out that this idea about the value of autonomy is hard to combine with a utilitarian moral theory. From a utilitarian perspective, maximizing welfare is by definition the sole aim of morality. Respect of autonomy as such is not allowed to outweigh welfare considerations. Autonomy might be included in the concept of welfare. In that case, however, it is no separate consideration.

What many find repugnant about the Repugnant Conclusion is its implication about “the comparative value of a population with very high welfare and a population with very low welfare.”²⁵ It is its implication that a population with very low welfare could somehow be better in terms of welfare than a population with very high welfare. This sting of the Repugnant Conclusion has not been taken away by the above-mentioned considerations. What remains counter-intuitive for many is the idea that a perfectly happy population might be required to give up significant amounts of welfare, in order to allow for the addition of extra persons, resulting in a larger population of people that are barely happy. This intuition has also been spelled out on a smaller scale. Should a couple, which expects to be happier without children, be required to have children, provided that the happiness of those children would more than compensate for the lost happiness of the couple? This is what Total Utilitarianism requires - at least if no other way of raising total welfare is

²⁵ Arrhenius (forthcoming).

available.

On a deontic level the Repugnant Conclusion has been considered more repugnant than on a purely axiological level.²⁶ If we compare two populations as such in terms of welfare, the idea that the population with the greatest overall welfare is the better one in terms of welfare might be unavoidable. It says something about the value of states of affairs. It does not mean that people in the first world have the duty to bring more beings into existence and strive for the second world. The repugnance of the Repugnant Conclusion might lie exactly in the implication that maximizing happiness by making people that are happy should count on a par with maximizing happiness by making people happy. The idea behind that criticism of the Total View would be that existing beings (and future beings) should have priority over contingent beings. That is exactly what the Prior Existence View claims. So, the Prior Existence View can be seen as an expression of the idea that maximizing happiness by making people that are happy should *not* count on a par with maximizing happiness by making people happy.

4. The Expected Misery Argument

In the literature, the Prior Existence View is typically presented as an “initially appealing” theory.²⁷ The Prior Existence View does not take into account the possible happiness of contingent beings in decisions about whether to create them. Therefore, that view can avoid the Repugnant Conclusion on the deontic level. The Prior Existence View can accept that only the expected effects on the welfare of the (possible) parents - and on other already existing persons - determine whether or not to bring a potential child into existence.

²⁶ Broome (2004): 140.

²⁷ Hurka (2001): 587.

Yet, exactly this alleged advantage has an implication that has been considered problematic. The Prior Existence View does not take into account the possible welfare of contingent beings. That a contingent child would be happy is no reason for having it. But then, is the knowledge that a contingent child would have a miserable life not a moral reason *against* having it either? The point is that the Prior Existence View cannot say this without inconsistency. This problem has famously been evoked by Parfit's Case of the Wretched Child. Parfit describes this case as follows:

The Wretched Child: Some woman knows that, if she has a child, he will be so multiply diseased that his life will be worse than nothing. He will never develop, will live for only a few years, and will suffer pain that cannot be wholly relieved.²⁸

The idea is that in this case, we would surely want to take into account the expected welfare of the child, which would be negative, as a reason against causing this child to exist. According to the Prior Existence View, only actually existing beings should be taken into account in moral considerations, along with beings that *will definitely exist*. Now, the problem is that once this child exists, the choice not to cause it to exist is not open anymore. And when the couple chooses beforehand not to cause this child to exist, the potential child is not a child that will definitely exist. So, the Prior Existence View lacks any basis on which the potential child's welfare can be taken into account.

If a couple knows before conception that any child that it could have would (due to a genetic deficiency, say) have a short and very miserable life, then supposedly most people are inclined to think that the

²⁸ Parfit (1984): 391.

couple is blameworthy if it has this child. This is brought forward as an attack against the Prior Existence View. It cannot account for the idea that there is something morally wrong with having this child. Maybe not everybody shares this intuition. At least it is playing a prominent role in the evaluation of the Prior Existence View. Therefore, I will address it here.²⁹

Either adherents of the Prior Existence View cannot accommodate the idea that the expected misery for a child should count as a reason against bringing it into existence. Or, they need to justify that the misery of a potential miserable child counts in decisions about whether or not to bring it into existence, while the happiness of a potential happy child does not count. This challenge is referred to as the Asymmetry.³⁰ The Asymmetry refers to the discrepancy between on the one hand accepting the potential suffering of a potential child as a reason for *not* having it and on the other hand *not* accepting the potential happiness of a potential child as a reason *for* having it. With other words: the future misery of a potential miserable child counts as a reason for not having this child, while the future happiness of a potential happy child does not count as a reason for having this child. Taking it into account in one case (the miserable child) and not taking it into account in the other case (the happy child) implies an asymmetry.

It should be noted, though, that the Asymmetry does only arise if one horn of the dilemma is gripped. It arises only if proponents of the Prior Existence View want to count the misery of the contingent miserable child as a reason against having it. The other horn of the dilemma would be not to take the misery into account. Then, there would be no asymmetry. However, this would go against the idea that it is morally blameworthy to bring into existence the miserable child.

²⁹ Singer (1993): 123, Singer (forthcoming 2011): 109, Arrhenius (2000): 161, Mulgan (2006): 5.

³⁰ Anglin (1978): 36; Parfit (1984): 391.

5. Can Prior Existence Utilitarianism account for the idea that there is something morally wrong about having the miserable child?

In this section I argue that contrary to what has been claimed, Prior Existence Utilitarianism can account for the intuition that bringing into existence the miserable child is morally blameworthy.

First of all, we need to specify what seems to be wrong with the case. The idea might be, more precisely, that parents that would like to have such a child are bad persons or people with a bad character. Indeed the parents seem to be very egoistic, thinking only about themselves and not about their child. Maybe they are keen on having such a child in order to care for it, but then still this would be very egoistic. A related idea might be that such parents would never be good parents and should not have a child at all. Alternatively, the project of having the child as such might be considered the reason why there is something morally wrong with having the child.

Considering these possible specifications of the idea that there is something morally wrong with the parents wanting to have the miserable child, I will argue that Prior Existence Utilitarianism can at least account for the first set of justifications, namely those referring to the character of the parents. After all, Prior Existence Utilitarianism, just as other forms of utilitarianism, can accept character as an evaluative focal point. After having made that point, I will explore whether and in how far Prior Existence Utilitarianism can also account for the idea that the project of having the child as such is morally wrong.

To begin with, it should be noted that the decision of not having the child can usually be justified even without taking the negative welfare of the possible child into account. After all, the welfare of others, such as the parents, siblings of the ‘society in general’ would be affected

negatively by having the child. This shows that in practice, the trouble would be restricted to cases in which the parents would be happy to have the child and others would not be negatively affected to a degree that outweighs the parents' welfare gain. (Except, of course, the child itself. But that does not count.) The problem is restricted to cases in which the parents want to conceive the wretched child. The question is how the Prior Existence View can deal with such a case.

As a second preliminary remark, it should be noted that the Prior Existence View would have no problems with taking into account the suffering of the child, as soon as the child is actually suffering. By that time the child is an actual sentient being, and as such it will get full moral consideration according to Prior Existence Utilitarianism. Unfortunately, the only thing that could be done at that moment in order to prevent further suffering of the miserable child would be to kill the child as soon as it starts to be sentient. Depending on when sentience and in this case the suffering starts, this might be before or after the child is born. Even if the child would not be allowed to suffer, the commonsense idea might still be that there is something morally wrong about having this child. But then, what does "having the child" mean? Utilitarians can draw a distinction between "three stages in the process of bringing a miserable child into the world that are importantly different in their moral value:

- (i) intending to conceive such a child,
- (ii) conceiving such a child, and
- (iii) keeping such a child alive."³¹

Utilitarians have reasons to evaluate those stages separately.

³¹ Sapontzis (1987): 191.

(i) is a bizarre project and expresses bad moral character if the prospective parents intend to keep the child alive. If, however, they intend to abort the potentially miserable child before it develops sufficiently to suffer, than their intention is to satisfy their bizarre compulsion to conceive in a way that does no direct harm from any utilitarian viewpoint.³²

Stage (ii) is in itself morally neutral, because a fertilized egg is not a sentient being and therefore as such it does not count as a moral object according to utilitarianism. There might be indirect reasons for or against conceiving such a child. However, as soon as there is a child that is able to feel anything (stage iii), this child is an actual moral object, an existent sentient being, and thus counted by the Prior Existence View. The point of this distinction is to isolate what seems to be the most disturbing aspect of the case: the child should not be allowed to suffer. The Prior Existence View does not allow the suffering of the child.

The challenging question is whether the Prior Existence View can accommodate the idea that there is something morally wrong with having the child, apart from possible negative effects on third parties, and even before the child actually exists and suffers. Before the child exists as a sentient being no harm is done. What can the Prior Existence View say about conceiving the child or even the intention to do so? If the parents' intention is to keep the child, the parents can be accused of having a bad character:

But if such people did exist, they would have a perverted idea of reproducing and parenting and would show, by keeping the child alive for two miserable years, their willingness to use others merely as means to their own satisfaction. [...] Prior existence utilitarianism can account for this intuition, since this is an evaluation of character, and prior

³² Sapontzis (1987): 192.

existence utilitarianism no more precludes making character evaluations than does the total population view.³³

Let me explain how utilitarianism can take character into account. Utilitarianism's 'theory of the right' requires maximizing welfare. Now, the question rises how, more specifically, we should go about in maximizing welfare. What should be our focus? Should we always choose the *act* that maximizes welfare? Should we follow the *rules* that generally tend to maximize welfare? Should we adopt *motives* that are likely to maximize welfare? Should we do all of these? In order to get a clear idea of what utilitarianism requires, we must be clear about the 'evaluative focal points', the focus of moral evaluation.³⁴ As explained in chapter 2, section 4.3, direct utilitarianism can be distinguished from indirect utilitarianism:

Direct utilitarianism applies the principle of rightness directly to all focal points. There are many possible forms of direct utilitarianism, depending on what is taken as legitimate evaluative focal point or focal points. There are direct utilitarian theories, which accept only one evaluative focal point, for instance acts. There are also theories, which accept several direct evaluative focal points.

Indirect utilitarianism applies the theory of the right (e.g. "maximize welfare") only to one central evaluative focal point. Other evaluative focal points are evaluated only indirectly in terms of their relation to the central one. Rule utilitarianism is an example of an indirect utilitarian theory. Rules are evaluated directly in terms of the goal of maximizing welfare. ("Choose the rule that maximizes welfare.") Acts are judged right whenever they conform to the right rules.

It is possible for utilitarians to include character as an evaluative focal point (among other evaluative focal points), either directly or

³³ Sapontzis (1987): 191.

³⁴ Kagan (2000): 134.

indirectly. 'Character' has been defined as "the complex of mental and ethical traits marking a person".³⁵ 'Character', therefore, can be understood as "the set of qualities that make somebody distinctive, especially somebody's qualities of mind and feeling".³⁶ Such an understanding of 'character' is implied in a claim, such as: "It is not my character to behave that way".³⁷ Thus, character is linked to behaviour, in particular also moral behaviour. 'Disposition' is a synonym for character, meaning "an inclination or tendency to act in a particular way".³⁸ Character is considered as something rather constant and predictable, but also improvable. Aristotle has famously claimed that moral instruction should aim at inculcating good character.³⁹ Contemporary thinkers have followed up on that idea, notably Benett, who claims: "The central task of education is virtue".⁴⁰ On the other hand, critics have claimed that character does not really exist; instead, people's actions are typically influenced by other, more contingent factors.⁴¹

If character were accepted as a direct evaluative focal point, then utilitarianism would ask people to develop a kind of character that is most likely to maximize welfare. Now, one might wonder whether and in how far it is possible to influence one's character, or that of others, as might be intended in moral education. Now, what kind of character would generally help to maximize welfare? Would this be a character that is insensitive to the needs and suffering of others? Would this be a character that sanctions the use of others as mere means for the own satisfaction? This is ultimately an empirical question. Yet I am confident

³⁵ Merriam Websters Online Dictionary

³⁶ Encarta.msn.com/dictionary

³⁷ Encarta.msn.com/dictionary.

³⁸ Encarta.msn.com/dictionary.

³⁹ Aristotle (1984, e.g. 1103, 7-30)

⁴⁰ Benett (1993): 13-14. See also McDowell (1979): 333 and (1996), Sherman (1989): 157-99, Nussbaum (1999): 174, 187).

⁴¹ Doris (2002).

that such a quality of mind and feeling and such an inclination to act would generally not maximize welfare. Therefore, the intention of the couple to conceive and keep alive the wretched child can be condemned on utilitarian grounds as a sign of bad character. This can be done directly, if character is taken as a direct evaluative focal point, be it the only one or one among others.⁴²

As we have seen, utilitarianism, including Prior Existence Utilitarianism, can condemn the character of the couple that intends to have the miserable child. Hence Prior Existence Utilitarianism can justify the commonsense idea that there is something morally wrong about the parents having the child. Prior Existence Utilitarianism can point out the parents have a bad character in so far as they intent to have this child and let it suffer. In order to do this, one need not take into account the welfare of this particular contingent child, which is not possible according to the Prior Existence View. Rather, one needs to judge such a character on the basis of its general tendency to maximize welfare. The idea is that a character that is compatible with letting such a miserable child suffer is not a character that generally maximizes welfare. In that sense, Prior Existence Utilitarianism can subscribe to the common-sense judgments that those parents cannot be expected to be good parents and should not have children at all.

Before I proceed by exploring whether Prior Existence Utilitarianism can also support the idea that the project of having the child as such is morally wrong, I will say a bit more about the above-mentioned character evaluation. Evelien Pluhar has denied that Prior Existence Utilitarianism can make such a character evaluation. I will argue that this denial is mistaken.

First, Pluhar claims that a couple that would deliberately conceive the Wretched Child believing that ‘God makes no mistakes’ could not be

⁴² For a reflection on character utilitarianism, see Railton (1988).

accused of bad character.⁴³ Apparently, Pluhar thinks that the couple could not be accused of having bad character in this case, because it does not have evil intentions. However, this argumentation is mistaken. Utilitarianism does not determine the goodness of somebody's character on the basis of his or her intentions. A character is evaluated in terms of its general tendency to maximize welfare. A character that is compatible with having the miserable child is not likely to generally lead to the best consequences in terms of neutral welfare maximization.

Secondly, Pluhar claims that even if the couple acted from monstrous intentions, their character could not be dismissed *on utilitarian grounds*. According to Pluhar, utilitarians “can judge character traits to be good or bad solely by reference to their consequences.”⁴⁴ Pluhar specifies: “Specifically, one must ask how much harm (disutility) such traits are apt to cause.”⁴⁵ Pluhar claims that for a utilitarian, it is only the action that can be directly evaluated. Character can only be indirectly evaluated. Thus, the evaluation of the character follows the evaluation of the action. Now, according to Pluhar, Prior Existence Utilitarianism cannot condemn the action or project of having the miserable child before the child actually exists. Therefore she assumes that the corresponding character cannot be judged wrong.⁴⁶

However, utilitarianism has more possibilities of taking character into account than Pluhar assumes. Pluhar assumes a version of utilitarianism that accepts actions as a direct evaluative focal point, and character as an indirect one. In that case, actions would be directly evaluated in terms of the welfare they produce. We would be required to choose the action that maximizes welfare. Character would be evaluated indirectly, depending on the action. We would need to know whether an

⁴³ Pluhar (1995): 214.

⁴⁴ Pluhar (1995): 214.

⁴⁵ Pluhar (1995): 214.

⁴⁶ Pluhar (1995): 215.

action is right or wrong, and could then conclude whether the character of the agent was right or wrong. As Pluhar puts it: “*For a utilitarian, the bad characters of the parents do not “give the project a strong immoral value”: it is the project that makes their characters bad.*”⁴⁷ If that version of utilitarianism is assumed, then, in the case at hand, there would indeed be no way of condemning the character, if the action to which it leads is morally acceptable. That, however, would be a peculiar kind of utilitarianism, with actions as a primary and character as a secondary focal point. A person’s character could only be evaluated on the basis of the consequences of each particular action. That is not the way character evaluations are typically made. One could wonder what the use of making character evaluations in that way would be, if character would not be a more general category than a particular action.

Let me briefly investigate more promising possibilities. One alternative is to turn it around and argue for character as primary evaluative focal point and evaluate acts only indirectly. In that case, an act would be right if it “sprang from” the right disposition, i.e. if the agent had the right character. For instance, let us assume that the disposition of caring particularly for one’s own child is generally a good thing, i.e. generally maximizes welfare. Imagine a case in which a mother rescues her own child from a burning building rather than several others. According to the version of utilitarianism we are now considering, with character as a primary and actions as a secondary focal point, rescuing her own child rather than several others would be the morally right thing to do. However, I think that the focus on character, rather than acts would be suited as a decision procedure, if at all, rather than being an attractive account of rightness. As a decision procedure, i.e. a rule of thumb indicating how to go about maximizing welfare, focusing on character might recommend itself, because it avoids making complex welfare calculations before every action. Furthermore, focusing on

⁴⁷ Pluhar (1995): 214-215. Italics mine.

character allows uncoupling one's behavior from that of others, which might be advantages for solving collective action dilemmas.⁴⁸ However, as an account of rightness it seems to me incomplete. It is more plausible to choose character as one direct evaluative focal point, next to other direct evaluative focal points, such as acts.

Thus, yet another possibility would be to accept both character and actions as primary evaluative focal points. When several direct evaluative focal points are accepted, evaluative conflict might occur. This is the case when, for instance, a right disposition leads to a wrong action, as in the above-mentioned case of the mother rescuing her own child. A disposition that usually maximizes welfare and is therefore right leads to an action that fails to maximize welfare and is therefore wrong, on this version of utilitarianism. In extraordinary circumstances, acting from a right disposition can lead to a wrong action. The fact that the disposition leads to a wrong action in a particular case does not make the disposition wrong. After all, a disposition is judged on the basis of its *general* tendency to produce good or bad outcomes.

I do not think that the occurrence of evaluative conflict diminishes the plausibility of an account of what is morally right. Rather, proposals for taking away evaluative conflict have been shown to be problematic.⁴⁹ The acknowledgement of evaluative conflict does even make utilitarianism intuitively more plausible. In the case of the mother rescuing her own child, one could say that her act was morally wrong, because she failed to maximize welfare. However, she acted from a disposition, which is the right one to have, because it generally (or so we assumed) maximizes welfare. So, there is a real tragic conflict, which we experience in the case. Utilitarianism, accepting several focal points, can acknowledge this conflict and explain it. I take this to be an attractive

⁴⁸ See Jamieson (2005) on this issue.

⁴⁹ Louise (2006).

feature, rather than a disadvantage, for this account of rightness.⁵⁰

Summing up so far, Prior Existence Utilitarianism can account for the idea that it is morally blameworthy to bring into existence a child that would surely have a miserable life. Prior Existence Utilitarianism can morally condemn the character of the parents. Let me now explore whether and in how far it can also condemn the *project* of having the miserable child.

The problem for Prior Existence Utilitarianism with respect to condemning the project consists in the fact that the child's existence depends on the parents' decision about whether to have it. As such, the possible child is a contingent being, and Prior Existence Utilitarianism does not take its welfare into account. Prior Existence Utilitarians cannot argue away this implication. However, Sapontzis has argued that the incapacity to take the contingent child's welfare into account should not bother us, because once the child is actually suffering, Prior Existence Utilitarianism will take this suffering into account. This is true, and it takes away what seems to be problematic about the case, namely the idea that Prior Existence Utilitarianism would need to accept the existence of a needlessly suffering child.

Singer, however, is not satisfied with that solution, as brought forward by Sapontzis:

But this would mean that it is not wrong to decide to conceive a miserable child, although it is wrong to decide to keep the child alive once it exists. What if one knows, at the time the child is conceived, that one will have no opportunity for having an abortion or for carrying out euthanasia after the child is born? Then there will be a miserable child, so it would seem that a wrong has been done. But in Sapontzis's view, there appears to be no time at which that wrong can

⁵⁰ See Ord (2008).

be done. I cannot see that this suggestion solves the problem.⁵¹

Singer's problem with Sapontzis' view is that "there appears to be no time at which that wrong can be done". So, the challenge seems to be to indicate when the act becomes wrong.

What about the wrongness of the action? Conceiving the child is not wrong as such, because it does not make the outcome worse, it harms nobody. However, the omission of preventing the sentient fetus or child from suffering is wrong. At that moment the child counts morally and its uncompensated suffering should be prevented. However, ought implies can. What if the parents cannot have an abortion or carry out euthanasia at that point? If there is no way of preventing the child from suffering, one must conclude that letting the child suffer is not morally wrong. The assumption is that those options are not present. It seems that the couple cannot be blamed for letting the child suffer if there are no alternatives.

However, this seems odd in this case, as the couple had the possibility not to conceive the child in the first place, knowing that it will suffer and that its suffering cannot be prevented, once it is conceived. Many people are inclined to say that the couple should abstain from conception. However, according to the Prior Existence View, the possible welfare of the possible child should not count in decisions that affect whether the child will exist. Deciding whether to conceive the child is such a decision. So, the possible welfare of the possible child cannot be taken into account. So the Prior Existence View cannot accommodate the intuition that conceiving the child is the wrong action in this particular case. This implication of the Prior Existence View has been criticized:

⁵¹ Singer (1995): 228, note 20.

Surely there is something wrong with a morality that can ground no objection to doing A, the sole effect of which is to cause X to come into existence, and then tells us that we must do what we can to end the existence of X at the earliest possible stage.⁵²

I have shown at the beginning of this section that it is not the case that Prior Existence Utilitarianism cannot ground any objection to conceiving the child. I have shown that a utilitarian can object to conceiving the child. Conceiving the child is a sign of bad character, when the intention is to keep the child. Thus, there is a ground for objecting to it. This objection is available even before the child exists. Thus, contrary to what has been claimed, it is not implied by Prior Existence Utilitarianism that it cannot account for the judgment that there is something morally wrong about having the miserable child. Prior Existence Utilitarianism can subscribe to the judgment that wantonly causing such miserable beings to exist deserves moral blame. However, as we have also seen, Prior Existence Utilitarianism cannot condemn the project as such. Usually, this incapacity would not lead to acceptance of any actual suffering. Once the child suffers, the child exists as a sentient being and Prior Existence Utilitarianism will take its suffering into account. In the very farfetched case that Singer sketches - featuring a couple that wants to have the miserable child, knowing for sure before conception that the child will be miserable and not being able to have an abortion or euthanasia - Prior Existence Utilitarianism could indeed not condemn the action of the parents even though there would be a suffering child. It could, however, condemn the parents for exhibiting a bad character.

⁵² Singer (1998): 385.

6. Conclusion

The most infamous implication of Total Utilitarianism is the Repugnant Conclusion. At the deontic level, this means that a huge population consisting of very happy beings should develop into a much larger population consisting of beings whose welfare is only slightly above zero, if this is the most efficient way to raise the overall amount of happiness. Prior Existence Utilitarianism can avoid the Repugnant Conclusion, because it does not take into account the possible welfare of contingent beings. However, exactly the omission of taking into account the possible welfare of contingent beings seems to lead to an implication that many consider no less repugnant. Prior Existence Utilitarianism seems to be unable to account for the intuition that there is something morally wrong with bringing into existence a being that would surely have a miserable life. However, contrary to what has been claimed, Prior Existence Utilitarianism can account for that intuition. Even though Prior Existence Utilitarianism cannot condemn the project of having the miserable child, it can morally condemn the parents for having it, because doing so they would exhibit a bad character.

9. VEGANISM VERSUS ANIMAL-FRIENDLY ANIMAL HUSBANDRY

1. Introduction

In this last chapter before the general conclusions I will further explore the implications of Total Utilitarianism and Prior Existence Utilitarianism with regard to animal husbandry.

It has been argued that if one really cares about animals, one should better consume their products. If nobody wanted to eat their products, the animals would not be kept at all. They would not even exist. Provided that the animals have pleasant lives, the farmer seems to do the animals a favor by having them around. The animals might, as it were, be grateful to be there, just as I might be grateful that my parents had me. Thus, the idea is that animal-friendly animal husbandry is good for the animals, even if the animals are ultimately killed. A short pleasant life is better for an animal than no life at all. This influential argument in favor of animal-friendly animal husbandry and against vegetarianism and veganism has been dubbed the Logic of the Larder. I will point out how

this argument relates to Total Utilitarianism and Prior Existence Utilitarianism (section 2).

Does utilitarianism support the goal of animal-friendly animal husbandry, after all? It is now time to point out what Total Utilitarianism and Prior Existence Utilitarianism would imply for the practice of animal husbandry. Total Utilitarianism implies the Replaceability Argument and seems to be compatible with the moral goal of animal-friendly animal husbandry. However, Total Utilitarianism's alleged support for animal-friendly animal husbandry must be qualified in important ways. Prior Existence Utilitarianism, in turn, is not compatible with the Replaceability Argument. As Prior Existence Utilitarianism does not allow for the killing and replacement of animals, it is more in line with the goal of vegan agriculture (section 3).

Total Utilitarianism can in principle be compatible with the goal of animal-friendly animal husbandry, which implies that it is morally all right to use and kill animals provided that this does not infringe on the animal's welfare in an unacceptable way. Prior Existence Utilitarianism, in contrast, is more in line with the goal of vegan agriculture, because it implies, somewhat simplistically stated, that animals should not be used or killed (section 4).

2. The Illogic of the Larder

The argument that animals benefit from animal husbandry because otherwise they would not exist at all has been dubbed the Logic of the Larder (section 2.1). I will point out how this argument relates to Total Utilitarianism and to Prior Existence Utilitarianism (section 2.2). While Prior Existence Utilitarianism is incompatible with the fundamental assumptions of this argument, Total Utilitarianism might in principle be compatible with it, but might nevertheless attack the argument on practical grounds (section 2.3).

2.1 Meat eaters as real animal lovers

In 1914, Henry S. Salt has dubbed an argument that has been used in his own time and long before “the Logic of the Larder”. The argument claims that we do animals a favor by keeping them for their meat, eggs and milk, for if we did not keep them for those purposes, the animals in question would not exist. Since Salt, and in spite of his fierce rebuttal of it, the Logic of the Larder is still used. In the field of animal ethics, Salt’s term for the argument has become common. Salt has chosen that name, because the argument implies “that the real lover of animals is he whose larder is fullest of them”.¹ ‘Larder’ means ‘store-room’. The Logic of the Larder has been directed against the Jewish abstinence from pork, as well as more generally against vegetarian and vegan diets and agriculture. Most relevantly for this thesis, the Logic of the Larder has been used to defend animal-friendly animal husbandry. It has also been used to defend other uses of animals.

The Logic of the Larder is still defended, if only informally.² An explicit utilitarian defense of the Logic of the Larder is Richard Hare’s “Why I am only a Demi-vegetarian”. Hare defends the consumption of meat from ‘happy animals’, because he considers a short and happy life more valuable for the animal than no life at all. Hare explains this with respect to the free-roaming pigs in his neighborhood: “[T]hey are mainly pigs, who would certainly not be kept except for the bacon market”.³ Roger Scruton, a famous Christian thinker and traditional conservationist claims that livestock farming is beneficial for “all those, including the animals, who are part of it”. This is because “the life that is sacrificed would not exist, but for its sacrifice. A great number of animals owe

¹ Salt (1914).

² For a list of defenders of the Logic of the Larder, see Matheny and Chan (2005): 579.

³ Hare (1993).

their lives to our intention to eat them.”⁴ The Logic of the Larder is thus not a particular utilitarian argument.

A definition of the Logic of the Larder would be:

The Logic of the Larder: Killing animals is morally permissible and even obligatory, provided that:

- the animals have pleasant lives
- the animals would not otherwise have existed.

The Logic of the Larder goes further than the Replaceability Argument in the sense that it does not only aim at justifying animal-friendly animal husbandry as morally permissible. It even promotes animal-friendly animal husbandry as morally superior to vegan agriculture. According to the Logic of the Larder, animal-friendly animal husbandry cannot only be morally neutral, but even positive, morally good or obligatory. Remember that the Logic of the Larder has been brought forward against vegetarian and vegan diets. It is not about how to compensate for the welfare loss due to the killing. Rather, it is about how to maximize welfare by killing animals that exist only for that purpose.

On empirical grounds alone the Logic of the Larder is unconvincing. After all, a vegan diet would save a lot of space, energy and resources that could be used more efficiently in order to allow animals with pleasant lives to live. So, even if creating happy beings were a morally worthwhile aim to strive for, animal husbandry would not be an efficient means to that end.⁵ There are more fundamental problems with the argument that will be discussed in the following section.

⁴ Scruton (2006).

⁵ Matheny and Chan (2005).

2.2 The argument's relation with both versions of utilitarianism

The Logic of the Larder, which is about benefiting the animals kept in animal husbandry, must assume that causing an animal to exist can benefit that animal. As such, the Logic of the Larder is incompatible with Prior Existence Utilitarianism. After all, Prior Existence Utilitarianism must deny that assumption. It must instead assume that causing an animal to exist cannot benefit or harm that animal.⁶

The Logic of the Larder is compatible with Total Utilitarianism. If proponents of Total Utilitarianism also believe that causing a being to exist can benefit or harm that being, then they will accept the Logic of the Larder. If proponents of Total Utilitarianism do not believe that causing an animal to exist can benefit that animal, then they will still hold that adding happy animals to the world is obligatory, if it is the most efficient way of maximizing the quantity of welfare in the world. They will, as it were, accept an impersonal version of the Logic of the Larder: even if adding happy animals to the world does not benefit those animals, it is a morally good thing to do, in as far as it maximizes the overall quantity of welfare in the world.

When we ask ourselves whether it is morally acceptable to keep and kill animals, the possible animals in question do not yet exist. Whether they will exist depends on our answer to the very question we are contemplating. Thus, those animals are contingent beings. Proponents of the Logic of the Larder presuppose that the possible welfare of contingent beings counts. This is the same presupposition on which the Replaceability Argument is based: the Total View. It is denied by the Prior Existence View, which does not count the possible welfare

⁶This necessary assumption has been introduced in chapter 5 and defended in chapter 6.

of contingent beings. Thus, the Logic of the Larder is compatible with Total Utilitarianism but incompatible with Prior Existence Utilitarianism.

Total utilitarians might, however, dismiss the Logic of the Larder for practical reasons. For instance, Peter Singer dismisses an argument of the nineteenth-century British philosopher Leslie Stephen, who wrote:

Of all the arguments of Vegetarianism none is so weak as the argument from humanity. The pig has a stronger interest than anyone in the demand for bacon. If all the world were Jewish, there would be no pigs at all.⁷

Singer responds to Stephen's argument by pointing out that animals do usually not lead pleasant lives in animal husbandry. Furthermore, Singer explains that proponents of this argument should not defend animal husbandry at all, because "with the possible exception of arid areas suitable only for pasture, the surface of our globe can support more people if we grow plant foods than if we raise animals."⁸ Furthermore, according to Singer, a greater number of happy animals could live if one reduced the number of human beings in favor of a greater number of happy mice.

Remarkably, Singer does not distinguish between the Logic of the Larder and the Replaceability Argument. Singer claims that one "may call this 'the replaceability argument'".⁹ In my view, Stephen's view as cited by Singer expresses the Logic of the Larder, and not the Replaceability Argument. As I explained, those arguments are related, but not identical. Strictly speaking, the Logic of the Larder needs not be about replacement at all. It is rather about adding happy beings to the

⁷ Stephen, quoted in Singer (1993): 121. See also Singer (forthcoming 2011): 105.

⁸ Singer (1993): 122.

⁹ Singer (1993): 121. Singer (forthcoming 2011): 106.

world and thereby benefiting them. Singer's reply is directed against what I presented as the Logic of the Larder, rather than against what I presented as the Replaceability Argument.

Singer links the Logic of the Larder to the Total View. Singer accepts the Total View and sees nothing wrong with the underlying logic of the Logic of the Larder argument as such. The underlying logic is that adding happy beings to the world is morally just as good as making existing beings happier. Furthermore, the underlying logic is that causing a happy being to exist benefits that being. Indeed, Singer claims in this context: "[...] it is difficult to explain why we do not do something good when we knowingly bring a happy being into existence."¹⁰ Thus, Singer subscribes to the underlying logic of the Logic of the Larder argument. Singer's above-mentioned practical counter arguments do "merely" show that this logic does not serve the case for animal husbandry.

Note that the assumption that bringing a happy being into existence benefits this being is in contradiction with Singer's suggestion, as discussed in chapter 4, that welfare can only be neutral or negative. Singer made this latter suggestion in order to be able to restrict the scope of the Replaceability Argument. I already indicated in chapter 4 that this move would be at odds with the rest of Singer's work, where he assumes that welfare can indeed be positive. Singer's acceptance of the underlying logic of the Logic of the Larder is a clear example. Taking serious the view that welfare can only be negative or neutral would allow Singer to dismiss the Logic of the Larder on a fundamental level. As we have seen in chapter 4, it would also imply that most of us had better not lived.

Unlike Singer, who accepts the Total View and thus seems to accept the underlying idea of the Logic of the Larder, proponents of the Prior Existence View will find fault with the Logic of the Larder argument at a more fundamental level. They can do so, without having

¹⁰ Singer (1993): 123.

to accept that most of us had better not lived. Prior Existence Utilitarians will deny that existence can be better or worse for an animal than non-existence. Furthermore, they will hold that outcomes should be evaluated in terms of what they do for sentient beings, rather than what they do for the overall quantity of welfare as such. Note that Total utilitarians, who accept the underlying logic, need an account about whether it principally holds for human beings as well. Singer, for instance, describes a case of people being brought into existence and kept in a state of happy infants as a reservoir of spare body parts. Singer claims that this implication would “reduce the appeal” of the argument.¹¹ As far as I can see, though, a Total utilitarian accepting Singer’s view on the harm of death seems to have nothing in principle to bring forward against this implication.

Summing up, an influential argument in defense of animal-friendly animal husbandry claims that animal-friendly animal husbandry actually benefits the animals in question, because if it were not for being kept in animal husbandry, those animals would not exist at all. This argument is not compatible with Prior Existence Utilitarianism. After all, the argument assumes that existence can be better for a being than non-existence, an assumption that the Prior Existence utilitarian must deny. Total utilitarians, in contrast, need not find fault with this underlying assumption of the argument. Total utilitarians do favor adding happy beings to the world provided that this maximizes welfare, either because they do not evaluate outcomes in terms of harms and benefits to sentient beings but in terms of the total overall quantity of welfare, or because they believe that causing a being to exist can benefit that being, or both.

¹¹ Singer (forthcoming 2011): 107.

3. Implications concerning animal husbandry

So, what, more precisely, are the implications of Prior Existence Utilitarianism and Total Utilitarianism for the practical issue with which I started: animal husbandry. What would follow from both versions of utilitarianism for the practice of animal husbandry?

As we have seen, Total Utilitarianism implies the Replaceability Argument. Therefore, it seems to allow for the killing and replacement of animals and thus be compatible with the goal of animal-friendly animal husbandry. However, there are some caveats concerning this conclusion.

First of all, as we have seen in chapter 4, the Replaceability Argument only permits the killing and replacement of an animal under certain conditions. One condition is that the animal has a pleasant life. More precisely, it is crucial that the killed animal could have had a pleasant future if it were not killed. This is a condition of the Replaceability Argument, because if the animal had an unpleasant future, there would be no need to replace it, from a utilitarian perspective. A second condition of the Replaceability Argument is that the killing does not have any unbalanced negative side effects, such as causing suffering for the animals or for others. As has been elaborated in chapter 4, it is doubtful whether those conditions are fulfilled in many practices of what is considered animal-friendly animal husbandry. Usually, the process of the killing has a negative impact on the animal's welfare. Furthermore, it is not clear that the overall welfare level of an animal that is kept for human consumption is always positive. After all, even in more animal-friendly practices of animal husbandry animals might suffer from boredom, stress, fear and/ or pain and might lack sufficient opportunities for having positive experiences. However, practices of animal husbandry are conceivable that score better in terms of animal welfare than current practices of animal-friendly animal husbandry. Thus, as far as the conditions of the Replaceability Argument are concerned,

certain forms of animal-friendly animal husbandry might in principle be justified.

Secondly, as we have seen in chapter 4, even if the Replaceability Argument is accepted and if the relevant conditions are fulfilled, it is controversial which beings would be replaceable and which would not. As elaborated in chapter 4, Singer tries to restrict the scope of the Replaceability Argument to the effect that persons are irreplaceable. By 'persons' Singer indicates those beings that are self-aware and have future oriented desires and a conception of life and death. Only those beings, according to Singer, have a preference for continued life. Only persons are therefore significantly harmed by death. As we have seen in chapter 4, Singer appeals to the moral ledger model, which implies that unfulfilled desires count negatively on the welfare scale and the fulfillment of a desire can only restore the negative count to zero. Therefore, killing a person would leave unsatisfied desires, such as the desire to go on living, while any newly created being could never score better than zero on the welfare scale. According to Singer, personhood comes in degrees, as does replaceability. I have discussed the implications of Singer's arguments for limiting the scope of the Replaceability Argument. I have argued that it is pessimistic, dependent on a particular account of the harm of death, and ad-hoc in several respects. Furthermore, it fails to restrict replaceability in the desired way. Those who, in spite of all this, would like to accept Singer's account of the scope of the Replaceability Argument should note that Singer is prepared to ascribe personhood to the animals that are usually killed for consumption, or at least to give them the benefit of the doubt.

For instance, Singer discusses evidence about the capacities of pigs and other animals and ends up with the following remark:

We think of dogs as being more 'human' than pigs, but we have already seen that pigs can plan ahead and grasp whether another pig does or

does not know the location of food. Are we turning persons into bacon?¹²

Likewise, Singer discusses the capacities of chicken, which in flocks of up to ninety birds appear to recognize one another as individuals, and who have shown evidence of self-control and future planning.¹³ Furthermore, there is growing insight into the capacities of fish, and even some invertebrates, such as the veined octopus have been observed to show behavior that suggests forward planning.¹⁴ In case of doubt about personhood, Singer suggests to give a being the benefit of the doubt when considering whether or not to kill it.¹⁵

Finally, Singer has repeatedly argued against animal production and consumption on grounds that are not related to animal welfare. As has been mentioned in chapter 2, a complete moral evaluation of the practice of animal husbandry would not be restricted to the effects on the animals involved. It would also include effects on human beings all over the globe, and also on future generations. Topics such as environmental effects and world food distribution need to be considered in this regard. This has led Singer and others to condemn animal husbandry on utilitarian grounds.

Hence, it seems that even Total utilitarians, which accept the Replaceability Argument, would not necessarily support the goal of animal-friendly animal husbandry, all things considered.

What would Prior Existence Utilitarianism imply for the practice of animal husbandry? The ideal form of agriculture, according to Prior Existence Utilitarianism, would be vegan agriculture. After all, Prior

¹² Singer (forthcoming 2011): 102.

¹³ Idem.

¹⁴ Idem.

¹⁵ Singer (forthcoming 2011): 103.

Existence Utilitarianism, as a form of utilitarianism, strives for the maximization of welfare. Animal welfare is taken fully into account. So, Prior Existence Utilitarianism would condemn practices of animal husbandry that involve serious animal suffering. Furthermore, killing an animal which otherwise could have continued a happy life counts as a welfare loss. According to Prior Existence Utilitarianism bringing another animal into existence cannot compensate for this welfare loss. Therefore, Prior Existence Utilitarianism rules out the routine killing of animals that is implied by the practice of animal husbandry. So, it seems like Prior Existence Utilitarianism would be opposed to animal husbandry altogether.

Vegan agriculture refers to an agricultural system that does not use animals for food production. Thus, it would not produce meat, eggs, dairy or other animal products for which animals would be caused suffering or would be killed. Prior Existence Utilitarianism seems to rule out animal husbandry, as I defined it in this thesis. That, however, does not mean that it would be opposed to all forms of consumption of animal products or to all uses of animals. For instance, if one could manage to receive a bit of a cow's milk without causing any suffering or death for the cow or her calf, this would not be ruled out. If one would eat an animal that had died a natural death or was killed in an accident, this would be permitted. One might eat non-sentient animals, if they exist. (If there is any reasonable doubt as to whether an animal is sentient, it should get the benefit of the doubt.) Furthermore, one might eat meat that is grown without the use of animals. There are attempts to produce meat from stem cells, but this is not an available alternative yet. Animal husbandry, which involves the infliction of harm and killing of animals would be ruled out.

Summing up, Prior Existence Utilitarianism does not support the ideal of animal-friendly animal husbandry, because it does not allow for the routine killing of animals that could otherwise have had a pleasant future. Prior Existence Utilitarianism does not allow the killing of happy

animals, so to say. Total Utilitarianism principally allows for the killing and replacement of animals. However, due to the conditions of the Replaceability Argument and due to the whole range of effects of animal husbandry on the welfare of humans and nonhumans, and due to the capacities of animals, which according to some interpretations might exclude them from replaceability, even proponents of Total Utilitarianism might all things considered not support animal-friendly animal husbandry. For those proponents of Total Utilitarianism, as well as for all proponents of Prior Existence Utilitarianism, the practice of animal-friendly animal husbandry cannot be morally justified. Vegan agriculture is considered a morally preferable alternative.

4. Conclusion

Total Utilitarianism can in principle be compatible with animal-friendly animal husbandry. Total Utilitarianism implies that it is morally all right to use and kill animals provided that this does not infringe on the animal's welfare in an unacceptable way. Total utilitarians accept the Replaceability Argument. Furthermore, they might accept that animal husbandry is morally preferable to vegan agriculture, either because granting animals pleasant lives benefits them, or because having more beings with pleasant lives around is a good thing, or both. Prior Existence Utilitarianism, in contrast, is more in line with the goal of vegan agriculture. After all, according to Prior Existence Utilitarianism, non-existent animals cannot benefit from being caused to exist. Therefore, it is not a morally good thing to bring additional animals into the world. Furthermore, Prior Existence Utilitarianism considers killing animals that could otherwise have continued pleasant lives morally problematic. According to Prior Existence Utilitarianism, bringing another animal into existence cannot compensate the welfare loss that is caused by the killing of an animal that could otherwise have had a pleasant future

10. GENERAL CONCLUSIONS

It is now time to draw conclusions (section 1) and to make a note of remaining questions (section 2).

1. Conclusions of the thesis

The Total View and the Prior Existence View are two competing coherent utilitarian views on whose welfare to take into account in the aggregation of welfare. Total Utilitarianism considers the welfare of all actual and all possible beings. Prior Existence Utilitarianism takes into consideration the welfare of all actual and all necessary beings. These are those that do already exist and those that will definitely exist, independently of the moral choice that is contemplated. Prior Existence Utilitarianism does not take into account the welfare of contingent beings, i.e. those possible beings whose existence depends on the moral choice that is contemplated.

Total Utilitarianism and Prior Existence Utilitarianism in principle yield different verdicts about the permissibility of killing

animals that could otherwise have had a pleasant future. Killing such “happy animals” is usually implied by the production of meat, dairy and eggs, in so-called “animal-friendly” animal husbandry systems. Killing “happy animals” is also part of aquaculture, sports hunting, sports fishing, and the breeding of animals for animal experimentation. In those practices, animals that could otherwise have had a pleasant future are routinely killed on a massive scale. It is typical for those practices that other animals replace the animals that are killed. What is more, those latter animals would not exist if the other animals were not killed. Furthermore, the lives of the newly created animals are at least as pleasant as the future lives of the killed animals would have been.

How do Total Utilitarianism and Prior Existence Utilitarianism evaluate such practices of killing and replacing animals? When considering whether it is permissible to kill an animal that could otherwise have had a pleasant future, Total Utilitarianism takes into account the effects of that choice on the welfare of the existing animal, which one considers to kill. It also takes into account the effects of the choice on the welfare of the possible future animal that would be brought into existence if and only if the other animal would be killed. Provided that the life of this newly created animal contains at least as much welfare as the future life of the killed animal would have contained, killing and replacing the animal yields as much welfare overall as letting the animal live. In such a case, Total Utilitarianism would in principle allow for the killing and replacement of an animal. Prior Existence Utilitarianism, in contrast, does not take into consideration the possible welfare of the contingent animal that might live if and only if the other animal will be killed. According to Prior Existence Utilitarianism bringing into existence another animal cannot compensate for the welfare loss, which is caused by killing an animal. Thus, these different utilitarian views on whose welfare to consider have different implications concerning the permissibility of killing.

Total Utilitarianism accepts the Replaceability Argument, as explained above. It accepts that it is permissible to kill an animal that could otherwise have had a pleasant future, provided that another animal is put in the killed animal's place, whose life is at least as pleasant as the killed animal's future life would have been and which would not otherwise have existed. Furthermore, the killing should not have any uncompensated negative side effects. The Replaceability Argument does not only hold for those animals that are typically used for human consumption. Those who accept the Total View, and thus the Replaceability Argument must accept that it holds for all animals, including human beings.

Peter Singer has tried to restrict the scope of the Replaceability Argument by arguing that the more future-directed projects a being has, the less replaceable it is. Those beings who have many and strong future-directed desires are called 'persons', whereby Singer points out that personhood - and therefore also replaceability - comes in degrees. When animals such as pigs, cows, chickens and fish are ascribed personhood, or at least given the benefit of the doubt in this respect, as Singer proposes, those animals would be excluded from replaceability. This, however, works only if one accepts Singer's account of welfare and his 'moral ledger model'. This means that one needs to accept a desire-satisfaction account of welfare and assume that unsatisfied desires have a negative effect on welfare, which can only be neutralized in case the desire is fulfilled. Under those assumptions, killing a person would usually leave many strongly held desires unsatisfied, such as the desire to continue life and other future-directed projects. Those unsatisfied desires would leave negative welfare, which, according to the moral ledger model, cannot be compensated by bringing another person into existence. This is because the welfare of the other person could never be better than neutral, and it can be neutral only in the exceptional case in which all his or her desires would be fulfilled.

Thus, in order to exclude persons from replaceability, Singer proposes the desire-satisfaction account of welfare together with the moral ledger model. As we have seen, under those assumptions welfare can never be positive. It can only be negative or, in the unlikely case in which all one's desires are fulfilled, neutral. This account of welfare, as I have argued, seems overly pessimistic. Unlike other accounts of welfare, this particular account of welfare cannot capture the experience that welfare can also be positive, and that, for instance, the satisfaction of at least some desires does more than only wipe out a state of deprivation. Furthermore, the particular account of welfare that Singer proposes in order to exclude persons from replaceability is ad-hoc in several respects. For instance, it does not fit with the rest of Singer's work, where he seems to assume that welfare can either be positive, negative or neutral. Furthermore, it would seem that every being with desires would be irreplaceable, because it would be likely to leave unsatisfied desires when being killed. It seems arbitrary to claim that persons are irreplaceable, while those who have less future directed desires are replaceable, to a certain degree.

An implication of the account of welfare that Singer proposes in order to exclude persons from replaceability is that most of us had better not lived. After all, we are all likely to leave at least some unsatisfied desires when we die. Thus, on that version of utilitarianism, bringing any person into the world seems to be morally forbidden, because remaining childless would better serve the maximization of welfare by avoiding the likely negative welfare scores due to unsatisfied desires. According to that account of welfare, it would have been better if all of us, both human and non-human persons, would never exist. Furthermore, as soon as any being has all of his or her preferences fulfilled, he or she would be replaceable by another being that would be equally 'lucky'.

Singer has recently tried to avoid this implication by accepting preference-independent values. If those values are still about welfare, this move does not help. Either welfare can be positive and hence

replaceability is possible, or welfare can only be negative and neutral and hence the implication is that those of us who are not replaceable had better not lived at all. When the desire-independent values that Singer proposes to accept are also welfare-independent, the proposed view would fall outside utilitarianism. It would no longer hold that welfare is the sole ultimate value. Furthermore, this move would not help. Again: either the value of a person's life can be positive and hence replaceability is possible, or the value of a person's life cannot be positive, and hence most of us had better not lived.

Even if adherents of the Total View accept Singer's model for restricting the scope of the Replaceability Argument, there will always be beings, both human and non-human, who fall short of the criterion of full personhood, and who would therefore be replaceable. For instance, babies and newborn infants would be replaceable, according to Singer's model. This is because they lack any strong future-directed desires. That means that they could in principle be killed even if they would otherwise have had a pleasant future (thus, we are not talking about euthanasia here).

Prior Existence Utilitarianism does not accept the Replaceability Argument, which has been considered the most controversial part of Singer's moral theory. It would not sanction the routine killing and replacement of animals. Neither a human nor a non-human animal is deemed replaceable. In that sense, Prior Existence Utilitarianism accords all sentient beings a stronger protection. That would be good news for all who have criticized utilitarianism for not taking animals - including human beings - seriously enough.

Related to this implication of taking animals (both human and non-human) more seriously, Prior Existence Utilitarianism, as opposed to Total Utilitarianism, is a truly person-affecting moral theory. Instead of being concerned with the abstract quantity of welfare, Prior Existence Utilitarianism is concerned with effects on sentient beings. It is

concerned with harms and benefits for sentient beings. Prior Existence Utilitarianism always requires producing maximal welfare (defined as net benefits) for sentient beings, but - unlike Total Utilitarianism - it never requires producing sentient beings in order to maximize welfare.

This means that Prior Existence Utilitarianism implies the Person-Affecting Restriction on what matters in the evaluation of outcomes. This is the view that outcomes should be evaluated in terms of harms and benefits for sentient beings. I have defined this view in the wide sense. The Wide Person-Affecting Restriction is concerned with harms and benefits for sentient beings, *not* as particular individuals, but rather as sentient beings *whoever they are*. This wide interpretation of the Person-Affecting Restriction fits better with utilitarianism and it is necessary in order to avoid the Non-Identity Problem. Thus, the focus in the evaluation of outcomes according to the view that is implied by Prior Existence Utilitarianism is on aggregate net benefits.

Total Utilitarianism, in contrast, implies the Impersonal View on what matters in the evaluation of outcomes. According to that view, outcomes are not morally evaluated in terms of how they affect people (or sentient beings), but rather in terms of the abstract amount of welfare that they contain. The Impersonal View is an intrinsic aspect view on the evaluation of outcomes. The outcome that contains the greatest amount of welfare is considered best.

A famous implication of this intrinsic aspect view on the evaluation of outcomes is that a population consisting of well-off people can be required to multiply, even if this reduces the welfare of the already existing people and even if everyone, including the newly created people, will end up with a welfare level that is barely positive. The Impersonal View has this implication, because it is only concerned with the quantity of welfare that an outcome contains, and not with benefits and harms for sentient beings. If the resulting population is large enough, the sum of everybody's welfare can be greater than the overall amount of

welfare that was present in the original population consisting of a smaller number of very happy beings. This implication, that a population consisting of very well-off beings can be required to develop into a much larger population consisting of barely happy beings has been dubbed the Repugnant Conclusion.

Prior Existence Utilitarianism does not lead to the Repugnant Conclusion. Prior Existence Utilitarianism would not sanction adding additional beings to a population, at the expense of those who already exist or will necessarily exist. This is because Prior Existence Utilitarianism does not include the welfare of contingent beings in the aggregation of welfare.

This has led to the challenge, famously exemplified by Parfit's case of the Wretched Child, according to which Prior Existence Utilitarianism cannot account for the expected misery of a contingent miserable child. In Parfit's case a couple knows for sure before conception that any child it could have would have a short and miserable life. According to the Prior Existence View, the welfare of this contingent child cannot be taken into account in the decision about whether or not to have that child. However, there would be other reasons for not having this child, such as the welfare of the prospective parents and other beings that already exist or will necessarily exist.

In the very unlikely case that the parents would want to have this child, and others would not be negatively affected, what could proponents of Prior Existence Utilitarianism say? They could say that conceiving the child with the intention of having it in spite of its suffering could be condemned on utilitarian grounds as a sign of bad character. After all, being prepared to let others suffer in that way for one's own satisfaction does usually not contribute to the maximization of welfare. Prior Existence Utilitarianism, just like any other form of utilitarianism, can directly evaluate character, next to acts and other evaluative focal points. Accepting several evaluative focal points has the

consequence that utilitarianism is able to account for moral dilemmas. For instance, when a mother faces the choice between rescuing her own child or several strangers from a burning building, it seems that caring particularly about one's own child might be a good character, or disposition, in utilitarian terms. However, in extraordinary cases as this one, it might lead to a sub-optimal outcome in terms of welfare. The mother cannot both, have the optimal character and choose the optimal action, both judged in utilitarian terms. Being able to acknowledge that kind of dilemmas seems to be an advantage of the acceptance of several direct evaluative focal points, such as actions and character.

Thus, proponents of Prior Existence Utilitarianism can claim that there is something morally wrong with parents that want to have such a miserable child. Furthermore, once the child is actually suffering, Prior Existence Utilitarianism does take this suffering fully into account, because by then the child actually exists. As any other form of utilitarianism, Prior Existence Utilitarianism might then require ending the suffering of the child by an act of abortion or euthanasia. Thus, Prior Existence Utilitarianism can avoid the Repugnant Conclusion and it can deal with the Expected Misery Argument.

Proponents of Total Utilitarianism need not accept the Impersonal View on what matters in the evaluation of outcomes. Just as proponents of Prior Existence Utilitarianism, they might accept the Person-Affecting Restriction, according to which outcomes should be evaluated in terms of harms and benefits to sentient beings. This, however, is only possible for proponents of Total Utilitarianism, if they accept that bringing a being into existence can harm or benefit that being. Only then it makes sense to take into account the possible welfare of contingent beings in the aggregation of welfare, from a person-affecting perspective. Therefore, if an adherent to Total Utilitarianism is to accept that outcomes should be evaluated in terms of harms and benefits to sentient beings, she must claim that bringing a being into existence can benefit that being. This view, however, is hard to defend.

Bringing a being into existence cannot benefit or harm that being, because the welfare level of an existing being cannot be compared to the welfare level of that being before it existed. This is because in case of non-existence, a being has no welfare level whatsoever. The temptation of putting the welfare level in case of non-existence at zero should be resisted. Zero-welfare is a situation in which positive and negative welfare hold the balance, and that is different from the situation of non-existence to which the welfare level is inapplicable.

Unless proponents of Total Utilitarianism can make a convincing case for non-existent beings having a welfare level, they cannot claim that bringing a being into existence can make this being better or worse off in terms of welfare. If this claim is impossible, the possibility of maximizing welfare by making happy beings, as accepted by the Total View, cannot be explained in person-affecting terms. Therefore, it seems that Total Utilitarianism does indeed fit better with the Impersonal View on the evaluation of outcomes.

A major criticism against utilitarianism is that it is overly impersonal, not interested in sentient beings, but only in the maximization of welfare as an abstract quantity. I have shown that Prior Existence Utilitarianism is a truly person-affecting version of utilitarianism. Prior Existence Utilitarianism does not imply that making happy beings is on a par with making beings happy. It does not treat sentient beings as replaceable. It is focused on maximizing aggregate net benefits for sentient beings. An underlying rationale for such a form of utilitarianism can be the principle of equal consideration. If that rationale underlies the duty to neutrally maximize welfare, which is defined as aggregate net benefit, arguably the respect for sentient beings is accounted for at the most basic level of utilitarianism. Total Utilitarianism, in contrast, fits better with the teleological rationale, according to which welfare as such should be maximized, because it is good, no matter whether maximizing welfare benefits or harms sentient beings.

Finally, let me get back to animal husbandry. A still very popular argument in defense of animal-friendly animal husbandry is called the Logic of the Larder. It holds that animals actually benefit from animal-friendly animal husbandry, because the animals have pleasant lives and they would not exist at all if it were not for being killed. This argument might seem to be a utilitarian argument. The underlying idea is that a short and pleasant life is better for the animal than no life at all.

Proponents of Total Utilitarianism accept the Logic of the Larder, either because they accept that bringing a being into existence can benefit that being, or because they evaluate outcomes in terms of the abstract quantity of welfare, which they contain, or both. Again, those who accept the Logic of the Larder must accept it in principle for all animals, both human and non-human. For instance, Singer has evoked a thought experiment of babies being produced for the purpose of organ donation. Provided that those babies are granted pleasant lives, and are killed painlessly, and provided that nobody else suffers significantly from this practice, using those babies as a source of spare organs might not only be morally justified, but even morally required as an option that maximizes welfare. This would just be another application of the Logic of the Larder.

Prior Existence Utilitarianism does not accept the Logic of the Larder. In order to be a coherent utilitarian view, Prior Existence Utilitarianism must deny that bringing an animal into existence can benefit that animal. After all, Prior Existence Utilitarianism does not take into account the welfare of contingent beings. This is only consistent with the utilitarian duty to neutrally maximize welfare, if one evaluates outcomes in terms of harms and benefits for sentient beings *and* if one accepts that bringing a being into existence cannot harm or benefit that being. Only then does it make sense to disregard the possible welfare of contingent beings in striving for neutral welfare maximization. Thus Prior Existence Utilitarianism does not accept that bringing into existence can benefit a being. Killing a being that could otherwise have

had a pleasant life, harms that being. Therefore, Prior Existence Utilitarianism does not accept the Logic of the Larder.

In conclusion, whether practices of killing and replacing animals, such as meat and dairy production, aquaculture, sport hunting, sport fishing or breeding animals for animal experimentation are acceptable from within utilitarianism depends on whether one accepts Total Utilitarianism or Prior Existence Utilitarianism. I have here presented the assumptions and implications of both versions of utilitarianism. In short, Prior Existence Utilitarianism is not impersonal but truly person-affecting. It provides animals, both human and non-human, a stronger protection against killing.

Hence, in opposition to what is typically brought forward in the classrooms and in the literature, there is more to utilitarian animal ethics than Singer's position. Utilitarianism is correctly recognized as being the moral theory that, historically, has contributed most to the recognition of animal suffering as an evil. However, the utilitarian concern with animals is not restricted to the avoidance of suffering. Utilitarianism has the resources to oppose the routine killing of animals, as practiced in animal husbandry, and in the other above-mentioned practices of animal use.

2. Remaining questions

I have explored the assumptions and implications of Total Utilitarianism and Prior Existence Utilitarianism. This might make one wonder which, if any, of those versions of utilitarianism can be justified. I have not asked, nor answered the question of justification in this thesis. In order to answer the question about justification, one needs to point out what it means for a moral theory to be justified and what it takes to justify a moral theory, or any particular version of it. This is highly controversial. Hence, arguing about the justification of any of those theories would unavoidably lead into meta-ethical disputes about justification. There are

different, more or less demanding views on what counts as a justification. Either way a sufficient justification of the theory is necessary in order for it to give us reasons. Then, there is the question what the force of those moral reasons should be, in relation to other, non-moral reasons. Those questions have been left untouched here, as they fall outside the scope of this thesis. However, answering these questions is necessary before we know how we ought to act.

Next to those very fundamental questions about justifications - the roots, so to say, of the moral theory - there are questions about the offshoot of the theory. (Being well aware that without roots there is no offshoot.) I have gone some way towards indicating Total Utilitarianism's and Prior Existence Utilitarianism's offshoot by discussing their practical implications. For instance, I have indicated what they would imply for the practice of animal husbandry. However, in order to be action guiding, it would be necessary to become more concrete. It is one thing to point out what the ideal would be according to each version of utilitarianism, be it, for instance animal-friendly animal husbandry or vegan agriculture. It is another thing to point out which concrete actions it takes by whom in order to aspire to the realization of that ideal. A question that has already been put on the agenda in animal ethics in this regard is whether for instance the vegan ideal can best be strived for by focusing on better welfare for kept animals, or rather by incrementally abolishing the use of animals for certain purposes, or both.¹ According to utilitarianism obviously one should do what maximizes welfare for all concerned on the long run. However, it is disputed what this entails. In order to become very practical, it is not sufficient to make clear what the ideal should be, but also to point out how it can best be realized.

¹ Sandoe (2008): 80-84.

APPENDICES

A. Different-number choices

As we have seen in chapter 7, Prior Existence Utilitarianism can successfully deal with standard non-identity cases. Standard non-identity cases are presented as same-number choices: in both outcomes, the same number of beings exists. So, for instance, the 14-years old girl will have one child, and the question is only when she will have it, and the timing affects the genetic identity of the child. Different number choices present a further challenge for the Prior Existence View.

In different number choices, a different number of beings exist in each outcome. So, for instance, the timing of the conception might determine not only the welfare level and genetic identity of the woman's next child, but it determines also whether she will have twins or rather a single child. In this appendix, I will explore whether and, if so, how the Prior Existence View can deal with different number choices.

Decisions that affect the identity and welfare of possible beings, might also affect their number. For instance, policy choices that we make now, might affect the welfare level of the future generation, and also the identities of who will live. If this is the case, those policy choices are also

very likely to affect the size of the future generation and thus *the number of people that will live*. Therefore, that kind of policy choice seems to show the prevalence and relevance of different number choices. However, as I have pointed out in chapter 7, that kind of policy choice does not present us with genuine non-identity cases. As Roberts has argued, the same people might exist in both outcomes, so it is not necessarily the case that they would not otherwise have existed. Furthermore, as Parfit himself says, that kind of policy choice can better be construed as a Same Number choice, because it is unclear what the impact is on the number of beings that will live: “There would of course be some difference in numbers. But this we can morally ignore, since it would either not now be predictable, or would not be the feature which is morally important.”¹

So, let me make up a better example of a different number choice. We can slightly change the case of the 14-years-old girl and assume that the girl is on a medication now that would cause her to have twins if she conceived now. In a couple of years, she would not need this medication anymore. So if she conceived later, she would have only one child. So, the choice would be between one happy child later or two less happy children now. Here is another, similar, case, based on a real story:

In Vitro Fertilization: A woman is having an in vitro fertilization and she has to choose whether to implant one or rather two of the fertilized eggs. Thus, she will either have a single child or twins.² If she has twins, each of them will be somewhat worse off than a single child would be.³ As she would soon pass the age-

¹ Parfit (1981): 122.

² In real life, things are more complex: each implanted egg might develop into twins, and it is also unsure whether it will develop into any child at all.

³ We can assume that this is due to complications during gestation or birth, or because of somewhat limited resources of the mother to provide for them. Again, this is a

boundary of 40 years, this would be her last chance of having children.

Now, what would it mean to do the best for “her child/ children”? How should the fact of the different numbers be taken into account? Any view about the moral status of possible beings should also tell us how to deal with different number choices.

Can the Wide Person-Affecting View handle different number choices? I have defended the following approach to different people (same number) choices: “The woman will have her next child in both cases. So, it is a necessary being. Genetic make-up is irrelevant for her obligations concerning her next child. So, she should do what is best for her next child, whoever it will be.” Yet, what should we say if the choice affects whether she will have one child or twins? We might say: “The woman will have her next child/children in both cases. So, it is a (or they are) necessary possible being(s). Genetic make-up and number are irrelevant concerning her obligations towards her next child/children. So, she should do what is best for her next child/ children.” This second way of framing the situation seems to be possible. Caspar Hare, for one, mentions as an advantage of his account of de-dicto betterness, as discussed in chapter 7, that it is applicable to same-number choices as well as to different-number choices.⁴ Unfortunately, Hare does not show what exactly de-dicto betterness in different-number choices amounts to. Neither does Clark Wolf show how his treatment of standard non-identity cases, where he employs a loose notion of identity (as discussed in chapter 7) can be applied to different number choices. Wolf recognizes the problem and says that he believes there is “a natural way”

simplified assumption. In real life, those things need not influence the welfare level of the children.

⁴ Hare (2007).

to accommodate it.⁵ Yet, unfortunately, he offers no clue about what this way would be.

It is not obvious how different-number choices should be handled, departing from the Wide Person-Affecting Restriction. The most straightforward way seems to be to treat number like genetic identity as morally irrelevant, and to focus exclusively on overall welfare. So, in the above-mentioned case of en-vitro fertilization, the woman should, all else being relevantly equal, make the choice that makes “her next child/children” better off. What would this amount to, when the number of children in each outcome differs? In case of twins, a method of aggregation has to be used in order to determine the overall welfare of the outcome.

Which method of aggregation should be used? The most obvious possibilities would be either summing up as suggested by the total view, or taking the average, as suggested by the average view. Taking the average might require to have one child, rather than two, if that single child’s welfare level is higher. Although the average view yields intuitively acceptable results in some cases, it has very implausible implications. It requires killing beings whose welfare level is beneath average, if this raises the average level of welfare. It also implies that having a child whose welfare level is beneath average is, *ceteris paribus*, unacceptable, even if the child’s welfare level were very high. Counter-intuitive implications are no reason for rejecting a view. However, killing all but the happiest being seems contrary to what utilitarianism is all about: maximizing welfare for sentient beings. The alternative method of aggregation, summing up, might imply that having twins is better, even if every child as such were worse off than the single child in the alternative outcome would be. This implication somewhat resembles the Repugnant

⁵ Wolf (2009): 113, note 23.

Conclusion: The total amount of welfare is bigger, although each individual is worse-off.

Does the choice of the total view as a method of aggregation mean that the Prior Existence View also implies the Repugnant Conclusion? Among necessary beings, the greatest amount of welfare might indeed be found in a group of beings, of which each individual is worse off than the individual(s) in an alternative outcome. In order to illustrate this, imagine a slightly different version of the In-Vitro Fertilization case. Imagine that the woman would be somewhat worse off if she had twins. Imagine that the aggregate welfare of the woman and the twins would be bigger than the aggregate welfare of her and the single child. The Prior Existence View would count the welfare of all, because the woman and 'her next child/children' are all necessary beings. Differences of genetic identity and number are supposed to be irrelevant. The Prior Existence View in conjunction with the total view as a method of aggregation would favor having the twins.

I think that this is an acceptable implication. It differs crucially from the Repugnant Conclusion. The Total View in conjunction with the total view as a method of aggregation implies that contingent beings must be brought into existence provided that this maximizes overall welfare. This will be required, even if existing and other necessary beings will be worse off as a consequence. The Prior Existence View does not have that implication. The Prior Existence View only counts the welfare of necessary beings, not of contingent ones. Thus, it allows aggregation of welfare only amongst necessary beings. Those *who live or will live anyway* should be as well off as possible. Identity does not matter, neither does number. It sounds unfamiliar to talk about "those who will live anyway", while at the same time acknowledging the variability of their identity and number. This is the *wide* interpretation of the Person-Affecting Restriction. Choosing the twins in this case is compatible with striving for the best outcome overall. It is thus in line with utilitarianism and the Wide Person-Affecting Restriction. So, aggregating as proposed by the

total view seems to be a plausible thing to do in case of different number choices.

Prior Existence Utilitarianism is not a theory about the value of states of affairs as such; it is only interested in the evaluation of outcomes of possible actions. In order to decide which outcome is better in the In-Vitro Fertilization case, both possible resulting populations are compared in terms of welfare. In fact, all kinds of populations can be compared with each other in terms of welfare. For instance, the welfare of a population of extraterrestrials on planet A can be compared to a population of extraterrestrials on planet B. This, however, is not the business of Prior-Existence Utilitarianism. Utilitarianism is a moral theory, a theory about what we ought to do. It is not a theory for determining the value of states of affairs. Of course, axiological considerations, i.e. determining the value of states of affairs, can be relevant for utilitarianism. But only in so far as the states of affairs are outcomes of possible actions. Utilitarianism requires the maximization of welfare, and the Prior Existence View determines the scope of moral objects. As Arrhenius puts it:

Interpreted as normative theories, the restriction to outcomes of alternative actions is perfectly legitimate. Since a normative theory concerns what we ought to do, and since “ought” implies “can” (in some sense), a normative theory is not in the business of ordering populations that are not outcomes of alternative actions in some choice situation.⁶

The Replaceability Argument does not present us with a different number case, and not even with a typical non-identity case at all. The Replaceability Argument implies a different people choice, or better: a

⁶ Arrhenius (forthcoming).

different beings choice. If the farmer chooses *not* to kill cow A, cow A will live. If the farmer chooses to kill cow A and replace her, a different cow, cow B, will exist instead. So, in both outcomes, different beings live. However, unlike typical non-identity cases, one of those beings is a presently living cow, while the other is a possible future cow. Therefore, the Replaceability Argument does not imply a typical non-identity case. In a typical non-identity case, only the genetic identity and the welfare level of the possible beings differ. Everything else is relevantly similar. In case of the Replaceability Argument, one of the beings actually lives and would have to be killed, in order to make it possible for the other to live. This makes that case crucially different.

Genuine and interesting different number cases are rare; yet, I have explored whether the Prior Existence View can deal with different number cases. This is because a moral theory should be able to deal with such cases. I have shown that Prior Existence Utilitarianism can deal with different number choices.

B. Childless George case

In chapter 8 I have shown that Prior Existence Utilitarianism can deal with the Expected Misery Argument. Here, I will briefly discuss a similar case and show that Prior Existence Utilitarianism can deal with that case as well. Here is the case:

Childless George

George passes up an opportunity to conceive a child who would, George has every reason to think, be relentlessly miserable from the start to the end of his

or her short life. By doing so, George imposes some small cost on necessary people.⁷

It might seem that the Prior Existence View cannot account for the judgment that George has done the right thing. After all, the possible child is a contingent being and therefore its welfare does not count. The welfare of necessary beings does count. As the option of not causing the child to exist is somewhat worse for necessary people, it seems that the Prior Existence View would favor having the child. This would be counter-intuitive, because the child's suffering seems clearly worse than having to bear some small cost.

What could proponents of the Prior Existence View say about that case? Well, they could say that George should indeed have conceived the child in order to avoid that cost to necessary beings. However, as soon as the child would have developed into a sentient being, an abortion would have been required in order to prevent the suffering of an actual being. So, there would be no suffering. Still, one might regret that the child is taken into account only once it is an actual sentient being. This makes it necessary to have an abortion, which, intuitively, might seem worse than refraining from conception. It should be noted, however, that painless abortions are bad for the unborn child, if at all, because death deprives the child of the good things it would otherwise have experienced. In case of this child, abortion is not harmful at all, but it prevents the misery the child would otherwise have experienced. It should also be noted that in real life, abortion usually implies more costs for necessary beings than non-conception. Therefore, usually, non-conception is preferable to conception plus abortion.

⁷ Hare (2007) has evoked this case against a different view on the moral status of potential beings, but it might also be evoked against the Prior Existence View. See also Narveson (1967) who drew attention to similar cases.

In case of Childless George, non-conception causes costs to necessary beings as well. Thus, one would need to compare the costs of non-conception to the costs of conception plus abortion, and do what causes the least costs for necessary beings. Neither option causes any costs for the contingent child. After all, the contingent child does not suffer. If it can suffer it is already a necessary child. Furthermore, even if comparisons between existence and non-existence (or death) were possible (which, as I argued in chapter 6, they are not), in this case non-existence or death would be better for the child than the miserable life it would otherwise have had.

Thus, the Prior Existence View can arrive at the judgment that George has done the right thing. In the Childless George Case, Prior Existence Utilitarianism requires minimizing costs to necessary beings, without causing any costs to others.

C. Prior Existence Utilitarianism and transitivity

In chapter 8 and 9, I have explored some implications of Total Utilitarianism and Prior Existence Utilitarianism. Here, I will briefly hint at another possible implication of Prior Existence Utilitarianism, which might deserve further exploration. It is about transitivity.

One of the standard criteria for evaluating moral theories is their consistency; and one aspect of consistency is transitivity. Transitivity means that if the theory considers choice A all things considered morally preferable to B, and B all things considered morally preferable to C, then A must be considered all things considered morally preferable to C as well. It typically counts as a major objection to an ethical theory if it yields intransitive moral judgments. If a theory implies that for some choices A, B and C, A is all things considered morally preferable to B and B is all things considered morally preferable to C, yet C is all things

considered morally preferable to A, then this is usually considered an unacceptable implication and proof of the theory's falsity.

Recently, however, the assumption that a moral theory has to be transitive in order to be rational has been questioned. Following Larry Temkin and Stuart Rachels, Alex Friedman argues that transitivity cannot be reasonably required of a moral theory. According to Friedman: "any ethical theory complex enough to be even minimally plausible allows us to generate intransitive sets of preferences".⁸ Intransitivity can occur if a theory considers several factors morally relevant and if the weight of each factor varies depending on what is compared. Friedman shows that intransitivity occurs in a much wider range of cases than was initially thought. What is more, many common cases of intransitivity are so-called 'hard' cases: they imply that whatever one does there is something better one could have done. As intransitivity, according to Friedman, is unavoidable, transitivity should not be required of moral theories.

Utilitarianism appears to be able to avoid intransitivity, because it accepts only a single relevant factor. The single relevant factor for moral evaluations, according to utilitarianism, is the amount of welfare that an outcome entails. If utilitarianism can indeed avoid intransitive moral judgments, the failure of alternative theories to avoid it might be conceived as an argument in favor of utilitarianism.

However, Friedman argues that the most plausible version of utilitarianism does have intransitive implications in particular cases. It yields intransitive requirements in 'different number cases'. These are cases in which the moral choice that is contemplated affects how many people will live. In such cases utilitarianism in combination with the total view on aggregation proves to be implausible, according to Friedman, because it leads to the Repugnant Conclusion. Accepting the average view as a method of aggregation will not help, according to Friedman,

⁸ Friedman (2009): 277.

because the average view has very implausible implications as well.⁹ A combination of the total view and the average view on aggregation might be able to avoid the counter-intuitive implications of both views, and thus, according to Friedman, this would be the most plausible utilitarian view in different number choices. However, combining two different methods of aggregation means that one accepts two different standards or factors for evaluating outcomes. Therefore, this version of utilitarianism would have intransitive implications. Thus, Friedman concludes that any minimally plausible version of utilitarianism cannot avoid intransitivity.

I will not take a position here as to how problematic intransitive implications are for a moral theory; it is, however, interesting to point out that Prior Existence Utilitarianism can avoid intransitivity. After all, Prior Existence Utilitarianism makes what is morally right dependent on a single factor: welfare. Furthermore, as we have seen, Prior Existence Utilitarianism avoids the Repugnant Conclusion. This is because it does not count the possible welfare of contingent beings in the aggregation of welfare. It can stick to the total view on aggregation and thereby also avoid the counter-intuitive implications of the average view. It focuses solely on the total amount of welfare, and thus does not combine several methods of aggregation.

Prior Existence Utilitarianism can avoid intransitivity at the level of what one is morally required to do, i.e. the deontic level. This is, arguably, the level at which intransitivity would be most disturbing.¹⁰ After all, intransitivity in the sense that whatever one does, there is a better option, implies that there is no rational basis for choice. This would probably lead to skepticism about practical reasoning. Even if intransitivity remains at the axiological level, avoiding it at the deontic level seems to be relevant. So, if transitivity is indeed a reasonable

⁹ I have discussed that in chapter 2, section 6.

¹⁰ Temkin (1996): 209.

requirement of a moral theory, being able to avoid intransitivity at least on the deontic level is a big advantage of Prior Existence Utilitarianism.

Recently, it has been argued that the Person-Affecting Restriction is at odds with transitivity. It has been claimed that either the Person-Affecting Restriction or the requirement of transitivity has to be given up. I have argued that there is a strong link between the Prior Existence View and the Person-Affecting Restriction. The Prior Existence View can be motivated by the Person-Affecting Restriction, and under certain assumptions, is even implied by it. If the Person-Affecting Restriction were indeed bound up with intransitivity, this would hold for Prior Existence Utilitarianism as well.

Temkin brings forward an example that is supposed to show that adherence to the Person-Affecting Restriction might be at odds with transitivity. First, Temkin describes the situation that I have already discussed in chapter 2, section 4.4.

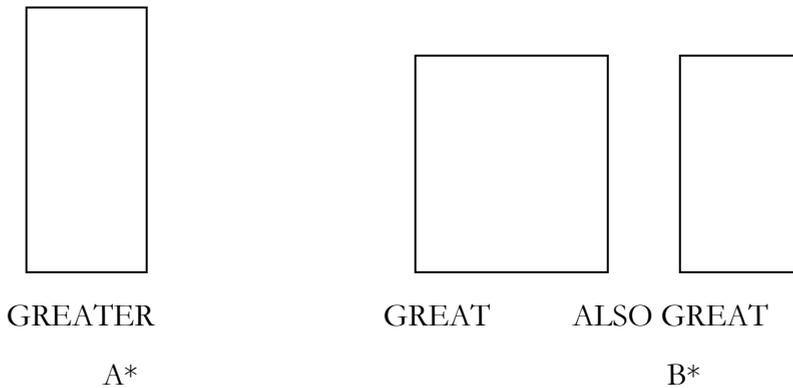


Figure 2: Greater versus Great and Also Great

According to Person-Affecting Utilitarianism, the GREATER population would not be required to have additional children (GREAT) at costs to themselves (ALSO GREAT). Then, Temkin describes that situation further. We are asked to imagine that, by accident, extra

children are born into the situation containing GREATER. The extra children (OK) are barely well off. They do not, in any way, adversely affect GREATER.

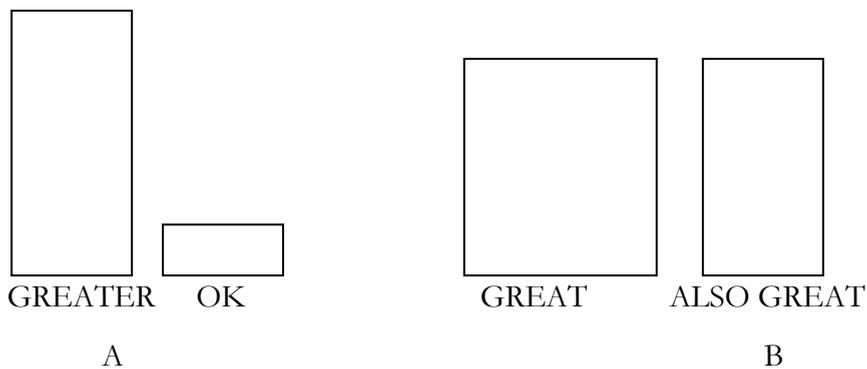


Figure 5: Greater and OK versus Great and Also Great

Appealing to the Person-Affecting Restriction, according to Temkin:

“[...]while A would not be better than GREATER [...], neither would it be worse, since there is no one for whom it is worse. [...] then, A would be better than B, since A is not worse than GREATER which is better than B. Note, this argument assumes transitivity, which we may ultimately reject if the argument succeeds!”¹¹

¹¹ Temkin (1999): 782.

Temkin attempts to show with this example that adhering to the Person Affecting Restriction is a threat to transitivity. After all, those adhering to the Person-Affecting Restriction would agree that GREATER is preferable to GREAT plus ALSO GREAT and that OK plus GREATER is not worse than GREATER. Thus, transitivity would require them to claim that A is not worse than B. But this seems odd, since B seems clearly preferable to A.

I think that Temkin's example does not show that proponents of the Person-Affecting Restriction will have problems with sticking to transitivity. After all, if we are to compare A to B, we must be aware of the fact that both outcomes consist of necessary people. The original population (GREATER) will have children anyway: they will have children in A and in B. So, those children are necessary beings. Only their identity, number and welfare depend on the moral choice that is contemplated, but not their existence. They, whoever they will be and whatever number they will be, will exist anyway. In a choice between necessary beings, the total view on aggregation will be applied. Thus Prior Existence Utilitarianism will imply that B is preferable to A. If A* were an option, and the choice would be between A*, A, and B, then both A and A* would be permissible choices, above B. This is because those would be the best options for necessary beings. Only the original population would consist of necessary beings. But A* proves not to be an option in Temkin's new version of the case. Only A and B are options. A, however, is clearly worse than B. This implication is not counter-intuitive, and there is no need to reject transitivity, or even to contemplate its rejection.

It makes a difference for the discussion of those examples whether one accepts the Narrow or the Wide Person-Affecting Restriction. My concern was to show that assuming the Wide Person-Affecting Restriction, it is possible to deal with those examples in an intuitively appealing way and without being at odds with transitivity.

D. Arrhenius on the relation between the Prior Existence View and the Person-Affecting Restriction

In chapter 5 I have argued that the Prior Existence View needs two assumptions in order to be a coherent utilitarian view on whose welfare to count in the aggregation of welfare. It needs the assumption that in the evaluation of outcomes the focus should be on harms and benefits to sentient beings rather than on the quantity of welfare as such. This view on the evaluation of outcomes is called the Person-Affecting Restriction. In order to be a coherent utilitarian view, the Prior Existence View does also need the assumption that causing a being to exist cannot benefit or harm that being. I have argued that those two assumptions together imply the Prior Existence View.

Contrary to this line of thought, Arrhenius has argued that the Person-Affecting Restriction does not imply the Prior Existence View, or what he calls ‘Necessitarianism’, *even if* it is true that bringing a being into existence cannot harm or benefit that being. Arrhenius attempts to show that, irrespective of whether or not coming into existence can make a being better or worse off, the Person-Affecting Restriction does not imply any version of Necessitarianism, nor Actualism, nor Presentism.¹² (Actualism and Presentism are alternative views on whose welfare to count in the aggregation of welfare that have not been discussed in this thesis.) My own concern here is only with Necessitarianism. Arrhenius defines the Person-Affecting Restriction as follows:

- (a) If outcome A is better than B, then A is better than B for at least one individual.

¹² Arrhenius (forthcoming).

- (b) If outcome A is better than B for someone but worse for no one, then A is better than B.

This is in line with the narrow interpretation of the Person-Affecting Restriction. I made the distinction between the Wide and the Narrow Person-Affecting Restriction in chapter 7.

Arrhenius brings forward only one argument in order to defend the claim that the Person-Affecting Restriction implies neither Actualism, nor Necessitarianism, nor Presentism.

The example he uses focuses on Presentism. We are asked to assume that outcome A consists of the x-people enjoying 10 units of welfare, outcome B consists of the x-people enjoying 5 units of welfare.



Figure 6: Two outcomes with the x-people

Arrhenius explains that a strict presentist would count A and B as equally good, or perhaps as incommensurable, if the x-people were future people. This is correct. As a strict presentist is only concerned with presently living people, it is not evident how he would judge outcomes in which only the welfare of future people differs. Thus, this evaluation by the Presentist would not be in line with the Person-Affecting Restriction. After all, according to the Person-Affecting Restriction, if outcome A is better than B for someone but worse for no one, then A is better than B.

However, does this argument show that the Prior Existence View can be incompatible with the Person-Affecting Restriction? No,

this argument does not apply to the Prior Existence View. If we assume two different outcomes, in which the x-people enjoy 10 or 5 units of welfare respectively, we cannot say that the Prior Existence View would count the outcomes as equally good or as incommensurable *if the x-people were contingent people*. After all, if the x-people exist in all possible outcomes, they are by definition not contingent people. They are necessary people. If it were two different populations, x and y, then the Narrow Person-Affecting Restriction would be indifferent between those outcomes. However, the Wide Person-Affecting Restriction in conjunction with utilitarianism would prefer outcome B, because the y people are better off in B than the x-people are in A. Then, outcome B would be better than outcome A. So, Arrhenius has not shown that the Person-Affecting Restriction can be incompatible with the Prior Existence View, and thus he has not shown that the Person-Affecting Restriction does not imply it.

Interestingly, Arrhenius acknowledges that the Person-Affecting Restriction, in conjunction with the assumption that existence cannot harm or benefit a being, does imply a view called Strict Comparativism. This is another view about the moral status of potential beings. According to Strict Comparativism:

We should completely disregard the welfare of uniquely realizable people, that is, people that only exist in one out of two compared outcomes. We should only count the welfare of non-uniquely realizable people, that is, people who exist in both of the compared outcomes.

Strict Comparativism is not the same as the Prior Existence View. Arrhenius correctly points out that there are differences between both

views.¹³ I wonder whether those differences are such that the Person-Affecting Restriction in conjunction with the assumption that existence cannot harm or benefit a being does imply Strict Comparativism but not the Prior Existence View. In many cases the contingent people are also the uniquely realizable people and the necessary people are the non-uniquely realizable people. However, there might be differences. Arrhenius explains:

A person that is uniquely realizable relative to all pairs of outcomes in the situation of which she is part is also a contingent person, but a contingent person is not necessarily uniquely realizable in respect to all pairs of outcomes in a choice situation since she can exist, for instance, in two out of three outcomes.

Thus, a contingent person might exist in two or more of the possible outcomes, while a uniquely realizable person exists only in one. I do not see how this should make a difference for the issue at hand, i.e. whether the Person-Affecting Restriction in combination with the view that existence cannot benefit or harm a being implies the Prior Existence View. Arrhenius does not point out any relevant difference. Thus, Arrhenius has not convincingly justified his claim that the Person-Affecting Restriction would not imply the Prior Existence View, even if causing a being to exist cannot harm or benefit that being. The implication works as I have indicated in chapter 5.

¹³ Arrhenius, forthcoming.

E. A more technical argument for the Wide Person-Affecting Restriction

In chapter 7, I have defended the Wide Person-Affecting Restriction. This view holds that outcomes should be evaluated in terms of harms and benefits to sentient beings. However, according to the Wide as opposed to the Narrow Person-Affecting Restriction, the focus should not be on particular individuals but rather on sentient beings, whoever they are. I have defended the Wide Person-Affecting Restriction on the basis of Wolf's argument that identity is vague. Wolf's assumption is that the property or set of properties on which identity is based can vary in degree. As explained, another way of expressing that same point is saying that identity is not perfectly fragile. If identity were perfectly fragile, than any minuscule change in the relevant property or set of properties would lead to a different being. Wolf assumes that identity is not perfectly fragile. Caspar Hare shares this assumption and builds an argument on it that leads to the defense of the Wide Person-Affecting Restriction. Here, I will present and discuss this argument.

Consider the case of the 14-year old girl; let's call her 'Charlotte'. Let us assume that Charlotte can either have James now, or Jane eleven years later. Due to the circumstances in which Charlotte finds herself, she knows that James would have a welfare level barely above zero. Jane would have a high level of welfare. I will follow Hare in calling James' life 'miserable' and Jane's life 'happy'. However, it should be kept in mind that unlike the really miserable life in Parfit's Wretched Child case, James' welfare level is still positive. Hare argues that Charlotte should wait. Besides the assumption about identity not being perfectly fragile, Hare departs from a couple of other assumptions, which seem rather uncontroversial, at least for utilitarians. The second assumption is that parents ought to prefer, other things being equal, that any given child they have be better off:

Personal Dominance:

Charlotte ought to prefer a state of affairs in which she has a baby, to an all-other-things-equal state of affairs in which she has the same baby and he or she is worse off.¹⁴

Given this second assumption, the question rises what is a ‘state of affairs’, and how can someone exist in different states of affairs. There are two rival theories about sameness across states of affairs. Hare works out his argument for both theories in turn, in order to leave the choice between them open. His argument works for both theories, though in a different way. I will show how Hare’s argument works for both theories about sameness across states of affairs, first the ‘counterpart theory’, and then the ‘real identity across states of affairs theory’.

So, let us first assume the counterpart theory of identity across states of affairs. The counterpart theory assumes that states of affairs are possible worlds. Ordinary entities, according to this theory, only exist in the actual world. However, relevantly similar entities, which would exist in different possible worlds, are considered counterparts. Assuming this counterpart theory about sameness across states of affairs, the above-mentioned Personal Dominance assumption would be spelled out as follows:

Personal Dominance [according to counterpart theory]:

For worlds W_i , W_k , in which Charlotte-counterparts have counterpart babies B_i , B_k , if B_i is better off than B_k and all other things are appropriately equal, then Charlotte ought to prefer that W_i be actual.¹⁵

¹⁴ Hare (2009): 113.

¹⁵ Hare (2009): 114.

When exactly babies in different worlds are counterparts depends on the view about personal essence that one accepts. For instance, *genetic essentialism* would focus on the genetic make-up of the babies, *essentialism about origins* would focus on the conditions under which the babies come into existence, and *psychological essentialism* would focus on the babies' present and future psychologies. Again, we do not need to decide between those theories.

Assuming, as I explained, that identity is not perfectly fragile, Hare constructs a morphing sequence¹⁶ of intermediary worlds W_1, W_2, \dots, W_n such that

Morphing:

James in W_{James} has a counterpart in W_1 , who has a counterpart in W_2, \dots , who has a counterpart in W_n , who is a counterpart of Jane in W_{Jane} .¹⁷

The third assumption is that wellbeing is fine-grained, i.e. that it can be slightly better or worse. We find that each person in the sequence from James to Jane is slightly better off than the previous person. Hare calls this 'up-slope morphing'.

Up-Slope Morphing:

James in W_{James} is worse off than his counterpart in W_1 , who is worse off than his counterpart in W_2, \dots , who is worse off than her

¹⁶ Morphing animation software, invented in the 1980's works as follows: if you put in two images, it produces a sequence of intermediary images. If you play those images in order, you see one object smoothly transforming into the other, for instance the president into a chimp.

¹⁷ Hare (2009): 115.

counterpart in W_n , who is worse off than her counterpart in W_{Jane} , Jane.¹⁸

It follows from the Personal Dominance Assumption, that Charlotte ought to prefer each subsequent world above its predecessor, because in each subsequent world Charlotte's baby is better off than its counterpart in the previous world.

The final assumption is that preferences are transitive. It follows that Charlotte rationally ought to prefer the world in which she has Jane above the world in which she has James.

Transitivity:

If you prefer that W_1 rather than W_2 be actual, and you prefer that W_2 rather than W_3 be actual, then you prefer that W_1 rather than W_3 be actual.¹⁹

So, Hare has provided an argument based on the assumption that identity is vague and on four rather uncontroversial additional assumptions. According to Hare, claiming that the 14-year-old girl in Parfit's case might as well have the child immediately proves to be morally and/ or rationally deficient.

Let us now see how the argument works for the alternative theory about identity across states of affairs. Let us assume the 'real identity across states of affairs theory'. This theory holds that there can be real identity across states of affairs. Departing from this theory, the Personal Dominance Assumption would read as follows:

¹⁸ Hare (2009): 116.

¹⁹ Hare (2009): 116.

Personal Dominance [assuming real identity across states of affairs]:

For states of affairs S_1 , S_2 , if Charlotte has one baby in both, and her one baby in S_1 is her one baby in S_2 , and that baby is better off in S_1 , then, all other things being appropriately equal, she ought to prefer that S_1 come about.²⁰

Now, it matters whether the baby that Charlotte would have in S_2 is indeed the same baby that she would have in S_1 . This, again, depends on one's view about personal essence. It is plausible to assume that essence is somewhat fragile. That means that big enough differences of the identity determining property or set of properties result in different babies. However, it is also plausible to assume that essence is not perfectly fragile. So, some difference in the identity determining property or set of properties might be compatible with the baby in S_2 still being the same baby as the baby in S_1 . That means: "*All actual people could have been ever so slightly different along any given dimension.*"²¹

Now, Hare introduces a simplified version of the above-mentioned morphing sequence. Imagine that Charlotte, let us say by pressing one out of three buttons, could bring into existence one out of three possible children: either she would have the miserable James, or the happy Jane, or an intermediate option, which is a moderately happy, moderately male, moderately female baby. Imagine that Charlotte actually presses the first button. The states of affairs that she would have brought about by pressing each of the three buttons are:

²⁰ Hare (2009): 117.

²¹ Hare (2009): 118.

Charlotte actually presses the first button. The alternatives for each button are:

1. *B1 is pressed, and miserable James comes into existence.*
2. B2 is pressed and moderately happy James comes into existence.
3. B3 is pressed and a happy child other than James comes into existence.

This is so, because essence is somewhat fragile - child 1 and child 3 are not the same; but essence is not perfectly fragile - child 1 and child 2 are the same. Imagine that, alternatively, Charlotte actually presses the second button:

Charlotte actually presses the second button. The alternatives for each button are:

1. B1 is pressed and miserable Janus comes into being.
2. *B2 is pressed, and moderately happy Janus comes into being.*
3. B3 is pressed and happy Janus comes into being.

As essence is not perfectly fragile, child 1 and child 3 would also be Janus. Finally, assume that Charlotte actually presses the third button:

Charlotte actually presses the third button. The alternatives for each button are:

1. B1 is pressed, and a miserable baby other than Jane comes into being.

2. B2 is pressed and a moderately happy Jane comes into being.
3. *B3 is pressed and happy Jane comes into being.*

The child in B2 would be Jane, because essence is not perfectly fragile. The child in B1 would not be Jane, because essence is somewhat fragile.

Hare proposes that in such cases, the rational thing for Charlotte to do is to press the third button. Why is this so? In order to explain this, Hare introduces the notion of *pair-wise superiority*. Option A is pair-wise superior to option B when the following conditions hold:

- (i) Supposing you actually take A, you would have brought about a less preferable state of affairs by taking B.
- (ii) Supposing you actually take B, you would have brought about a preferable state of affairs by taking A.²²

Hare proposes the following procedure:

Step 1: Choose an option.

Step 2: If there are pair-wise superior options choose one, otherwise keep the option
you have.

Step 3: Continue, until there are no pair-wise superior options.²³

²² Hare (2009): 121. Hare adds a third condition, which is only relevant for evidential decision theorists. This third condition is: (iii) The state of affairs you will bring about, supposing you actually take B, is not preferable to the state of affairs you bring about, supposing you actually take A. For actual decision theorists the first two conditions are sufficient.

²³ Hare (2009): 121.

An option is an attractor if, no matter how you apply this procedure (no matter which option you start with, no matter which pair-wise superior options you choose along the way) you will always get this option.²⁴ Hare assumes that if you are rational, and an option is a stable attractor, you will take it. In Charlotte's case this means:

So, when Charlotte is in a position to push buttons B1, B2, B3, if she has the preferences that she ought, morally to have (in particular her preferences conform to Personal Dominance) and she is rational (in particular she picks stable attractors) then she will press B3.²⁵

Charlotte, if she has the preferences she morally ought to have, and if she is rational, will also choose B3 when the only options are B1 and B3. The following principle about rationality can explain this:

The practical insignificance of irrelevant alternatives:

If, given the option of bringing about S_1, S_2, S_3 , you will willingly bring about S_3 , then given the option of bringing about S_1^*, S_3^* (complete states of affairs relevantly just like S_1 and S_3 , but in which the option of bringing about S_2 is not available to you), you will willingly bring about S_3^* .²⁶

²⁴ Again, Hare adds extra conditions which are required by evidential decision theory: Say that option O is stable when both of the following hold: (i) There is no option K such that, supposing you actually take O, you would bring about a preferable state of affairs by taking K. (ii) There is no option K such that the state of affairs you will bring about, supposing you actually take K, is preferable to the state of affairs you will bring about, supposing you actually take O.

²⁵ Hare (2009): 122.

²⁶ Hare (2009): 122.

Hare concludes that the 14 year-old girl in Parfit's example will prefer to wait, if she is rational and has the preferences she morally ought to have, just because the child she would have if she waited would be better off than the child she would have if she conceived immediately.

At this point, I briefly want to discuss a criticism that has been brought forward against Hare's argument. Graham claims that the second version of Hare's argument fails. Graham argues that the Principle of the Practical Insignificance of Irrelevant Alternatives (in short PIIA) is not applicable to cases like those Hare discusses. According to Graham, "[...] PIIA may fail to hold when the desirability of one of a person's options depends on what other options that person has."²⁷ As an example of such a case, Graham imagines that he has made a solemn promise to give \$ 10 to the Tom De Lay reelection campaign, whenever he has the option of eating a worm. Now, imagine that Graham is facing three options:

S_{Worm} - in which he does eat a worm and does not give \$10 to the campaign

S_{Give} - in which he does not eat a worm and does give \$ 10 to the campaign.

S_{Keep} - in which he does not eat a worm and does not give \$ 10 to the campaign.

Faced with all three options, Graham claims that he should chose S_{Give} . After all, because of the promise, he must give money to the campaign, whenever he faces the option of eating a worm, and he does face the

²⁷ Graham (2007): 5.

option. If, however, only options S_{Give} and S_{Keep} were open to Graham, it is not the case that he rationally ought to actualize S_{Give} . He has no option of eating a worm, and therefore he is not required to give the money away.

The problem that utilitarians are likely to have with this example is that it crucially depends on the non-utilitarian claim that having made a promise is directly morally relevant to what Graham ought to do. According to a utilitarian, faced with those three options Graham ought to do what will maximize utility. This is likely to be S_{Keep} , no matter what the options are.

In order to show that PIIA fails in cases like those Hare discusses, Graham has to point out that in those cases the desirability of an option depends on what other options are available. Graham claims that this is the case. According to Graham, the desirability of S_{James} depends on whether or not S_{Janus} is an option. If it were an option, then choosing S_{James} would wrong someone: it would make some individual worse off than he could have been. If only S_{James} and S_{Jane} are possible, choosing S_{James} does not wrong anyone. Thus, according to Graham, the desirability of S_{James} depends on what the other options are, and therefore PIIA is not applicable in those cases. Here, again, the problem a utilitarian would have with this argument is that it depends on the non-utilitarian notion of ‘wronging someone’. For a utilitarian, whether a particular individual could be better off would be irrelevant, unless her or she accepts the Narrow Person-Affecting Restriction. However, this should not be assumed, as the question at issue is precisely whether a utilitarian needs to do so.

Furthermore, even if Graham could plausibly dismiss the second version of Hare’s argument, his critique is not directed at, nor applicable

to the first version of Hare's argument. Thus, anyway, Graham has not established the failure of Hare's conclusion.²⁸

Hare's argument grounds the choice for the better-off baby in a person-affecting notion of morality. Instead of having to claim that the choice of the better-off baby is better *simpliciter*, Hare can claim that the choice is better 'for the baby', in a wide sense. Hare employs a Wide Person-Affecting Restriction. As a reaction to Hare's argument, it has been noted:

[...] if we really can construct a morphing sequence from pretty much any person or group of persons to any other person or group, the distinction between person-affecting and non person-affecting principles pretty much disappears.²⁹

²⁸ For Hare's critical reply to Graham's argument, see: http://experimentalphilosophy.typepad.com/2nd_annual_online_philoso/2007/05/caspar_hare.html Hare explains: "But notice that caring about things being better for any particular baby you have is not the same as caring about whether you make things better for your particular baby. Charlotte's preferences conform to Dominance, so she prefers a state of affairs in which she has baby Q, with well-being n, and she could have done better for him (WAY better, even), over a state of affairs in which she has baby Q, with well-being n—dx, and she could not have done better for him.

Furthermore, notice that, even if there is this second thing that she cares about (her making things better for her baby) it is not only her caring about this second thing that explains why she presses B3 in the three option case. As I argued, if she is rational and her preferences conform to Dominance and she has no further preferences, then she will press B3 in the three-option case. So she can't say 'I press B3 in the three option case, because S2 has a feature (my failing to make things better for my child) that I find undesirable. S4 does not have this feature, so it's okay for me not to press B3 in the two option case.' She would press B3 in the three-option case even if she did not find the feature undesirable."

²⁹ Norcross: idem.

Indeed, such a Wide Person-Affecting Restriction would yield the same results as an Impersonal View, *as long as only necessary (present or future) beings are involved*. For instance, in the case described by Hare, the mother will have her next baby anyway, so “her next baby” is a necessary being. Only the (genetic) identity and the welfare of the baby will depend on the choice in question, i.e. when to have it. So, the Wide Person-Affecting Restriction yields the same results as the Impersonal View in most cases. It admits of inter-personal aggregation.³⁰ That is in line with a core idea of utilitarianism. However, the Wide Person-Affecting Restriction, as opposed to the Impersonal View, is compatible with the Prior Existence View on whose welfare to count in the aggregation of welfare. Therefore, it will have different implications about the evaluation of outcomes when contingent beings are involved.

³⁰ It is a tricky issue whether a morphing sequence could indeed be constructed from one person to several persons. The case seems to be more complicated than the sequence, constructed by Hare, from male to female. See: the discussion between Hare and Norcross at: http://experimentalphilosophy.typepad.com/2nd_annual_online_philoso/2007/05/caspar_hare.html

BIBLIOGRAPHY

Anglin, W. (1996). In Defense of the Potentiality Principle. Sikora, R. I., & Barry, B. (Eds.). *Obligations to Future Generations*. The White Horse Press.

Anscombe, G.E.M. (1981). Modern Moral Philosophy. *Collected Philosophical Papers, Volume Three, Ethics, Religion and Politics*. Minneapolis: University of Minnesota Press.

Aristotle (1984/ C4BCE) *Metaphysics 1072a*, tr. Ross.

Arrhenius, G. (forthcoming). *Population Ethics*. Oxford University Press.

Arrhenius, G. (2003). The Person-Affecting Restriction, Comparativism, and the Moral Status of Potential People. *Ethical Perspectives* 10

(3-4) 185-195.

Arrhenius, G. (2000). *Future generations: a challenge for moral theory*. Uppsala University.

Balashov, Y. (1999). Zero-Value Physical Quantities. *Synthese* 119, 253-286.

Bekoff, M., & Pierce, J. (2009). *Wild Justice: The Moral Lives of Animals* (1st ed.). University Of Chicago Press.

Benatar, D. (2006). *Better Never to Have Been: The Harm of Coming into Existence*. Oxford University Press.

Bentham, J. (1963). *An introduction to the principles of morals and legislation*; Hafner Pub. Co.

Birnbacher, D. (2006). Utilitarismus. M. Düwell, C. Hübenthal, M. Werner (eds.) *Handbuch Ethik*, Stuttgart: Metzler.

Birnbacher, D. (2005). Philosophie des Glücks. *E-journal Philosophie der Psychologie*, www.jp.philo.at/texte/BirnbacherD1.pdf

Bognar, G. (2010). Authentic Happiness. *Utilitas*, 22(03), 272-284.
doi:10.1017/S0953820810000191

Bordes, E.C.de (2010). Dieren in het Geding. Een Juridisch-Historische Analyse van het Verbod op Dierenmishandeling. (Animals in Jeopardy: A Legal History of the Prohibition of Maltreatment of Animals in the Netherlands – with a summary in English).
Doctoral thesis, Utrecht University.

Bracke, M.B.M., Hopster, H. (2006). Assessing the Importance of

- Natural Behavior for Animal Welfare. *Journal of Agricultural and Environmental Ethics* 19, 77-89.
- Bradley, B. (2009). *Well-being and death*. Oxford University Press.
- Bradley, B. (2009). Analyzing Harm. Unpublished Manuscript.
- Bradley, B. (2006). Against Satisficing Consequentialism. *Utilitas*, 18(02), 97-108. doi:10.1017/S0953820806001877
- Brink, D. O. (1989). *Moral realism and the foundations of ethics*. Cambridge University Press.
- Broome, J. (1999). *Ethics out of economics*. Cambridge University Press.
- Broome, J. (2004). *Weighing lives*. Oxford University Press.
- Brosnan, S. F., & de Waal, F. B. M. (2003). Monkeys reject unequal pay. *Nature*, 425(6955), 297-299. doi:10.1038/nature01963
- Brownlee, K. (2010). Moral Aspirations and Ideals. *Utilitas* (22), 3.
- Buchanan, A., Brock, D. W., Daniels, N., & Wikler, D. (2001). *From chance to choice: genetics and justice*. Cambridge University Press.
- Bykvist, K. (2006). The Benefits of Coming into Existence. *Philosophical Studies*, 135(3), 335-362. doi:10.1007/s11098-005-3982-x
- Carlson, E. (2003). On Some Recent Examples of Intransitive Betterness. K. Segerberg and R. Sliwinski (eds.) *Logic, Law, Morality*, Uppsala Department of Philosophy.
- Coalition Accord of the Dutch Government, 7 February 2007.
- Cochrane, A. (2007). Animal Rights and Animal Experiments: An

- Interest-Based Approach. *Res Publica*, 13(3), 293-318.
doi:10.1007/s11158-007-9037-8
- Cochrane, A. (2009). Ownership and Justice for Animals. *Utilitas*, 21(04), 424-442. doi:10.1017/S0953820809990203
- Cohnitz, D., Häggqvist, S. (2009). The Role of Intuitions in Philosophy. *Studia Philosophica Estonica* 2, 1-14.
- Crisp, R. (2008). Well-Being. Stanford Encyclopedia of Philosophy.
- Crisp, R. (2006). Hedonism Reconsidered. *Philosophy and Phenomenological Research*, 73(3), 619-645. doi:10.1111/j.1933-1592.2006.tb00551.x
- DeGrazia, D. (2007). Must We Have Full Moral Status Throughout Our Existence? A Reply to Alfonso Gomez-Lobo. *Kennedy Institute of Ethics Journal*, 17(4), 297-310.
- DeGrazia, D. (2003). Identity, Killing, and the Boundaries of Our Existence. *Philosophy and Public Affairs*, 31(4), 413-442.
doi:10.1111/j.1088-4963.2003.00413.x
- Diederer, P. (2003). Burger, Laat de Consument met Rust. H. Dagevos, L. Sterrenberg (eds.) *Burgers en Consumenten: Tussen Tweedeling en Twee-eenheid*. Wageningen Academic Publishers.
- Dorsey, D. (2010). Consequentialism, Metaphysical Realism, and the Argument from Cluelessness. Unpublished paper.
- Earlenbaugh, J., & Molyneux, B. (2009). Intuitions are inclinations to believe. *Philosophical Studies*, 145(1), 89-109. doi:10.1007/s11098-009-9388-4

- Faucitano, L., & Schaefer, A. L. (2008). *Welfare of pigs from birth to slaughter*. Wageningen Academic Publisher.
- Fehige, C., & Wessels, U. (1998). *Preferences*. Berlin: Walter de Gruyter.
- Feinberg, J. (ed.) (1984). *The Problem of abortion*. 2nd ed. Belmont, CA: Wadsworth Publishing Company.
- Foer, J.S. (2009). *Eating Animals*. New York: Little, Brown and Company.
- Francione, G.L. (2000). *Introduction to Animal Rights. Your Child or the Dog*. Philadelphia: Temple University Press.
- Francione, G. L. (1996). *Rain without thunder: the ideology of the animal rights movement*. Philadelphia: Temple University Press.
- Frey, R. G., & Morris, C. W. (1993). *Value, welfare, and morality*. Cambridge University Press.
- Friedman, A. (2009). Intransitive Ethics. *Journal of Moral Philosophy* 6, 277-297.
- Gauthier, D. (1962). *Practical Reasoning: The Structure and Foundations of Prudential and Moral Arguments and Their Exemplification in Discourse*. Oxford University Press.
- Gesang, B. (2000). Der Nutzenbegriff des Utilitarismus. *Erkenntnis* 52, 373-401.
- Glover, J. (1990). *Causing death and saving lives. The Moral Problems of Abortion, Infanticide, Suicide, Euthanasia, Capital Punishment, War and Other Life-or-death Choices*. Penguin.

- Greene, J.D., Sommerville, R.B., Nystrom, L.E., Darley, J.M., & Cohen, J.D. (2001). An fMRI Investigation of Emotional Engagement in Moral Judgment”, *Science* 293, 2105-2108.
- Gren, J. (2004). *Applying utilitarianism: the problem of practical action-guidance*. Acta Universitatis Gothoburgensis.
- Gruen, L. (2003). The Moral Status of Animals. Stanford Encyclopedia of Philosophy,
<http://plato.stanford.edu/entries/moral-animal/>.
- Haidt, J. (2001). The Emotional Dog and Its Rational Tail: A Social Intuitionist Approach to
Moral Judgment. *Psychological Review* 108, 814-834.
- Haidt, J., Rozin, P., McCauley, C. R., & Imada, S. (1997). Body, Psyche, and Culture: The
Relationship between Disgust and Morality. *Psychology and Developing Societies*, 9, 107-
131.
- Hare, C. (2009). The ethics of morphing. *Philosophical Studies*, 145(1), 111-130. doi:10.1007/s11098-009-9385-7
- Hare, C. (2008). Obligation and Regret When There is No Fact of the Matter About What Would Have Happened if You had not Done What You Did. Unpublished Paper.

- Hare, C. (2007). Voices From Another World: Must We Respect The Interests Of People Who Do Not, And Will Never Exist? *Ethics* 117 (3).
- Hare, R. (1993). Why I am Only a Demi-Vegetarian. D. Jamieson (ed.) *Singer and His Critics*,
Oxford: Blackwell Publishers.
- Hare, R. (1981). *Moral Thinking. It's Levels, Method and Point*. Oxford University Press.
- Harman, E. (2004). Can We Harm and Benefit In Creating? *Philosophical Perspectives*, 18(1),
89-113. doi:10.1111/j.1520-8583.2004.00022.x
- Harman, E. (2009). Harming as Causing Harm. M. Roberts & D. Wasserman (Eds.), *Harming Future Persons. Ethics, genethics and the Nonidentity Problem*. Dordrecht: Springer.
- Hart, H.L.A., Honoré, T. (1985). *Causation in the Law, Second Edition*. Oxford: Clarendon Press.
- Haybron, D. M. (2008). *The pursuit of unhappiness: the elusive psychology of well-being*. Oxford University Press.
- Hentzepeter (2008). Kalverliefde. *Smaakmakend*, voorjaar 2008, 23-26.
- Heyd, D. (1994). *Genethics: moral issues in the creation of people*. University of California Press.
- Heyd, D. (n.d.). The Intractability of the Non-Identity Problem. In Roberts, Melinda & Wasserman, David (Eds.), *Harming Future*

Persons. Ethics, genetics and the Nonidentity Problem. Dordrecht: Springer.

- Holtug, N. (2010). Killing and the Time-relative Interest Account. *The Journal of Ethics*. doi:10.1007/s10892-010-9087-6
- Holtug, N. (2008). Good for whom? *Theoria*, 69(1-2), 4-20.
doi:10.1111/j.1755-2567.2003.tb00750.x
- Holtug, N. (2004). Person-Affecting Moralities. J. Ryberg & T. Tännsjö (eds.) *The Repugnant Conclusion*, Dordrecht: Kluwer Academic Publishers.
- Holtug, N. (2001). On the Value of Coming into Existence. *The Journal of Ethics*, 5(4), 361-384-384.
- Hooker, B. (1996). Ross-Style Pluralism Versus Rule-Consequentialism. *Mind*, New Series, 105(420), 531-552.
- Howard, C. J., Kumar, A., Mitloehner, F., Stackhouse, K., Green, P. G., Flocchini, R. G., & Kleman, M. J. (2010). Direct Measurements of the Ozone Formation Potential from Livestock and Poultry Waste Emissions. *Environmental Science & Technology*, 44(7), 2292-2298. doi:10.1021/es901916b
- Hübenthal, C. (2006). Teleologische Ansätze. M. Düwell, C. Hübenthal, M. Werner (eds.) *Handbuch Ethik*, Stuttgart: Metzler.
- Hurka, T. (1990). Two kinds of satisficing. *Philosophical Studies*, 59(1), 107-111. doi:10.1007/BF00368395
- Hurka, T. (1983). Value and Population Size. *Ethics*, 93(3), 496-507.

- Jamieson, D., Elliot, R. (2009). Progressive Consequentialism. *Philosophical Perspectives* 23 (1), 241-251.
- Jamieson, D. (2007). When Utilitarians Should Be Virtue Theorists. *Utilitas* 19 (2).
- Jamieson, D. (1984). Utilitarianism and the Morality of Killing. *Philosophical Studies* 54 (2).
- Jamieson, D. (ed.). (1999). *Singer and his critics*. Oxford: Blackwell.
- Jamieson, D. (1991). Method and Moral Theory. P. Singer (ed.) *Companion to Ethics*. Oxford: Blackwell.
- Kagan, S. (2000). Evaluative Focal Points. B. Hooker, E. Mason, D. Miller (eds.) *Morality, Rules, and Consequences*, Edinburgh: Rowman & Littlefield.
- Kagan, S. (1998). *Normative ethics*. Westview Press.
- Kaldewaij, F. (2008). Animals and the Harm of Death. S.J. Armstrong & R.G. Botzler, (eds.), *The Animal Ethics Reader, 2nd edition*. London: Routledge.
- Kavka, G. S. (1982). The Paradox of Future Individuals. *Philosophy and Public Affairs*, 11(2), 93-112.
- Korthals, M. (2002). *Voor het Eten. Filosofie en Ethiek van Voeding*. Meppel: Boom.
- Kittay, E. F. (2008). At the Margins of Moral Personhood. *Journal of Bioethical Inquiry*, 5(2-3), 137-156. doi:10.1007/s11673-008-9102-9

- Kymlicka, W. (2002). *Contemporary Political Philosophy. An Introduction, second edition*, Oxford University Press.
- Lang, G. (2008). Consequentialism, Cluelessness, and Indifference. *Journal of Value Inquiry* 42, 477-485.
- Leake, J. (2005). McDonald's Chickens Sent to Gas Chamber. *The Sunday Times*, January 9, 2005.
- Lenman, J. (2000). Consequentialism and Cluelessness. *Philosophy and Public Affairs* 29, 342-370. doi: 10.1111/j.1088-4963.2000.00342.x
- Liao, S.M. (2007). Time-Relative Interests and Abortion. *Journal of Moral Philosophy* 4 (2), 242-256.
- Louise, J. (2006). Right Motive, Wrong Action: Direct Consequentialism and Evaluative Conflict. *Ethical Theory and Moral Practice*, 9(1), 65-85.
- Lund, V., & Olsson, I. (2006). Animal Agriculture: Symbiosis, Culture, or Ethical Conflict? *Journal of Agricultural and Environmental Ethics*, 19(1), 47-56.
- Luper, S (2009). Review of Ben Bradley: Well-being and Death. *Notre Dame Philosophical Reviews* (<http://ndpr.nd.edu/review.cfm?id=16606>).
- Luper, S. (2007). Mortal Harm. *The Philosophical Quarterly*, 57(227), 239-251. doi:10.1111/j.1467-9213.2007.482.x
- Marquis, D. (2009). Singer on Abortion and Infanticide. J.A. Schaler, (ed.) *Singer under Fire. The Moral Inconoclast Faces His Critics*, Illinois:

Carus Publishing Company.

- Matheny, G., & Chan, K. (2005). Human Diets and Animal Welfare: the Illogic of the Larder. *Journal of Agricultural and Environmental Ethics*, 18(6), 579-594-594.
- McDowell, J. (1979). Virtue and Reason. *Monist* 62, 331-350.
- Meijs, J. (2007). Voorwoord. M.Ruis & J.B. Pinxterhuis (eds.) *Verantwoorde en Cummuniceerbare Argumenten bij Biologische Producten: Dierenwelzijn*, Wageningen: Animal Science Group.
- Mill, J. S., & Bentham, J. (1987). *Utilitarianism and Other Essays* (Reprint.). Penguin Classics.
- Mill, J.S. (1869/1997). *The Subjection of Women*. Dover: Thrift Editions. ISBN 0-486-29601-6.
- Nagel, T. (1970). *The Possibility of Altruism*. Oxford University Press.
- Narveson, J. (1996). Future People and Us. R.I. Sikora & B. Brian (eds.), *Obligations to Future Generations*, Philadelphia: Temple University Press.
- Narveson, J. (1973). Moral problems of population. *The Monist*, 57(1), 62-86.
- Narveson, J. (1967). Utilitarianism and New Generations. *Mind*, 301, 62-72. doi:10.1093/mind/LXXVI.301.62
- Norcross, A. (2008). Off Her Trolley? Frances Kamm and the Metaphysics of Morality. *Utilitas*, 20(01), 65-80. doi:10.1017/S0953820807002919

- Norcross, A. (2008). Two Dogmas of Deontology: Aggregation, Rights, and the Separateness of Persons. *Social Philosophy and Policy Foundation*. doi: 10.1017/S0265052509090049.
- Norcross, A. (2004). Puppies, Pigs, and People: Eating Meat and Marginal Cases. *Philosophical Perspectives* 18 (1), 229-245.
- Norcross, A. (1999). Intransitivity and the Person-Affecting Principle. *Philosophy and Phenomenological Research* 59 (3), 767-776.
- Norcross, A. (1997). Comparing Harms: Headaches and Human Lives. *Philosophy and Public Affairs*. doi:10.1111/j.1088 4963.1997.tb00079.x
- Nozick, R. (1974). *Anarchy, State, and Utopia*. New York: Basic Books.
- Nussbaum, M. C. (2006). *Frontiers of Justice. Disability, Nationality, Species Membership*. Cambridge MA: The Belknap Press of Harvard University.
- Ord, T. (2008). How to be a Consequentialist about Everything??. Paper presented at the Tenth Conference of the International Society for Utilitarian Studies at University of California.
- Parfit, D. (1984). *Reasons and Persons*. Clarendon Press, Oxford.
- Parfit, D. (1981). Future Generations: Further Problems. *Philosophy and Public Affairs* 11 (2).
- Persson, I. (1995). Peter Singer on Why Persons Are Irreplaceable. *Utilitas*, 7(01), 55-66. doi:10.1017/S0953820800001849
- Pettit, P. (1991). Consequentialism. P. Singer (ed.) *Companion to Ethics*.

Oxford: Blackwell.

- Pimentel, D., & Pimentel, M. (2003). Sustainability of meat-based and plant-based diets and the environment. *The American Journal of Clinical Nutrition*, 78(3), 660S -663S.
- Platform Biologica, EKO-Monitor 2002, 13, 3.
- Pluhar, E. (1995). *Beyond Prejudice. The moral significance of human and non-human animals*. London: Duke University Press.
- Rachels, J. (1999). When Philosophers Shot from the Hip. H. Kuhse & P. Singer (eds.) *Bioethics: An Anthology*. Oxford: Blackwell Publishers.
- Rawls, J. (1971). *A Theory of Justice*. Cambridge: Harvard University Press.
- Regan, T. (2004). *The Case for Animal Rights*, Berkeley: University of California Press.
- Reijnders, L., & Soret, S. (2003). Quantification of the environmental impact of different dietary protein choices. *The American Journal of Clinical Nutrition*, 78(3), 664S -668S.
- Report of the German-Dutch Debate on the Future of Agriculture (2002).
- Roberts, M. A. (1998). *Child versus childmaker: future persons and present duties in ethics and the law*. Rowman & Littlefield.
- Roberts, M. A. (2009). The Nonidentity Problem and the Two Envelope Problem: When is One Act Better for a Person than Another? M.A. Roberts & D.T. Wasserman (eds.) *Harming Future Persons: Ethics, Genetics and the Nonidentity Problem*.

- Roberts, M.A. (2009b). The Non-Identity Problem. *Stanford Encyclopedia of Philosophy*.
- Roeser, S. (2002). *Ethical Intuitions and Emotions. A Philosophical Study*. Dissertation at the Free University of Amsterdam.
- Rogers, J. (2010). In Defense of a Version of Satisficing Consequentialism. *Utilitas*, 22(02), 198-221.
doi:10.1017/S0953820810000099
- Ruis, M., Pinxterhuis, J.B. (2007). *Verantwoorde en communiceerbare argumenten bij biologische producten: dierenwelzijn*, Wageningen: Animal Science Group.
- Rollin, B.E. (1995). *Farm Animal Welfare. Social, Bio-Ethical and Research Issues*. Ames: Iowa State University.
- Ryberg, J., & Tännsjö, T. (2004). *The repugnant conclusion: essays on population ethics*. Dordrecht: Kluwer Academic Publishers.
- Salt, H. (1914). "Logic of the Larder". www.henrysalt.co.uk.
- Sandberg, J., & Juth, N. (2010). Ethics and Intuitions: A Reply to Singer. *The Journal of Ethics*. doi:10.1007/s10892-010-9088-5
- Sandoe, P., Christiansen, S.B. (2008). *The Ethics of Animal Use*. Oxford: Blackwell.
- Sapontzis, S. (1987). *Morals, Reason and Animals*. Philadelphia: Temple University Press.
- Schnall, S., Heidt, J., Clore, G.L., Jordan, A.H. (2008). Disgust as Embodied Moral Judgment. *PSPB* 34 (8), 1096-1109.

- Scruton, R. (2006). Eating Our Friends. *Right Reason* (May 26, 2006).
(Reprinted at
<http://catholiceducation.org/articles/environment/en0011.htm>.)
- Sidgwick, H. (1981/ 1907). *The Methods of Ethics*. Indianapolis: Hackett Publishing.
- Shiffrin, S. (1999). Wrongful Life, Procreative Responsibility, and the Significance of Harm.
Legal Theory 5, 117-148.
- Sikora, R. I. (1996) Is it Wrong to Prevent the Existence of Future Generations? I. Sikora & B. Barry (eds.). *Obligations to Future Generations*. The White Horse Press.
- Singer, P. (forthcoming 2011). *Practical Ethics* (3rd edition). New York: Cambridge University Press.
- Singer, P. (2009). *The Life You Can Save. Acting Now to End World Poverty*. New York: Random House.
- Singer, P. (2009). Reply to Marquis. J.A. Schaler (ed.) *Singer under Fire. The Moral Inconoclast Faces His Critics*. Illinois: Carus Publishing Company.
- Singer, P. (2005). Ethics and Intuitions. *Journal of Ethics* 9, 331-352.
- Singer, P (1998). Possible Preferences. Fehige, C., Wessels, U. (eds.) *Preferences*. Berlin: De Gruyter.
- Singer, P. (1995). *Animal Liberation* (2nd edition). London: Pimlico.

- Singer, P. (1993). *Practical Ethics* (2nd edition). New York: Cambridge University Press.
- Sinnott-Armstrong, W. (2006). *Consequentialism*. Stanford Encyclopedia of Philosophy.
- Smolkin, D. (1999). Toward A Rights-Based Solution to the Non-Identity Problem. *Journal of Social Philosophy*, 30(1), 194-208. doi:10.1111/0047-2786.t01-1-00012
- Sobel, D. (2009). Subjectivism and Idealization. *Ethics*, 119, 336-352.
- Stassen, E. (2006). Liever een Gangbare Koe. www.wb-online.nl.
- Sterrenberg, L et al (2001). *Hoe oordelen wij over de veehouderij?* Den Haag: Rathenau Instituut.
- Stocker, M. (1976). The Schizophrenia of Modern Ethical Theories. *Journal of Philosophy* 63,453-466.
- Sturma, D. (2006). Person. M. Düwell, C. Hübenal, M. Werner (eds.) *Handbuch Ethik*. Stuttgart: Metzler.
- Sumner, L.W. (1996). *Welfare, Happiness and Ethics*, Oxford: Clarendon Press.
- Sumner, W. (1996). Classical Utilitarianism and the Population Optimum. R.I. Sikora & B. Brian (eds.) *Obligations to Future Generations*, Philadelphia: Temple University Press.
- Swart, J.A.A. (2006). Dieren in Context. Presentation at the Dutch Society for Bioethics (NVBe), 2 November 2006, Utrecht.
- Tännsjö, T. (1998). *Hedonistic Utilitarianism*. Edinburgh University Press.

- Temkin, L.S. (2000). Equality, Priority and the Levelling Down Objection. M. Clayton & A. Williams, (eds.) *The Ideal of Equality*. New York: Palgrave.
- William. (eds.) *The Ideal of Equality*. London/ New York, Macmillan and St. Martin's Press.
- Temkin, L.S. (1999). Intransitivity and the Person-Affecting Principle: A Respons. *Philosophy and Phenomenological Research* 59 (3), 777-784.
- Temkin, L. S. (1999b). Intransitivity and the Person-Affecting Principle: A Response. *Philosophy and Phenomenological Research*, 59(3), 777-784. doi:10.2307/2653796
- Temkin, L.S. (1993). *Inequality*. Oxford University Press.
- Timmons, M (2002). *Moral Theory. An Introduction*. New York: Rowan & Littlefield.
- Turry, J. (2005). You Can't Get Away With Murder That Easily. A Response to Tim Mulgan. *International Journal of Philosophical Studies* 13 (4), 489-492.
- Valentyne, P. (2006). Against Maximizing Act-Consequentialism. J. Dreier (ed.) *Contemporary Debates in Moral Theory*, Malden MA: Blackwell Publishing.
- Velleman, J.D. (2008). Persons in Prospect. *Philosophy & Public Affairs*, 36, 221-288.

- Verburg, G. (2007). *Nota Dierenwelzijn*. The Hague: Ministry of Agriculture.
- Visak, T. Meijboom, F.L.B., Kalis, A., Brom, F.W.A. (2004). *Verantwoord Consumenten: De Dubbele Moraal als Vertrekpunt voor een Discussie over Normative Vooronderstellingen en Verantwoordelijkheid*, verkennend essay, NOW Programma Ethiek, Onderzoek en Bestuur.
- Warren, M.A. (1996). Do Potential People Have Moral Rights? R.I. Sikora & B. Brian (eds.) *Obligations to Future Generations*, Philadelphia: Temple University Press.
- Wasserman, D.T. (2009). Harms to Future People and Procreative Intentions. M.A. Roberts & D.T. Wasserman *Harming Future Persons. Ethics, Genethics and the Nonidentity Problem*. Dordrecht: Springer.
- Weinberg, R. (2007). Identifying and Dissolving the Non-Identity Problem. *Philosophical Studies*, 137(1), 3-18. doi:10.1007/s11098-007-9168-y
- Williams, B. (1985). *Ethics and the Limits of Philosophy*. Cambridge, Mass.: Harvard University Press.
- Williams, B. (1973). *Problems of the Self*. Cambridge University Press.
- Wolf, C. (2009). Do Future Persons Presently Have Alternate Possible Identities? M.A. Roberts & D.T. Wasserman *Harming Future Persons. Ethics, Genethics and the Nonidentity Problem*. Dordrecht: Springer.
- Woodward, J. (1986). The Non-Identity Problem. *Ethics*, 96(4), 804-831.

Wrigley, A. (2006). Genetic Selection and Modal Harms. *The Monist* 89 (4).

INDEX

A

Actualism · 129, 259, 260
Aggregation · 40
Animal experimentation · 7
Animal husbandry · 7, 10
Argument from Composition · 149
Argument from Disparity · 154
Argument from Parity · 149
Argument from Unification · 150
Arrhenius · 47, 127, 131, 136, 201, 204,
250, 259, 260, 261, 262

B

Balashov · 148
Bentham · 14, 32, 60, 194
Bradley · 31, 38, 76, 79, 81, 82, 83, 84, 85,
86, 149, 150, 151, 152, 153
Brink · 40

C

Character · 209, 211
Coming into existence · 23, 158
Cows · 99

D

De dicto · 166, 168
De re · 168
Different-number choices · 245

E

Egg production · 7
Evaluative focal points · 44, 45
Existence · 4, 22, 23, 24, 46, 119, 126,
129, 131, 133, 136, 142, 158, 161, 164,
179, 180, 185, 186, 187, 188, 189, 190,
195, 202, 203, 204, 205, 206, 207, 210,
211, 214, 215, 216, 217, 219, 220, 223,
224, 226, 227, 229, 230, 231, 233, 234,

237, 238, 239, 240, 241, 242, 243, 244,
245, 249, 250, 251, 252, 253, 255, 256,
258, 259, 261, 262
Expected Misery Argument · 188, 202,
240, 251

G

Genuine Modal Realism · 170

H

Hedonism · 37
Heyd · 145, 146, 177
Holtug · 46, 87, 140, 141, 142, 143, 147

I

Identity · 172, 173, 249
Impersonal View · 46, 48, 49, 62, 88, 159,
164, 167, 183, 184, 186, 198, 238, 240,
241, 274

K

Kant · 16
Kantianism · 28
Killing · 21, 63, 64, 67, 78, 81, 93, 105,
106, 120, 122, 197, 222, 234, 242
Kripke · 169

L

Lewis · 170
Logic of the Larder · 24, 94, 95, 126, 219,
220, 221, 222, 223, 224, 225, 242
Luper · 150, 151, 152, 153, 154

M

Meat · 6, 221
Meat production · 6
Milk · 6
Milk and dairy production · 6
Mill · 32, 51, 194
Moral theory · 190

N

Narveson · 144, 146, 252
Necessitarianism · 131, 259, 260
No Properties of the Non-Existent
 Objection · 142
Non-existence · 142
Non-Identity Problem · 23, 160, 161, 162,
163, 164, 165, 169, 172, 174, 175, 176,
178, 179, 180, 181, 184, 185, 238

P

Parfit · 48, 142, 156, 157, 161, 162, 163,
166, 169, 197, 198, 203, 204, 239, 246,
263, 266, 271
Perfectionist value · 37
Person-Affecting Restriction · 23, 46, 47,
48, 49, 134, 135, 159, 160, 161, 163,
164, 165, 166, 167, 168, 175, 176, 177,
178, 179, 180, 181, 182, 183, 184, 185,
238, 240, 248, 249, 256, 257, 258, 259,
260, 261, 262, 263, 272, 273, 274
Person-Affecting Utilitarianism · 48, 167,
256
Pigs · 9, 13
Pluhar · 15, 59, 93, 95, 103, 109, 131, 132,
210, 211, 212
Practical Ethics · 1, 5, 106
Presentism · 259, 260
Presentists · 128

Prior Existence View · 5, 22, 126, 129,
130, 131, 132, 133, 134, 135, 136, 137,
138, 158, 159, 161, 164, 177, 179, 180,
185, 186, 187, 189, 202, 203, 204, 206,
207, 210, 215, 223, 225, 233, 239, 245,
249, 250, 251, 252, 253, 256, 259, 260,
261, 262, 274

R

Replaceability Argument · 5, 6, 7, 21, 22,
58, 59, 68, 71, 89, 90, 91, 92, 93, 94,
95, 96, 97, 100, 102, 103, 104, 105,
107, 109, 110, 111, 112, 113, 114, 115,
118, 119, 120, 121, 122, 123, 124, 125,
127, 129, 130, 131, 136, 164, 177, 179,
180, 187, 189, 220, 222, 223, 224,
225, 227, 228, 229, 231, 235, 237, 250
Repugnant Conclusion · 24, 184, 188, 196,
197, 198, 199, 200, 201, 202, 217, 239,
240, 249, 254, 255
Rule · 45, 208

S

Salt · 131, 221
Singer · 1, 4, 5, 22, 51, 59, 60, 68, 74, 75,
76, 90, 91, 94, 103, 104, 105, 106, 107,
108, 109, 110, 111, 112, 113, 114, 115,
116, 117, 118, 119, 120, 121, 124, 130,
131, 163, 189, 191, 192, 194, 195, 204,
214, 215, 216, 224, 225, 226, 228, 229,
235, 236, 237, 242, 243
Skepticism · 56
Sport fishing · 7
Sport hunting · 7

T

Teleological · 30
Time-Relative Interest Account · 82, 83,
84, 85, 86, 87, 88, 117
Total View · 5, 22, 91, 124, 126, 129, 130,
132, 133, 134, 136, 164, 187, 189, 196,
197, 198, 199, 200, 201, 202, 223, 225,
233, 235, 237, 241, 249
Two Women · 162

U

Unequal interests · 20
Utilitarianism · 3, 4, 14, 15, 21, 23, 24, 27,
28, 29, 30, 31, 32, 33, 34, 40, 46, 48,
49, 51, 57, 59, 68, 69, 70, 93, 161, 167,
178, 185, 186, 187, 188, 189, 194, 195,
201, 205, 206, 208, 210, 211, 213, 214,
216, 217, 219, 220, 223, 224, 226, 227,
229, 230, 231, 233, 234, 235, 237, 238,
239, 240, 241, 242, 243, 244, 245, 250,
251, 253, 254, 255, 256, 258
Utility · 31, 144

V

Value of Existence View · 142
Vegan agriculture · 230, 231

W

Welfare · 13, 30, 36, 116, 134, 155
Wolf · 172, 173, 174, 175, 176, 247, 248,
263
Wrigley · 170, 171, 172, 173

SAMENVATTING

1. Inleiding

Er lijkt een brede consensus te bestaan over de morele ontoereikendheid van gangbare vormen van intensieve veehouderij. Hoewel de daad om allerlei redenen achterblijft bij het woord, geven burgers en consumenten desgevraagd aan dat de veehouderij diervriendelijker zou moeten zijn. Er zijn allerlei maatschappelijke bewegingen en initiatieven die streven naar diervriendelijke veehouderij, en dat vindt ook zijn weerklank in de politiek. Het realiseren van diervriendelijke veehouderij wordt voorgesteld als een moreel ideaal. Het achterliggende idee is dat wij niet zomaar alles met dieren mogen doen, en dat wij ons moeten bekommeren om het welzijn van dieren. Zo dient onnodig leed te worden voorkomen in onze omgang met dieren.

Er zit een spanning in dit breed gedragen morele ideaal. Aan de ene kant doen dieren er kennelijk toe, in moreel opzicht, en dienen wij rekening te houden met hun welzijn. Aan de andere kant is diervriendelijke veehouderij nog steeds een vorm van veehouderij, en houdt dus in dat dieren massaal en routinematig worden gedood om gebruik te maken van hun producten. Aangezien het gebruik van

dierlijke producten in ieder geval in onze streken niet nodig is om gezond te leven lijkt het gebruik van dierlijke producten vooral te zijn ingegeven door overwegingen die neer komen op gemak en plezier. Is het routinematige doden van dieren voor dit soort zaken te verenigen met het idee dat dieren er moreel gezien toe doen, en dat wij ons moeten bekommeren om hun welzijn? Kan het ideaal van diervriendelijke veehouderij ondersteund worden door een morele theorie?

In dit onderzoek zal ik verkennen hoe de utilistische morele theorie zich verhoudt tot het ideaal van diervriendelijke veehouderij. Het utilisme, één van de klassieke morele theorieën en de theorie die traditioneel het meest rekening heeft gehouden met het welzijn van dieren, lijkt op het eerste gezicht een geschikte theorie om het ideaal van diervriendelijke veehouderij te verdedigen. Volgens het utilisme geldt namelijk die handeling als moreel juist, die het maximale welzijn oplevert voor alle betrokkenen. Dierenwelzijn is daarbij net zo belangrijk als het welzijn van mensen. Tegelijkertijd is het utilisme niet principieel tegen het doden van dieren. Volgens het utilisme is geen enkele handeling als zodanig moreel juist of verkeerd. Alles hangt ervan af welke gevolgen de handeling heeft in een specifieke situatie. Dus ook het doden zou geoorloofd zijn, als het de handeling is die het welzijn van alle betrokkenen maximeert. Reden genoeg om nader te onderzoeken of het utilisme het ideaal van diervriendelijke veehouderij kan onderbouwen.

2. Utilisme en dierethiek

Het utilisme is één van de twee belangrijkste morele theorieën van de Verlichting. Als zodanig heeft het utilisme tot doel om ons op systematische en samenhangende wijze te vertellen wat moreel juist en verkeerd is, en waarom dat zo is. Het utilisme beoordeelt handelingen enkel op basis van de gevolgen en houdt in dat alleen die handeling moreel juist is die van alle mogelijke handelingen die men kon doen de beste gevolgen oplevert. Bij het beoordelen van wat de beste gevolgen

zijn wordt alleen gekeken naar welzijn. De juiste handeling is dan die handeling die het meeste welzijn oplevert. Daarbij telt het welzijn van alle betrokkenen even zwaar mee.

Welzijn kan op verschillende manieren nader ingevuld worden. De hedonistische invulling definieert welzijn in termen van leed en genot: hoe meer genot den opzicht van leed, hoe groter is iemands welzijn. Genot en leed worden daarbij heel ruim opgevat. Zo vallen verveling en stress ook onder 'leed', en valt het plezier dat men kan ontleen aan het lezen van een goed boek of het helpen van anderen ook onder 'genot'. Een alternatieve invulling van welzijn definieert welzijn in termen van het bevredigen van wensen. Wat volgens deze opvatting iemands welzijn bevordert is niet de ervaring van genot als zodanig, maar het feit dat je krijgt wat je wilde. Doorgaans worden aan die verlangens nog verdere eisen gesteld, bijvoorbeeld dat zij gebaseerd moeten zijn op de juiste informatie, wil de vervulling ervan gelden als hetgeen welzijn in wezen uitmaakt. De gevolgen voor het welzijn van alle betrokkenen moeten bij elkaar worden opgeteld, om de waarde van de uitkomst van een handeling te bepalen. Dit hoeft overigens niet daadwerkelijk voor elke handeling te gebeuren, wat onmogelijk zou zijn en zeker niet het algehele welzijn zou maximieren. In de praktijk kunnen eenvoudigere regels worden gevolgd, mits het de beste regels zijn, in het licht van dit criterium van het moreel juiste.

Op die manier kunnen utilisten overigens niet alleen handelingen beoordelen. Ook bijvoorbeeld regels, intenties, en karakters kunnen worden beoordeeld op basis van hun gevolgen in termen van welzijn. Als een uitkomst van bijvoorbeeld een handeling wordt beoordeeld in termen van welzijn, dan kan dat op twee manieren gebeuren: óf men kijkt alleen naar de totale hoeveelheid welzijn die een uitkomst bevat, óf men kijkt naar de mate waarin het welzijn van voelende wezens bevordert of geschaad wordt. Dit lijkt op hetzelfde neer te komen, maar dat is niet altijd het geval, zoals later duidelijk kan worden. Een achterliggende reden om het welzijn van alle betrokkenen te maximieren

kan zijn dat welzijn op zichzelf waardevol is, en dat men het daarom moet maximeren. Daarbij zou het in eerste instantie gaan om het maximeren van welzijn, en zou het alleen een bijverschijnsel zijn dat een ieders welzijn even zwaar meetelt. Een achterliggende reden voor het utilistische principe kan ook zijn dat men iedereen in gelijke mate dient te respecteren en dat dit het beste gerealiseerd kan worden door het welzijn van alle betrokkenen op een neutrale manier te maximeren.

Het utilisme is een omstreden morele theorie. Een belangrijk punt van kritiek is dat het utilisme te onpersoonlijk zou zijn. Het utilisme zou alleen maar oog hebben voor welzijn als een abstract goed, en niet werkelijk geïnteresseerd zijn in voelende wezens. Zo zijn er geen grenzen aan wat men met een individu mag doen, als dit het algehele welzijn maximeert. Het idee dat individuen vervangbaar zouden zijn, zolang de totale hoeveelheid welzijn maar niet afneemt, is een fel bekritiseerde implicatie van het utilisme, die later nog uitgebreid aan de orde zal komen.

Wat de veehouderij betreft zet het utilisme het dierenwelzijn prominent op de agenda, maar hoe te denken over het (pijnloos) doden van dieren? Als een dier een gelukkige toekomst gehad zou kunnen hebben, dan gaat dit mogelijke welzijn verloren door het dier te doden. Daarom is het doden van gelukkige dieren problematisch met het oog op het maximeren van welzijn. Daarnaast kan ook beargumenteerd worden dat het dier geschaad wordt door de dood indien de dood het dier toekomstig welzijn ontnemt. In vergelijking met de optie om het dier te laten leven is het doden van een dier alleen toegestaan als het daardoor ontnomen welzijn gecompenseerd wordt. Het welzijnsverlies door het doden van een gelukkig dier zou zijn toegestaan als dit de beste manier zou zijn om nog meer welzijnsverlies te voorkomen, of als dit de beste manier zou zijn om zo veel mogelijk welzijnswinst teweeg te brengen. Aangezien dit niet het geval lijkt te zijn in de veehouderij, wordt een andere manier om het welzijnsverlies te compenseren relevant: het vervangbaarheidsargument.

3. Zijn dieren minder geschaad door de dood?

Wat maakt dat de dood erg is voor een wezen? Om in te zoemen op de dood als zodanig, in tegenstelling tot mogelijke pijn of angst voor de dood, kunnen wij ervan uitgaan dat het wezen ongemerkt en onverwacht dood gaat, tijdens de slaap. De vraag is wat maakt dat dit erg is voor het wezen zelf, mogelijke gevolgen voor anderen buiten beschouwing gelaten.

Volgens één visie bestaat de schade die de dood is erin dat men wensen die men had vóór de dood niet meer kan vervullen. Naarmate wezens minder vooruit plannen, minder besef hebben van hun sterfelijkheid en daardoor minder hun eigen voortbestaan kunnen wensen zouden zij volgens deze visie minder geschaad worden door de dood. Volgens een andere visie is de dood schadelijk omdat hij alles ontnemt dat voor het wezen in de toekomst van waarde zou zijn geweest. Hierbij gaat het dus niet om het frustreren van verlangens, en niet om wat een dier had gewild. Het gaat om het ontnemen van waarde, om hetgeen waardevol was geweest voor het dier. Voor die tweede visie maakt het geen verschil hoeveel plannen en toekomstbesef het wezen heeft.

Een kanttekening die bij die laatste visie gemaakt wordt is dat het verlies aan waarde minder zwaar mee dient te tellen, naarmate er minder psychologische verbondenheid is tussen het wezen op het moment van de dood en het wezen dat die waarde in de toekomst zou hebben ervaren. Als wezens dus meer in het moment leven, en minder vooruit of terug kijken, dan zijn zij, volgens deze visie, minder psychologisch verbonden met hun 'toekomstige zelf', aan wie de dood de waardevolle ervaringen ontnemt. Daarom zou dat verlies voor dit soort wezens minder zwaar mee moeten tellen. Het kan echter betwist worden of deze visie overtuigend is.

Hoe dan ook kan volgens elke van de boven genoemde visies gesteld worden dat de dood de dieren die doorgaans op de borden van mensen terecht wel degelijk schaadt. Dieren als koeien, varkens, schapen, kippen, en vissen, lijken te beschikken over het vermogen om tot op zekere hoogte vooruit te plannen en om zich dingen te herinneren, en te kunnen leren. Daarom schaadt de dood deze dieren wel degelijk, zelfs volgens de laatste twee hier besproken visies.

4. Het vervangbaarheidsargument

Het vervangbaarheidsargument houdt in dat het welzijnsverlies dat ontstaat door het doden van een dier dat anders een gelukkige toekomst gehad zou kunnen hebben, gecompenseerd kan worden indien een ander dier het gedode dier als het ware zal vervangen als een drager van welzijn. Voorwaarde is dat het leven van het vervangende dier minstens net zo veel welzijn bevat als de toekomst van het gedode dier bevat zou kunnen hebben. Bovendien gaat het argument alleen op onder de voorwaarde dat het vervangende dier anders niet zou hebben bestaan. Daarnaast mag het doden en vervangen geen niet gecompenseerde negatieve bijwerkingen hebben, zoals angst of pijn van het te doden dier of van derden. Overigens geldt het vervangbaarheidsargument specifiek voor dieren die anders een gelukkige toekomst gehad zouden kunnen hebben, voor ‘gelukkige’ dieren dus. Indien het doden van een dier een toekomst met noodzakelijk negatief welzijn voorkomt, dan maximeert het doden, *ceteris paribus*, het algehele welzijn, en hoeft er geen vervanging plaats te vinden.

In verschillende praktijken worden dieren gedood en vervangen. Dat geldt voor de productie van vlees, maar ook voor de productie van zuivel en eieren, voor het telen van vis, voor het sporthengelen, de sportjacht, en het kweken van dieren voor dierproeven. Al worden de dieren in deze praktijken niet om morele redenen vervangen, maar om de

praktijk gaanden te houden, dan zou het vervangbaarheidsargument toch een morele rechtvaardiging kunnen geven voor deze praktijken.

Het is overigens maar de vraag of de landbouwhuisdieren echt een gelukkig leven hebben, aangezien in deze praktijk de dieren allerlei lichamelijk en emotioneel leed te verduren hebben. Ook is het de vraag of het doden van dieren plaats vindt zonder negatieve bijwerkingen voor de dieren die gedood worden zelf of anderen. In principe zijn er echter praktijken denkbaar waarin aan de voorwaarden van het vervangbaarheidsargument zou worden voldaan.

Als dieren onder de genoemde voorwaarden vervangbaar zouden zijn, geldt dat dan ook voor mensen? Peter Singer, een van de meest beroemde hedendaagse filosofen, die een utilistische (dier-) ethiek verdedigt, heeft gepoogd om het bereik van het vervangbaarheidsargument te beperken. Singer blijkt bereid tot vergaande aanpassingen van zijn morele theorie om te voorkomen dat normale volwassen mensen vervangbaar zouden zijn. Uiteindelijk blijkt echter dat allerlei strategieën om het bereik van het vervangbaarheidsargument te beperken mislukken. Met behulp van Singers meest recente strategie, die ervan uit gaat dat welzijn alleen neutraal of negatief kan zijn, zouden sommige wezens kunnen worden uitgesloten van vervangbaarheid, namelijk de wezens waarvan het leven resulteert in negatief welzijn. Singers strategie leidt tot de conclusie dat precies deze wezens die dan niet vervangbaar zouden zijn noodzakelijkerwijs negatief welzijn zouden hebben en dus beter niet geleefd zouden hebben.

5. Total View tegenover Prior Existence View

Het vervangbaarheidsargument is gebaseerd op een bepaalde visie op wiens welzijn mee telt. Het argument gaat alleen op als men bij het beantwoorden van de vraag of men een bepaald dier mag doden, reeds rekening houdt met het mogelijk toekomstige welzijn van een mogelijk

toekomstig dier dat men daarvoor in de plaats zou kunnen zetten. Dat mogelijk toekomstige dier bestaat nog niet als men de keuze van het wel of niet doden van het bestaande dier maakt. Of het dier wel of niet ooit zal bestaan hangt ervan af hoe de keuze die men net aan het maken is uitvalt. Daarom is dat mogelijk toekomstige dier een zogenaamd ‘contingent’ wezen. Het bestaat nog niet, en of het zal bestaan hangt af van de morele keuze die men net aan het maken is.

Het is omstrede binnen het utilisme of het welzijn van dit soort contingente wezens mee zou moeten tellen in het bepalen van de waarde van een uitkomst. Volgens de *Total View* dient het welzijn van contingente wezens wel degelijk mee te tellen. De *Total View* houdt in dat het welzijn van alle bestaande en alle mogelijke wezens mee geteld dient te worden. Volgens de *Prior Existence View* daarentegen dient het mogelijke welzijn van contingente wezens niet te worden meegeteld. De *Prior Existence View* telt alleen het welzijn van wezens die op het moment van de keuze reeds bestaan, of die onafhankelijk van de keuze zullen bestaan.

Een voorbeeld kan de verduidelijken. Laten wij aannemen dat een stel zich afvraagt of zij wel of niet een kind zouden moeten krijgen. Het stel zelf zou van het hebben van een kind minder gelukkig worden, omdat het hen zou belemmeren in hun favoriete vrijetijdsbesteding. Het kind zelf zou echter een gelukkig leven hebben. Het zou opgroeien onder gunstige omstandigheden, en er zou goed voor gezorgd worden. Volgens de *Total View* moet bij die keuze het mogelijke welzijn van het mogelijke kind meetellen. Het zou dus kunnen dat het welzijnsverlies voor de ouders gecompenseerd zou worden door het welzijn van het kind. Volgens de *Prior Existence View* telt in dit geval alleen het welzijn van de ouders. Dit wil overigens niet zeggen dat de *Prior Existence View* geen rekening zou kunnen houden met toekomstige generaties. Als het bestaan van de toekomstige wezens niet afhangt van de keuze die wij maken, dan dient het welzijn van die wezens wel degelijk mee te tellen, ook indien zij nog niet bestaan. Als bijvoorbeeld een vrouw die hoe dan

ook een kind zal hebben zich afvraagt waar zij zal gaan wonen, dan moet bij die keuze het toekomstige welzijn van haar toekomstige kind meegewogen worden.

Volgens de *Prior Existence View* gaat het vervangbaarheidsargument niet op, omdat het welzijn van het mogelijk toekomstige dier niet meetelt. Het is daarom niet mogelijk om op die manier het welzijnsverlies dat veroorzaakt wordt door het doden van een dier te compenseren. Het aanvaarden van het vervangbaarheidsargument, en daarmee het verdedigen van diervriendelijke veehouderij hangt er dus binnen het utilisme van af of men een beroep doet op de *Total View* dan wel de *Prior Existence View*.

Het dient echter verhelderd te worden of de *Prior Existence View* überhaupt een coherente utilistische visie is. Het lijkt namelijk niet passen binnen het utilisme om het welzijn van sommige wezens buiten beschouwing te laten. Het utilisme gaat er immers juist van uit dat het welzijn van alle betrokkenen in gelijke mate mee telt. De *Prior Existence View* is alleen onder twee voorwaarden te verenigen met de uitgangspunten van het utilisme. Als men ten eerste ervan uitgaat dat uitkomsten beoordeeld dienen te worden door te kijken naar de mate waarin het welzijn van de betrokkenen geschaad en bevorderd wordt, en als men ten tweede ervan uitgaat dat dat het laten bestaan van een wezen diens welzijn niet kan schaden of bevorderen, dan heeft het zin om het welzijn van contingente wezens buiten beschouwing te laten. Hieronder zullen beide aannames nader worden toegelicht en aannemelijk worden gemaakt.

6. Kan het laten bestaan van een wezen diens welzijn bevorderen?

Natuurlijk kan iemands leven goed, slecht of neutraal zijn in termen van welzijn. De vraag of het laten bestaan van een persoon diens welzijn kan bevorderen is echter lastig te beantwoorden. Als men wil weten of iets iemands welzijn bevordert, dan is het gebruikelijk om het welzijn van het

de betreffende wezen na die gebeurtenis te vergelijken met het welzijn ervoor. Nog beter lijkt het om het welzijn van het wezen na de gebeurtenis te vergelijken met zijn of haar welzijn in de alternatieve situatie waarin het deze gebeurtenis niet had plaats gevonden. Dat laatste is soms lastig om te weten: wat was er bijvoorbeeld gebeurd als ik een bepaalde baan niet had aangenomen? In het geval van iemand laten bestaan is het duidelijk wat de alternatieve stand van zaken zou zijn: het betreffende wezen zou niet bestaan. Het is dus zaak om een vergelijking te maken tussen het welzijn van het wezen als het bestaat en het welzijn van hetzelfde wezen als het niet zou hebben bestaan.

Kan men het bestaan en het niet bestaan van een wezen überhaupt vergelijken in termen van het welzijn van dat wezen? Men is het er over eens dat een niet bestaand wezen geen welzijn heeft. Eén positie is dat de afwezigheid van positief en negatief welzijn in het geval van niet bestaan te boek kan worden geslagen als nul welzijn op de denkbeeldige meetlat van welzijn die rijkt van negatief, via neutraal naar positief welzijn. Hier wordt een andere positie tegenin gebracht die inhoudt dat de afwezigheid van welzijn niet verward mag worden met een neutrale welzijnsscore. Vergelijkbare discussies spelen in de kwantum fysica, bijvoorbeeld omtrent de vraag of een foton een massa van nul of géén massa heeft. Het ontwarren van de structuur van de argumentatie schept helderheid in het actuele debat en laat zien dat de positie verdedigbaar is, die inhoudt dat de afwezigheid van welzijn iets anders is dan het scoren van nul op de schaal van welzijn.

Mocht men er wel van uit willen gaan dat het niet-bestaan neutrale waarde heeft in termen van welzijn (in plaats van geen waarde), dan rijst de vraag voor wie het niet bestaan dan neutrale waarde zou hebben. In het geval van niet bestaan is er immers niemand voor wie het niet bestaan waarde zou kunnen hebben. Dat geldt ook voor neutrale waarde. Om een vergelijking te maken tussen het bestaan en het niet bestaan in termen van het welzijn van een wezen zouden enkele logische vereisten die normaal gesproken op vergelijkingen van toepassing zijn

over boord moeten worden gegooid. Er wordt weleens beweert dat het over boord gooien van die vereisten ook al nodig is om te kunnen zeggen dat het iemands welzijn bevordert om hem of haar van de dood te redden, dus voor vergelijkingen tussen leven en dood. Dit is echter niet het geval. Dergelijke uitspraken kunnen ook worden gedaan zonder aan de dood een neutrale welzijns waarde toe te kennen. Het volstaat om te zeggen dat een lang gelukkig leven meer welzijn voor het betreffende wezen oplevert dan een kort gelukkig leven. Men vergelijkt dus twee levens in termen van welzijn, en niet het leven met de dood. Het lijkt daarom ad hoc om de logische vereisten wel over boord te gooien voor vergelijkingen tussen het bestaan en het niet bestaan in termen van welzijn. Het lijkt dus een mogelijke positie dat men iemands welzijn niet kan bevorderen of schaden door hem of haar te laten bestaan. Daarmee kan aan de eerste voorwaarde voor de *Prior Existence View* als coherente utilistische visie worden voldaan.

7. Person-Affecting Restriction en Non-Identiteitsprobleem

De tweede voorwaarde die de *Prior Existence View* nodig heeft om een coherente utilistische visie te zijn is het accepteren van de visie dat uitkomsten beoordeelt dienen te worden in termen van het nut en de schade voor personen die zij bevatten en niet in termen van de totale hoeveelheid welzijn die een uitkomst bevat als zodanig. De visie die inhoudt dat men een uitkomst niet dient te beoordelen op basis van de totale hoeveelheid welzijn als zodanig, maar dat men moet nagaan in hoeverre in deze uitkomst het welzijn van voelende wezens wordt bevorderd dan wel geschaad, is de *Person-Affecting Restriction*. Het gaat er volgens deze visie om hoe voelende wezens in hun welzijn geraakt worden. Een uitkomst kan volgens deze visie alleen beter (of slechter) zijn, als hij beter (of slechter) is vóór wezens. De visie die alleen naar de totale hoeveelheid welzijn kijkt, onafhankelijk van wat dat betekent voor het bevorderen en schaden van iemands welzijn heet de *Impersonal View*.

Het belangrijkste argument dat naar voren is gebracht tegen de *Person-Affecting Restriction* is het Non-Identiteitsprobleem. Dat kan verduidelijkt worden aan de hand van de casus van het veertienjarige meisje dat zich afvraagt of het nu een kind zal krijgen of over enkele jaren. Stel dat het meisje nu niet heel goed voor het kind zou kunnen zorgen, en dat het kind hierdoor een moeilijke start in het leven zou hebben, en daardoor enkele problemen als volwassene zou ondervinden. Stel verder dat het kind al met al toch een leven zou hebben dat de moeite waard is. Als het meisje later een kind zou krijgen zou dat kind een gelukkiger leven hebben. Stel dat de keuze verder niemand's welzijn zou raken, ook het meisje zelf zou door elke optie even goed af zijn. Het lijkt in dit geval, zeker volgens het utilisme, juist te zijn om te wachten met het krijgen van een kind. Er wordt immers niemand door geschaad, en het kind zal gelukkiger zijn. Het probleem is echter dat het meisje, indien het zou wachten, niet hetzelfde kind zou krijgen, maar een ander, niet identiek kind. Er zou dus geen enkel wezen beter af zijn door te wachten. Dat is het non-identiteitsprobleem, dat zich voordoet als de in hun welzijn geraakten in de verschillende uitkomsten enkel niet-identieke individuen zijn. Als een uitkomst allen beter kan zijn als hij beter is voor iemand, dan lijkt het in dit geval niet beter te zijn voor het veertienjarige meisje om te wachten. Dit lijkt tegen-intuïtief.

Het non-identiteitsprobleem kan omzeild worden door de *Person-Affecting Restriction* breder op te vatten. Volgens die brede opvatting gaat het bij wat beter is voor een wezen niet om een specifiek individu, maar om wezens, wie het ook mogen zijn. In die zin is uitkomst A beter dan uitkomst B indien uitkomst A beter is voor de A-wezens dan uitkomst B is voor de B-wezens. Deze ruimere invulling van de *Person-Affecting Restriction* kan verdedigd worden door te laten zien dat identiteit een begrip is dat op verschillende manieren kan worden ingevuld, en dat in verschillende contexten verschillende invullingen van identiteit relevant zijn. Als het gaat om de morele vragen die in een non-identiteitscasus doorgaans aan de orde zijn, dan lijkt een genetische invulling van

identiteit niet de meest relevante. In deze gevallen lijkt een *de-dicto* invulling van wat beter voor iemand is relevant. Volgens deze invulling geldt: Als men twee uitkomsten O1 en O2 vergelijkt, dan is uitkomst O1 is *de dicto* beter voor S als degene die S is in O1 beter af is in O1 dan degene die S is in O2 af is in O2. In het geval van het veertienjarige meisje betekent deze visie, in combinatie met het utilisme, dat het meisje moet wachten, omdat haar volgende kind (wie het ook moge zijn) beter af is als zij wacht.

Indien ruim ingevuld hoeft de het idee dat uitkomsten dienen te worden beoordeeld op basis van de mate waarin zij het welzijn van wezens bevorderen of schaden (in plaats van de hoeveelheid welzijn die zij bevatten als zodanig) niet stuk te lopen op het non-identiteitsprobleem. Bovendien blijkt die ruime invulling ook beter aan te sluiten bij het utilisme dat zich typischer wijs vooral bekommert om gevolgen voor het welzijn van wezens waarbij de specifieke identiteit van een wezen er als zodanig niet toe doet. De *Person-Affecting Restriction* kan dus verdedigd worden, en daarmee kan ook aan de tweede voorwaarde voor de *Prior Existence View* als coherente utilistische visie worden voldaan.

8. Repugnant Conclusion en Expected Misery Argument

Aangezien zowel de *Total View* als de *Prior Existence View* coherente utilistische visies zijn is het interessant om wat implicaties van beide visies te verkennen. De meest beruchte implicatie van de *Total View* staat bekend als de *Repugnant Conclusion*, de ‘weezinwekkende conclusie’. Stel er is een populatie van minstens 10 miljoen heel gelukkige mensen. Door meer kinderen te krijgen zou deze bevolking kunnen groeien, maar stel dat die groei steeds meer ten koste zou gaan van het welzijn van de reeds bestaande mensen, en dat de nieuwe mensen matig gelukkig zouden zijn. Uiteindelijk zou uitbreiding kunnen resulteren in een enorm omvangrijke bevolking, vele malen omvangrijker dan de oorspronkelijke bevolking,

waarbij een ieder een welzijnsniveau zou hebben dat net boven het neutrale niveau ligt. Als de bevolking maar groot genoeg zou zijn, dan zouden alle kleine beetjes welzijn bij elkaar opgeteld tot een groter totaal aan welzijn kunnen leiden dan de oorspronkelijke bevolking bevatte. Volgens de Total View zou het dan een goede zaak zijn indien de oorspronkelijke bevolking op deze manier het totaal aan welzijn zou vergroten. Dit lijkt intuïtief een onaantrekkelijke implicatie van de Total View.

De *Prior Existence View* kan deze implicatie vermijden, omdat het mogelijk toekomstige welzijn van contingente wezens niet mee telt. Het welzijnsverlies van de reeds bestaande wezens zou een halt toeroepen aan verdere uitbreiding van de bevolking. Precies de implicatie dat het welzijn van contingente wezens niet mee telt kan in andere situaties echter als een sterk argument tegen de *Prior Existence View* naar voren worden gebracht. Het meest bekende voorbeeld hiervoor is de casus van het ellendige kind: Een stel weet al zeker, vóór de conceptie dat zij alleen maar een kind zouden kunnen krijgen dat een kort leven vol niet te compenseren leed zou hebben, een ellendig leven dus met ronduit negatief welzijn. Het lijkt dat dit een reden zou moeten zijn om dit kind niet te krijgen. De *Prior Existence View* lijkt hier echter geen rekening mee te kunnen houden, omdat het gaat om een contingent wezen, wiens welzijn niet meetelt.

Nu kan men zich afvragen hoe reëel deze casus is, want zouden ouders dit kind willen krijgen, en zou dat echt geen gevolgen voor het welzijn van reeds bestaande of onafhankelijk van deze keuze bestaande toekomstige wezens hebben? Maar stel dat de ouders het kind desondanks zouden willen krijgen. De *Prior Existence View* kan het welzijn van dit kind pas meetellen zodra het kind bestaat als een voelend wezen. Zodra het kind daadwerkelijk alleen maar pijn ervaart en er met zekerheid geen beterschap mogelijk is, dan zou abortus of euthanasie volgens elke versie van het utilisme te verkiezen zijn. Maar zelfs al voordat het kind bestaat als een voelend wezen, kan het plan om het

kind te krijgen en te laten lijden veroordeeld worden om utilistische redenen, ook als men uitgaat van de *Prior Existence View*. Het utilisme kan namelijk ook het karakter van het stel beoordelen in termen van de tendentie daarvan om het welzijn te maximeren. Het karakter dat te verenigen is met het ter wereld brengen en laten lijden van zo'n kind is volgens het utilisme een slecht karakter. Het *Prior Existence Utilisme* heeft dus wel degelijk de mogelijkheden om een dergelijk geval moreel te veroordelen. (In deze context is het interessant om op te merken hoe utilisten omgaan met tegen intuïtieve implicaties van morele visies.)

9. Veganisme versus diervriendelijke veehouderij

Welke implicaties hebben *Total View Utilitarisme* en *Prior Existence View Utilitarisme* voor de praktijk van de veehouderij? Eén argument ten gunste van diervriendelijke veehouderij stelt dat deze praktijk juist goed is voor de dieren. De dieren hebben immers een gelukkig leven en zouden anders helemaal niet hebben bestaan. Een gelukkig leven, ook indien dit vroegtijdig wordt beëindigd, is nog altijd beter dan helemaal niet bestaan, aldus dit argument. Dit argument, ook wel de *Logic of the Larder* – de logica van de voorraadkast – genoemd, houdt in dat de echte dierenvriend een vleeseter is.

Dit argument kan om praktische redenen bekritiseerd worden. Zelfs al wil men zo veel mogelijk gelukkige dieren laten bestaan, dan zou men geen dierlijke producten moeten produceren of consumeren. Er kunnen met plantaardig voedsel over het algemeen meer mensen gevoed worden, en in de ruimte die vrij komt kunnen wilde dieren leven. In plaats van een stuk vlees zou men dan beter een aantal gelukkige muizen kunnen aanschaffen. En zou het argument ook van toepassing zijn op mensen? Meer fundamenteel veronderstelt het argument dat men het welzijn van een wezen kan bevorderen door het leven te laten bestaan. Zoals wij hebben gezien verwerpt *Prior Existence Utilisme* deze aanname, en is dus niet te verenigen met dit argument. *Total Utilitarisme*

is in principe wel compatibel met het idee dat een het welzijn van een dier bevordert kan worden door het dier te laten bestaan. Zelfs zou een het *Total View* utilisme dit niet onderschrijven, dan nog geldt volgens deze visie dat het laten bestaan van meer gelukkige wezens de totale hoeveelheid welzijn bevordert.

Aangezien het *Prior Existence* Utilisme het vervangbaarheidsargument niet accepteert, zal het routinematige doden van dieren in de veehouderij door deze versie van het utilisme niet verdedigd worden. Veganistische landbouw zonder veehouderij is een ideaal dat beter bij het *Prior Existence* Utilisme zou passen. Het *Total* Utilisme daarentegen accepteert wel het vervangbaarheidsargument, met alle implicaties die dit in principe ook voor mensen heeft. Wil men sommige wezens van vervangbaarheid uitsluiten, zoals Singer voorstelt, dan heeft dit de implicatie dat deze wezens allen maar negatief welzijn kunnen hebben, en dus beter niet geleefd zouden kunnen hebben.

10. Conclusie

Hoe het utilisme oordeelt over het breed gedeelde ideaal van diervriendelijke veehouderij, en andere praktijken waarin gelukkige dieren worden gedood en vervangen, hangt af van welke versie van het utilisme men voor ogen heeft. *Prior Existence* Utilisme is naast het *Total* Utilisme ook een coherente versie van deze morele theorie. *Prior Existence* Utilisme biedt mens en dier een betere bescherming tegen het doden. Bovendien is het een versie van het utilisme die minder vatbaar is voor een prominent kritiekpunt dat tegen het utilisme naar voren wordt gebracht, namelijk het bezwaar dat het utilisme te onpersoonlijk zou zijn en zich alleen zou bekommeren om welzijn als een abstracte hoeveelheid.

In tegenstelling tot wat vaak wordt beweert bekommert het utilisme zich dus niet noodzakelijkerwijs alleen om het voorkomen van onnodig leed, maar heeft het wel degelijk ook iets in te brengen tegen het

doden van dieren. Dit is relevant voor zogenaamd diervriendelijke vormen van veehouderij en voor andere praktijken van diergebruik.

ZUSAMMENFASSUNG

1. Einleitung

Es scheint eine recht allgemein geteilte Übereinstimmung darüber zu geben, dass heutige Formen intensiver Tierhaltung moralisch unzureichend sind. Obgleich in dieser Angelegenheit Worte und Taten aus allerlei Gründen nicht unbedingt im Einklang miteinander sind, lassen Bürger und Konsumenten bei Nachfrage wissen, dass die Tierhaltung tierfreundlicher sein sollte. Allerlei gesellschaftliche Organisationen setzen sich ein für eine tierfreundliche Tierhaltung, und auch an der Politik geht dieses Bemühen nicht unbemerkt vorbei. Das Herbeiführen einer tierfreundlichen Tierhaltung wird erfahren als ein moralisches Ideal. Der hieran zugrunde liegende Gedanke ist, dass wir nicht einfach alles mit Tieren machen dürfen und, dass wir uns um das Wohlergehen der Tiere kümmern sollten. Unnötiges Leiden sollte in unserem Umgang mit Tieren vermieden werden.

Dieses recht allgemein akzeptierte Ideal der tierfreundlichen Tierhaltung beinhaltet eine gewisse Spannung. Einerseits gelten Tiere als moralisch beachtenswert und sollte ihr Wohlergehen uns wichtig sein. Andererseits ist auch die tierfreundliche Tierhaltung noch immer eine

Form der Tierhaltung, was bedeutet, dass Tiere in großer Zahl routinemäßig getötet werden so dass Menschen ihre Produkte benutzen können. Da der Gebrauch tierischer Produkte jedenfalls in unseren Regionen nicht nötig ist um gesund zu leben, scheinen vor allem Überlegungen die auf Bequemlichkeit und Vergnügen beim Gebrauch beruhen eine Rolle zu spielen. Ist das routinemäßige Töten von Tieren für diesen Zweck zu vereinbaren mit dem Gedanken, dass Tiere moralisch relevant sind, und, dass wir für ihr Wohlergehen sorgen sollten? Gibt es eine Moraltheorie, die das Ideal der tierfreundlichen Tierhaltung unterbauen kann?

Diese Untersuchung setzt sich mit der Frage auseinander, wie die utilitaristische Moraltheorie sich zum Ideal der tierfreundlichen Tierhaltung verhält. Der Utilitarismus, eine der klassischen Moraltheorien, und zwar traditionell diejenige, die das Wohlergehen der Tiere am meisten mit einbezogen hat, scheint auf den ersten Blick eine geeignete Theorie zu sein, um das Ideal der tierfreundlichen Tierhaltung zu verteidigen. Utilitarismus vertritt die These, dass genau jene Handlungen moralisch richtig sind, die das maximale Wohlergehen für alle Beteiligten zur Folge haben. Das Wohlergehen von Tieren ist dabei gleichwertig an das Wohlergehen von Menschen. Gleichzeitig ist der Utilitarismus aber nicht prinzipiell gegen das Töten von Tieren. Utilitaristisch betrachtet, ist keine einzige Handlung als solche moralisch richtig oder falsch. Alles hängt davon ab welche Folgen die Handlung in der jeweiligen Situation hat. Also würde auch das Töten von Menschen wie von Tieren, erlaubt sein, sofern es dies die Handlung ist, die das Wohlergehen aller Beteiligten maximiert. Aus diesen Gründen ist es fraglich, ob Utilitarismus das Ideal der tierfreundlichen Tierhaltung unterstützen kann.

2. Utilitarismus und Tierethik

Utilitarismus ist eine der zwei wichtigsten Moraltheorien der Aufklärung. Als solche will uns Utilitarismus auf systematische und zusammenhängende Weise erklären, was moralisch richtig und falsch ist, und warum dies so ist. Utilitarismus beurteilt Handlungen einzig und allein auf Grund ihrer Folgen und beinhaltet also, dass nur diejenigen Handlungen moralisch richtig sind, die von allen möglichen ausführbaren Handlungen, die besten Folgen haben. Bei der Beurteilung davon, was die besten Folgen sind, geht es nur um das Wohlergehen. Richtig ist die Handlung, die das meiste Wohlergehen erzeugt. Dabei zählt das Wohlergehen aller Beteiligten in gleicher Weise mit.

Was genau mit „Wohlergehen“ gemeint ist, kann auf verschiedene Weise erläutert werden. Die hedonistische Definition fasst Wohlergehen im Sinne von mentalen Empfindungen von Leiden und Genuss im weitesten Sinne auf. Langeweile und Stress gelten dann auch als „Leiden“ und Genuss kann man auch empfinden durch das Lesen eines guten Buches oder dadurch anderen zu helfen.. Eine alternative Definition fasst Wohlergehen auf als Wunschbefriedigung. Gemäß dieser Auffassung besteht positives Wohlergehen nicht in einer mentalen Befriedigungserfahrung als solcher, sondern darin, dass man bekommt was man will. Üblicherweise werden an Wünsche weitere Bedingungen gestellt. Sie müssen zum Beispiel auf richtiger Information beruhen , wenn die Erfüllung dessen, was Wohlergehen ausmacht gelten soll. Die Folgen für das Wohlergehen aller müssen zusammengezählt werden um den Wert des möglichen Resultats einer Handlung zu bestimmen. Diese moralische Kosten-Nutzen Rechnung stößt in der Praxis an ihre Grenzen. Utilitarismus erkennt diese Grenzen und hantiert für die Praxis ein Überschlagskriterium.

Auf diese Weise können Utilitaristen nicht nur Handlungen beurteilen. Regeln, aber auch Absichten und Charaktereigenschaften können im Hinblick auf ihre Folgen für das allgemeine Wohlergehen

beurteilt werden. Resultate, beispielsweise von einer Handlung, können auf zwei verschiedene Weisen beurteilt werden. Entweder man achtet ausschließlich auf die absolute Quantität des Wohlergehens, die ein Resultat beinhaltet, oder man achtet auf das Ausmaß in dem dem Wohlergehen fühlender Wesen geschadet oder es gefördert wird.. Ein tieferer Grund um das Wohlergehen aller Beteiligten zu maximieren kann die Annahme sein, dass Wohlergehen an sich wertvoll ist, und man es deshalb maximieren sollte. Es ginge dann in erster Linie darum, Wohlergehen zu maximieren. Dass man das Wohlergehen aller, in gleicher Weise berücksichtigen muss, wäre davon nur eine Beiläufigkeit. Ein dem utilitaristischen Moralprinzip zugrundeliegender Gedanke kann jedoch auch sein, dass man jeden in gleicher Weise respektieren sollte, und das dies am besten geschehen kann, indem man das Wohlergehen aller in neutraler Weise maximiert.

Utilitarismus ist eine umstrittene Moraltheorie. Ein wichtiger Kritikpunkt ist, dass der Utilitarismus zu unpersönlich sei. Demnach ginge es dem Utilitaristen nur um Wohlergehen als ein abstraktes Gut, und nicht um fühlende Wesen. So stellt Utilitarismus keine Grenzen an das, was man mit einem anderen tun darf, solange dies das allgemeine Wohlergehen maximiert. Die Annahme, dass Individuen ersetzbar seien, unter der Bedingung dass nur die totale Menge Wohlergehen gleich bleibt, ist eine stark kritisierte Implikation des Utilitarismus, die später ausführlich besprochen werden wird.

Was die Tierhaltung angeht, so hat richtet Utilitarismus Aufmerksamkeit auch auf das Wohlergehen von Tieren. Wie aber steht es mit dem Töten von Tiere auf eine schmerzlose Art und Weise? Wenn ein Tier eine glückliche Zukunft gehabt haben könnte, dann geht dieses mögliche Wohlergehen verloren, wenn man das Tier tötet. Deshalb ist das Töten „glücklicher“ Tiere problematisch im Hinblick auf das Maximieren des allgemeinen Wohlergehens. Außerdem kann auch argumentiert werden, dass der Tod dem Tier schadet, insofern als der Tod dem Tier zukünftiges Wohlergehen entnimmt. Verglichen mit der

Möglichkeit das Tier am Leben zu lassen, ist das Töten des Tieres nur dann erlaubt, wenn das dadurch entnommene Wohlergehen kompensiert wird. Der Verlust an Wohlergehen durch das Töten eines glücklichen Tieres wäre erlaubt, wenn dies die beste Weise wäre, um noch mehr Verlust von Wohlergehen zu vermeiden, oder um so viel Wohlergehen wie möglich zu Stand zu bringen. Da dies in der Tierhaltung nicht der Fall zu sein scheint, wird eine andere Weise relevant, um Verlust von Wohlergehen zu kompensieren: das Ersetzbarkeits-Argument.

3. Schadet der Tod den Tieren?

Was sorgt dafür, dass der Tod schlecht ist für ein Wesen? Ist es die Angst vor Schmerzen oder einem schmerzvollen Tod? Wie ist das friedliche Entschlafen eines Wesen dann zu beurteilen? Die Frage, inwiefern der Tod schlecht ist für ein Wesen, berücksichtigt mögliche Folgen für andere nicht.

Einer Variation der Wunscherfüllungstheorie nach ist der Schaden durch den Tod darin gelegen, dass man Wünsche die man vor dem Tod hatte, nicht mehr erfüllen kann. Der Tod wäre demnach weniger schädlich für Wesen, die weniger voraus planen, und wegen mangelnder Einsicht in die eigene Sterblichkeit, den Wunsch auf das eigene Fortbestehen weniger oder gar nicht haben. Der Entbehrungstheorie zufolge ist der Tod schädlich insofern als er alles entnimmt was dem Wesen in der Zukunft von Wert gewesen wäre. Es geht hierbei also nicht um das Frustrieren von Wünschen, und nicht um was ein Wesen *gewollt* hätte. Es geht vielmehr um das Entnehmen von Wert, um das, was wertvoll gewesen wäre für das Wesen. Bei dieser zweiten Theorie ist es unwichtig, wie viele Pläne oder Bewusstsein von der Zukunft ein Wesen hat.

Eine Einschränkung die bei der zweiten Theorie gemacht werden kann, ist dass der Verlust an Wert weniger mitzählt, wenn weniger

psychologische Verbundenheit besteht zwischen dem Wesen im Moment des Todes und dem Wesen, das den Wert in der Zukunft erfahren hätte, wenn es nicht gestorben wäre. Wenn Wesen mehr im Moment leben, weniger voraus oder zurück blicken, dann sind sie, so der Gedanke, weniger psychologisch verbunden mit ihrem „zukünftigen Selbst“, welchem der Tod die wertvollen Erfahrungen entnimmt. Deshalb solle der Verlust für solche Wesen weniger zählen. Es kann jedoch bezweifelt werden, ob diese Einschränkung plausibel ist.

Gemäß jeder der oben genannten Theorien kann gesagt werden, dass der Tod den Tieren, die im Allgemeinen auf den Tellern der Menschen landen schadet. Kühe, Schweine, Schafe, Hühner und Fische scheinen das Vermögen zu besitzen, um in einem gewissen Maß voraus zu planen, sich an Dinge zu erinnern, und auch lernen zu können.

4. Das Ersetzbarkeits-Argument

Das Ersetzbarkeits-Argument beinhaltet, dass der Verlust an Wohlergehen durch das Töten eines Tieres, das ansonsten eine glückliche Zukunft gehabt haben könnte, kompensiert werden kann, indem ein anderes Tier das getötete Tier ersetzt. Bedingung ist, dass das Leben des Ersatzes mindestens so viel Wohlergehen enthält als die Zukunft des getöteten Tieres enthalten haben könnte. Außerdem gilt das Argument nur unter der Bedingung, dass das Ersatztier ansonsten nicht gelebt hätte. Darüber hinaus darf das Töten und Ersetzen keine nicht kompensierbare negative Nebenwirkungen haben, wie zum Beispiel Angst oder Schmerz für das getötete Tier. Das Ersetzbarkeits-Argument gilt gerade für Tiere, die eine glückliche Zukunft gehabt haben könnten, also für „glückliche“ Tiere. Als das Töten eines Tieres dafür sorgt, dass notwendigerweise negatives Wohlergehen vermieden wird, dann maximiert das Töten, *ceteris paribus*, das Wohlergehen, und dann ist Ersetzen nicht nötig.

In verschiedenen Umständen werden Tiere getötet und ersetzt. Das gilt für die Fleischproduktion, aber auch für die Produktion von Milchprodukten und Eiern, für das Züchten von Fisch, für das Sportangeln, die Sportjagd, und das Züchten von Tieren für Tierversuche. Auch wenn die Tieren in diesen Umständen nicht aus moralischen Überlegungen heraus ersetzt werden, sondern lediglich um die Art des Benutzens zu konstituieren, böte das Ersetzbarkeits-Argument eine moralische Rechtfertigung für das Töten der Tiere unter diesen Umständen.

Es ist fragwürdig, inwieweit die Tiere unter diesen Umständen ein glückliches Leben haben. So erfahren Tiere, auch in der Tierhaltung die als tierfreundlich gilt, körperliches und emotionelles Leid. Auch ist es fragwürdig ob das Töten ohne negative Nebeneffekte stattfindet, in vielen Fällen sicherlich nicht. Prinzipiell sind allerdings Umstände denkbar, in denen die Bedingungen des Ersetzbarkeits-Argumentes erfüllt sind.

Wenn Tiere unter den genannten Bedingungen ersetzbar wären, gelte dies dann auch für Menschen? Peter Singer, einer der bekanntesten aktuellen Philosophen, der eine utilitaristische (Tier-) Ethik vertritt, probiert den Umfang des Ersetzbarkeits-Argumentes einzuschränken. Singer ist bereit um weitgehende Anpassungen an seiner Moraltheorie anzubringen, um die Implikation zu vermeiden, dass normale erwachsene Menschen ersetzbar seien. Schließlich wird aber deutlich, dass allerlei Strategien um den Umfang des Ersetzbarkeits-Argumentes zu beschränken nicht erfolgreich sind. Mithilfe von Singer's neuester Strategie, die davon ausgeht, dass das Wohlergehen nur neutral oder negativ sein kann, wäre es möglich manche Wesen von der Ersetzbarkeit auszuschließen. Dies betrifft diejenigen Wesen, deren Leben in negativem Wohlergehen resultiert. Singer's Strategie führt jedoch zu der Schlussfolgerung, dass genau die Wesen, die dann nicht ersetzbar wären, notwendigerweise negatives Wohlergehen erfahren würden, und darum besser überhaupt nicht gelebt haben sollten.

5. *Total View* gegenüber *Prior Existence View*

Das Ersetzbarkeits-Argument beruht auf einer bestimmten Theorie über die Frage, wessen Wohlergehen mitzählt. Das Argument kann nur vertreten werden, wenn man beim Beantworten der Frage, ob man ein bestimmtes Tier töten darf, bereits das mögliche zukünftige Wohlergehen des möglichen zukünftigen Tieres, welches das getötete Tier ersetzen soll, mitberücksichtigt. Dieses mögliche zukünftige Tier gibt es noch nicht, wenn die Entscheidung über das Töten des lebenden Tieres getroffen wird. Ob es dieses mögliche Tier jemals geben wird, hängt davon ab, wie die Entscheidung über das Töten des lebenden Tieres ausfällt. Darum ist dieses mögliche zukünftige Tier ein sogenanntes „kontingentes“ Wesen. Es gibt dieses Wesen noch nicht, und ob es je leben wird hängt von der Entscheidung die man gerade treffen will ab.

In Bezug auf den Utilitarismus ist es strittig, ob das Wohlergehen von solchen kontingenten Wesen mitzählen soll, wenn man den Wert des Resultats einer Handlung bestimmt. Gemäß des *Total View*, soll das Wohlergehen von kontingenten Wesen durchaus mit berücksichtigt werden. *Total View* besagt, dass das Wohlergehen von allen lebenden Wesen und allen möglichen Wesen berücksichtigt werden soll. Gemäß des *Prior Existence View*, soll das Wohlergehen von kontingenten Wesen nicht mitzählen. *Prior Existence View* berücksichtigt nur das Wohlergehen von Wesen, die es im Moment der moralischen Fragestellung bereits gibt, oder die es unabhängig von der Entscheidung die man gerade treffen muss, geben wird.

Ein Beispiel kann das veranschaulichen. Nehmen wir an, dass ein Paar vor der Frage steht, ob es ein Kind bekommen soll. Das Paar selbst würde mit einem Kind weniger glücklich sein, da es im Nachgehen ihrer liebsten Freizeitbeschäftigung beeinträchtigt würde. Das Kind selbst würde jedoch ein glückliches Leben haben. Es würde unter günstigen

Umständen aufwachsen, und gut versorgt werden. Gemäß des *Total View* muss bei der Entscheidung des Paares das mögliche Wohlergehen des möglichen Kindes mit berücksichtigt werden. Es könnte darum sein, dass der Verlust an Wohlergehen für die Eltern kompensiert wird durch das Wohlergehen des Kindes. Dem *Prior Existence View* zufolge zählt in diesem Fall jedoch nur das Wohlergehen des Paares.

Der *Prior Existence View* unterstützt das Ersetzbarkeits-Argument nicht, da das Wohlergehen des möglichen zukünftigen Tieres nicht mitzählt. Es ist daher gemäß dieser Theorie nicht möglich um den Verlust an Wohlergehen durch das Töten eines Tieres zu kompensieren. Die Akzeptanz de Ersetzbarkeits-Argumentes, und somit die utilitaristische Verteidigung tierfreundlicher Tierhaltung hängt also davon ab, ob man sich beruft auf die *Total View* oder auf die *Prior Existence View*.

Es muss aber geklärt werden, ob der *Prior Existence View* überhaupt eine kohärente utilitaristische Theorie ist. Es scheint nämlich nicht zum Utilitarismus zu passen, das Wohlergehen mancher Wesen nicht mit zu berücksichtigen. Der Utilitarismus geht doch davon aus, dass das Wohlergehen aller Beteiligten in gleicher Weise mit zählt. Der *Prior Existence View* ist nur unter zwei Bedingungen zu vereinbaren mit den Ausgangspunkten des Utilitarismus. Wenn man davon ausgeht, dass Folgen von Handlungen, abhängig von dem Maß indem sie das Wohlergehen der Beteiligten schaden oder fördern beurteilt werden sollen, und, wenn man zudem davon ausgeht, dass das „ins Leben bringen“ eines Wesens dessen Wohlergehen nicht verbessern oder verschlechtern kann, dann ist es sinnvoll das Wohlergehen kontingenter Wesen außer Acht zu lassen. In den Kapiteln 6&7 werden beide Bedingungen weiter erläutert und verteidigt.

6.Fördert das ins Leben bringen eines Wesens dessen Wohlergehen?

Natürlich kann das Leben von jemandem gut, schlecht oder neutral sein, was das Wohlergehen betrifft. Dennoch ist die Frage, ob das ins Leben bringen eines Wesens dessen Wohlergehen fördern kann, schwer zu beantworten. Wenn man wissen möchte, ob etwas das Wohlergehen von jemandem fördert, dann ist es üblich, das Wohlergehen des betreffende Wesens nach dem Geschehen zu vergleichen mit dem Wohlergehen davor. Noch besser scheint es, um das Wohlergehen des Wesens nach dem Geschehen zu vergleichen mit der Situation in der das Geschehene nicht stattgefunden hätte. Letzteres ist oft schwer einzuschätzen: was wäre zum Beispiel passiert, wenn ich einen bestimmten Job nicht angenommen hätte? Wenn es darum geht, jemanden ins Leben zu bringen, ist es deutlich, was die Alternative wäre: das betreffende Wesen würde es nicht geben. Es ist also notwendig, das Wohlergehen eines Wesens das es gibt zu vergleichen mit dem Wohlergehen dieses Wesens, wenn es dieses Wesen nicht gegeben hätte.

Kann man das Leben und das nicht-leben eines Wesens überhaupt vergleichen, im Hinblick auf das Wohlergehen dieses Wesens? Es ist unstrittig, dass ein Wesen, das es nicht gibt kein Wohlergehen haben kann. Eine Position ist, dass die Abwesenheit von positivem und negativem Wohlergehen im Falle des Nicht-Lebens, angemerkt werden kann als ein neutrales Niveau von Wohlergehen. Dieser Position wird jedoch entgegen gestellt, dass die Abwesenheit von Wohlergehen nicht verwechselt werden sollte mit neutralem Wohlergehen. Vergleichbare Diskussionen gibt es auch im Bereich der Quantenphysik, zum Beispiel über die Frage ob ein Photon eine Masse von Null hat, oder keine Masse. Das Aufzeigen der Struktur der betreffenden Diskussionen verdeutlicht die aktuelle Debatte zur Frage ob nicht-leben ein neutrales oder eher kein Wohlergehen beinhaltet, und zeigt, dass die Position dass die Abwesenheit von Wohlergehen etwas anderes ist, als neutrales Wohlergehen zu verteidigen ist.

Sollte man doch davon ausgehen wollen, dass das Nicht-Leben einen neutralen Wert hat im Hinblick auf Wohlergehen (anstelle von keinem solchem Wert), dann stellt sich die Frage, für wen das Nicht-Leben einen neutralen Wert haben sollte. Im Falle des Nicht-Lebens gibt es nämlich niemanden, für den das Nicht-Leben einen Wert haben könnte, nicht einmal einen neutralen Wert. Um das Leben und das Nicht-Leben eines Wesens vergleichen zu können im Hinblick auf Wohlergehen, müssten einige logische Gesetze, die normalerweise gelten für Vergleichen, über Bord geworfen werden. Es wird behauptet, dass das von Bord werfen dieser Gesetze ohnehin notwendig ist, um sagen zu können, dass das Retten eines Lebens gut ist im Hinblick auf das Wohlergehen dieser Person.. Dies ist jedoch nicht der Fall. Solche Aussprachen kann man auch machen, ohne dem Tod einen neutralen Wert in Bezug auf Wohlergehen zu geben. Es genügt, um zu sagen, dass ein langes glückliches Leben mehr Wohlergehen für die betreffende Person beinhaltet als ein kurzes glückliches Leben. Man vergleicht also zwei Leben bezüglich Wohlergehen, und nicht das Leben mit dem Tod. Es erscheint darum ad hoc, um die logischen Gesetze bezüglich Wohlergehen doch über Bord zu werfen für Vergleiche zwischen Leben und Nicht-Leben. Die Position, die besagt, dass man das Wohlergehen eines Wesens nicht fördern oder ihm schaden kann, indem man das Wesen ins Leben bringt, scheint darum eine mögliche Position zu sein. Damit ist die erste Bedingung, welche die *Prior Existence View* bracht um eine kohärente utilitaristische Theorie zu sein, erfüllt.

7. *Person-Affecting Restriction* und Non-Identitätsproblem

Die zweite Bedingung die erfüllt sein muss, damit *Prior Existence View* eine kohärente utilitaristische Theorie sein kann, ist eine bestimmte Theorie über das Beurteilen van Handlungsfolgen zu akzeptieren. Handlungsfolgen werden in Bezug auf das Wohlergehen, das sie beinhalten beurteilt; aber wie genau soll das geschehen? Soll es

dabei um die Menge des Wohlergehens als solches gehen, oder um das Maß, in dem das Wohlergehen von Wesen gefördert oder ihm geschadet wird. Letzteres besagt die *Person-Affecting Restriction*. Es sollte dieser Theorie zufolge darum gehen, wie Wesen in ihrem Wohlergehen getroffen werden. Eine Handlungsfolge kann nur dann besser (oder schlechter) sein als andere, wenn sie besser (oder schlechter) ist *für* jemanden. Die Akzeptanz dieser Theorie über das Beurteilen von Handlungsfolgen ist nötig, damit die Prior Existence View eine kohärente utilitaristische Theorie sein kann. Die Theorie die im Gegensatz dazu nur auf die totale Menge des Wohlergehens achtet, unabhängig davon was dies bedeutet für das fördern oder schaden von Wohlergehen einzelner, heißt *Impersonal View*.

Das wichtigste Argument, das gegen die *Person-Affecting Restriction* eingebracht ist, ist bekannt als das Non-Identitätsproblem. Das kann veranschaulicht werden mit Hilfe des Falles des vierzehnjährigen Mädchens, welches sich fragt, ob es jetzt ein Kind bekommen solle, oder erst in ein paar Jahren. Angenommen, das Mädchen könnte jetzt noch nicht sehr gut für ein Kind sorgen, und das Kind dadurch einen schweren Start ins Leben haben und weiterführende Probleme im Erwachsenenalter. Im Übrigen wird in diesem Fall davon ausgegangen, dass das Kind trotzdem ein Leben haben würde, das im Ganzen gesehen positives Wohlergehen befasst. Wenn das Mädchen später ein Kind bekäme, hätte es selbst ein glücklicheres Leben. Eine weitere Annahme ist, dass das Wohlergehen von niemandem sonst hierbei getroffen ist. Dem Mädchen selbst ginge es in beiden Optionen gleich gut. Es scheint in diesem Fall, insbesondere für Utilitaristen, richtig zu sein um mit dem Kinder kriegen zu warten. Das würde nämlich niemandem schaden, und das Kind wäre glücklicher. Das Problem ist jedoch, dass das Mädchen, falls es warten würde, nicht dasselbe Kind bekäme,, sondern ein anderes, nicht identisches Kind. Es wäre also für kein einziges Wesen besser, zu warten. Dies ist das Non-Identitätsproblem, das vorkommt wenn die im Wohlergehen getroffenen in den verschiedenen Handlungsfolgen nicht-

identische Individuen sind. Wenn ein Resultat nur besser sein kann, wenn es besser ist für jemand anderen, dann scheint es in diesem Fall nicht besser zu sein, wenn das Mädchen wartet. Das erscheint als gegenintuitiv.

Das Non-Identitätsproblem kann umgangen werden, wenn man die *Person-Affecting Restriction* in einem weiteren Sinne auffasst. Dieser „weiten“ Auffassung zufolge geht es bei der Bedingung, dass ein Resultat nur besser sein kann, wenn es besser *für* ein Wesen ist, nicht um ein spezifisches Individuum, sondern um Wesen, wer sie auch sein mögen. In diesem weiten Sinne ist Resultat A besser als Resultat B wenn Resultat A besser ist für die A-Wesen als Resultat B für die B-Wesen. Diese weite Auffassung der *Person-Affecting Restriction* kann verteidigt werden, indem man aufzeigt, dass „Identität“ ein Begriff ist, der auf verschiedene Weise verstanden werden kann, und, dass in verschiedenen Kontexten verschiedene Definitionen von Identität relevant sind. Wenn es um die moralischen Fragen geht, die in Non-Identität Fällen, wie der Frage des vierzehnjährigen Mädchens relevant sind, dann erscheint eine genetische Definition von Identität nicht die relevanteste. In solchen Fällen erscheint eine *de-dicto* Definition davon, was es bedeutet um „besser für jemanden“ zu sein, relevanter. Dieser Definition zufolge gilt: Wenn man zwei Resultate O1 und O2 vergleicht, dann ist Resultat O1 *de-dicto* besser für S wenn es demjenigen, der S in O1 ist besser geht in O1 als es demjenigen der S in O2 ist geht in O2. Im Falle des vierzehnjährigen Mädchens bedeutet dies, in Kombination mit dem Utilitarismus, dass das Mädchen warten muss, da ihr nächstes Kind – wer das auch sein möge – mehr Wohlergehen hat wenn sie wartet.

8. *Repugnant Conclusion* und *Expected Misery Argument*

Da sowohl *Total View* wie auch *Prior Existence View* also kohärente utilitaristische Theorien sind, ist es interessant um einige Implikationen von beiden Theorien zu untersuchen. Als die berüchtigtste Implikation

des *Total View* gilt die *Repugnant Conclusion*, die abstoßende Schlussfolgerung. Nehmen wir einmal an, es gibt eine Bevölkerung die besteht aus mindestens 10 Millionen glücklichen Menschen. Wenn diese Menschen mehr Kinder bekämen, könnte diese Bevölkerung noch umfangreicher werden. Dieses Wachstum würde stets mehr auf Kosten der bereits lebenden Menschen gehen, sodass auch die neuen Menschen nicht sehr glücklich sein würden. Schließlich könnte Ausbreitung resultieren in einer sehr umfangreichen Bevölkerung, viele Male umfangreicher als die ursprüngliche Bevölkerung, wobei ein jeder ein Niveau von Wohlbefinden hat, das gerade etwas über dem neutralen Niveau liegt. Wenn diese Bevölkerung groß genug ist, dann könnte diese Bevölkerung mehr Wohlbefinden befassen als die ursprüngliche Bevölkerung. Gemäß des *Total View* wäre es dann eine gute Sache, wenn die ursprüngliche Bevölkerung auf diese Weise die Summe des Wohlbefindens vergrößern würde. Dies scheint eine unattraktive Implikation des *Total View* zu sein.

Prior Existence View kann diese Implikation vermeiden, da er das mögliche zukünftige Wohlergehen kontingenter Wesen nicht berücksichtigt. Der Verlust des Wohlergehens der bereits lebenden Wesen, würde weitere Ausbreitung eindämmen. Genau die Implikation, dass das Wohlergehen kontingenter Wesen nicht mit zählt, kann in anderen Situationen jedoch als ein starkes Argument gegen den *Prior Existence View* angeführt werden. Das bekannteste Beispiel hierfür ist der Fall des unglücklichen Kindes: Ein Paar weiß bereits bevor der Erzeugung eines Kindes, dass ein jedes Kind das sie bekommen könnten ein kurzes, leidvolles Leben haben würde, mit einem deutlich negativen Niveau des Wohlbefindens. Es scheint, dass dies ein Grund sein sollte, um kein Kind zu bekommen. Der *Prior Existence View* scheint diesen Grund jedoch nicht berücksichtigen zu können, da es hier um ein kontingentes Wesen geht, dessen Wohlergehen nicht mit zählt.

Nun kann man sich fragen, wie realistisch dieser Fall ist. Würde das Paar ein solches Kind überhaupt zeugen wollen, und hätte dies

wirklich keine Folgen für das Wohlergehen bereits Lebender, oder zukünftiger Wesen, die es, unabhängig von dieser Entscheidung, geben wird? Doch angenommen dem Paar würde es gefallen, ein solches Kind zu erzeugen. Der *Prior Existence View* kann das Wohlergehen dieses Kindes erst dann berücksichtigen, sobald es das Kind gibt, als ein fühlendes Wesen. Sobald das Kind tatsächlich nur Schmerzen erfährt, und mit Sicherheit keine Besserung möglich ist, wäre Abtreibung oder Euthanasie einer jeden Version des Utilitarismus gemäß, *ceteris paribus*, geboten. Aber sogar schon bevor es das Kind gibt, als ein fühlendes Wesen, kann das Vorhaben des Paares um das Kind zu bekommen und leiden zu lassen aus utilitaristischen Gründen verurteilt werden. Dies gilt auch wenn man vom *Prior Existence View* ausgeht. Der Utilitarismus kann nämlich auch den Charakter des Paares beurteilen in Bezug auf die Neigung, das Wohlergehen zu maximieren. Ein Charakter, der zu vereinbaren ist mit dem Erzeugen und Leiden lassen eines solchen Kindes, ist dem Utilitarismus zufolge ein schlechter Charakter. Der *Prior Existence* Utilitarismus hat also doch durchaus die Möglichkeit, einen solchen Fall moralisch zu verurteilen. In diesem Zusammenhang ist es interessant um anzumerken wie Utilitaristen umgehen mit gegenintuitiven Implikationen moralischer Theorien. Obwohl Utilitaristen den Ruf haben, diese nicht wichtig zu nehmen, ist das nicht unbedingt der Fall.

9. Veganismus gegenüber tierfreundlicher Tierhaltung

Welche Implikationen haben *Total* Utilitarismus und *Prior Existence* Utilitarismus für die Tierhaltung? Ein Argument zugunsten tierfreundlicher Tierhaltung besagt, dass diese gut ist für die Tiere. Die Tiere haben nämlich ein glückliches Leben, und hätten ansonsten überhaupt nicht gelebt. Ein glückliches Leben, auch wenn es frühzeitig beendet wird, ist immerhin besser als überhaupt nicht zu leben, so lautet das Argument. Dieses Argument, bekannt als die Logik des

Vorradschranks (*Logic of the Larder*) besagt, dass der echte Tierliebhaber ein Fleischesser ist.

Dieses Argument kann aus praktischen Gründen kritisiert werden. Selbst wenn man so viel wie möglich glückliche Tiere ins Leben bringen wollte, dann sollte man keine tierischen Produkte produzieren oder konsumieren. Es können mit pflanzlicher Nahrung im allgemeinen mehr Menschen ernährt werden, und auf dem Platz, der dadurch frei kommen würde, könnten wilde Tiere leben. Anstelle eines Stück Fleisches, könnte man sich dann besser ein paar glückliche Mäuse anschaffen. Und wäre diese Logik auch in Bezug auf Menschen an zu wenden? Auf einem tieferen Niveau geht dieses Argument davon aus, dass man das Wohlergehen eines Wesens fördern kann, indem man dieses Wesen ins Leben bringt. Wie wir in Kapitel 6 gesehen haben, verwirft der *Prior Existence* Utilitarismus diese Annahme, und ist darum nicht zu vereinbaren mit diesem Argument. *Total* Utilitarisme kann im Prinzip damit einstimmen, dass das Wohlergehen eines Tieres gefördert werden kann, indem man das Tier ins Leben bringt. Auch wenn der Total Utilitarist dies nicht findet, dann kann er oder sie trotzdem sagen, dass es die totale Menge des allgemeinen Wohlergehens vergrößert, wenn man ein glückliches Tier ins Leben bringt.

Da der *Prior Existence* Utilitarismus das Ersetzbarkeits-Argument nicht akzeptiert, wird diese Version des Utilitarismus das routinemäßige Töten von Tieren in der Tierhaltung nicht verteidigen. Veganer Landbau, ohne Tierhaltung, passt als Ideal besser zum *Prior Existence* Utilitarismus. *Total* Utilitarismus dagegen akzeptiert doch das Ersetzbarkeits-Argument, mit allen Implikationen die dies, im Prinzip auch für Menschen, hat. Wenn man dann manche Wesen als unersetzbar anmerken möchte, so wie Singer es vorschlägt, dann impliziert dies, dass diese Wesen nur negatives Wohlergehen haben können und darum besser nicht gelebt hätten.

10. Schlussfolgerung

Wie der Utilitarismus urteilt über das recht allgemein akzeptierte Ideal der tierfreundlichen Tierhaltung, und über andere Umstände, in denen glückliche Tiere getötet und durch andere ersetzt werden, hängt davon ab, welche Version des Utilitarismus man vor Augen hat. *Prior Existence* Utilitarismus ist neben dem *Total* Utilitarismus auch eine kohärente Version dieser Moraltheorie. *Prior Existence* Utilitarismus bietet Menschen und (anderen) Tieren einen besseren Schutz gegen das Töten. Außerdem geht es dabei um eine Version des Utilitarismus, die weniger fassbar ist für einen zentralen Kritikpunkt der gegen den Utilitarismus erhoben wird, nämlich, dass der Utilitarismus zu unpersönlich sei und sich nur um das Wohlergehen als eine abstrakte Größe bekümmere. Der *Prior Existence* Utilitarismus bekümmert sich um die Folgen für das Wohlergehen von fühlenden Wesen.

Im Gegensatz zu geläufigen Behauptungen geht es beim Utilitarismus nicht unbedingt nur um das Vermeiden von unnötigem Leiden, sondern hat der Utilitarismus auch etwas einzubringen gegen das Töten von Tieren. Das ist relevant für sogenannt tierfreundliche Formen der Tierhaltung, und für andere Formen des Tiergebrauches.

CURRICULUM VITAE

Tatjana Visak was born in Giessen (Germany) on December 12, 1974. She studied Political Sciences at Leiden University, the Netherlands, and graduated in 1999 with a master's thesis about animal ethics. After her graduation, she worked during one year as a junior researcher at Leiden University, Department of Political Sciences. From 2000 until 2005, she worked as a junior researcher/ lecturer at the Ethics Institute of Utrecht University, and from 2002 until 2003 also as a researcher at the Ethics Department of the Erasmus Medical Center of Erasmus University, Rotterdam. From 2005 until 2010 she worked as a PhD researcher at the Ethics Institute of Utrecht University, on a dissertation on utilitarian ethics. Since then, she works as a lecturer at the Philosophy Department of Leiden University and at the International School of Philosophy in Leusden, in collaboration with the Ethics Institute of Utrecht University, as the organizer of the 2012 Minding Animals Conference, which will be hosted by Utrecht University.

Quaestiones Infinitae

PUBLICATIONS OF THE DEPARTMENT OF PHILOSOPHY

- VOLUME 21. D. VAN DALEN, *Torens en Fundamenten* (valedictory lecture), 1997.
- VOLUME 22. J.A. BERGSTRA, W.J. FOKKINK, W.M.T. MENNEN, S.F.M. VAN VLIJMEN, *Spoorweglogica via EURIS*, 1997.
- VOLUME 23. I.M. CROESE, *Simplicius on Continuous and Instantaneous Change* (dissertation), 1998.
- VOLUME 24. M.J. HOLLENBERG, *Logic and Bisimulation* (dissertation), 1998.
- VOLUME 25. C.H. LEIJENHORST, *Hobbes and the Aristotelians* (dissertation), 1998.
- VOLUME 26. S.F.M. VAN VLIJMEN, *Algebraic Specification in Action* (dissertation), 1998.
- VOLUME 27. M.F. VERWEIJ, *Preventive Medicine Between Obligation and Aspiration* (dissertation), 1998.
- VOLUME 28. J.A. BERGSTRA, S.F.M. VAN VLIJMEN, *Theoretische Software-Engineering: kenmerken, faseringen en classificaties*, 1998.
- VOLUME 29. A.G. WOUTERS, *Explanation Without A Cause* (dissertation), 1999.
- VOLUME 30. M.M.S.K. SIE, *Responsibility, Blameworthy Action & Normative Disagreements* (dissertation), 1999.
- VOLUME 31. M.S.P.R. VAN ATTEN, *Phenomenology of choice sequences*

- (dissertation), 1999.
- VOLUME 32. V.N. STEBLETSOVA, *Algebras, Relations and Geometries (an equational perspective)* (dissertation), 2000.
- VOLUME 33. A. VISSER, *Het Tekst Continuüm* (inaugural lecture), 2000.
- VOLUME 34. H. ISHIGURO, *Can we speak about what cannot be said?* (public lecture), 2000.
- VOLUME 35. W. HAAS, *Haltlosigkeit; Zwischen Sprache und Erfahrung* (dissertation), 2001.
- VOLUME 36. R. POLI, *ALWIS: Ontology for knowledge engineers* (dissertation), 2001.
- VOLUME 37. J. MANSFELD, *Platonische Briefschrijverij* (valedictory lecture), 2001.
- VOLUME 37A. E.J. BOS, *The Correspondence between Descartes and Henricus Regius* (dissertation), 2002.
- VOLUME 38. M. VAN OTEGEM, *A Bibliography of the Works of Descartes (1637-1704)* (dissertation), 2002.
- VOLUME 39. B.E.K.J. GOOSSENS, *Edmund Husserl: Einleitung in die Philosophie: Vorlesungen 1922/23* (dissertation), 2003.
- VOLUME 40. H.J.M. BROEKHUIJSE, *Het einde van de sociaaldemocratie* (dissertation), 2002.
- VOLUME 41. P. RAVALLI, *Husserls Phänomenologie der Intersubjektivität in den Göttinger Jahren: Eine kritisch-historische Darstellung* (dissertation), 2003.
- VOLUME 42. B. ALMOND, *The Midas Touch: Ethics, Science and our Human Future* (inaugural lecture), 2003.
- VOLUME 43. M. DÜWELL, *Morele kennis: over de mogelijkheden van toegepaste ethiek* (inaugural lecture), 2003.

- VOLUME 44. R.D.A. HENDRIKS, *Metamathematics in Cog* (dissertation), 2003.
- VOLUME 45. TH. VERBEEK, E.J. BOS, J.M.M. VAN DE VEN, *The Correspondence of René Descartes: 1643*, 2003.
- VOLUME 46. J.J.C. KUIPER, *Ideas and Explorations: Brouwer's Road to Intuitionism* (dissertation), 2004.
- VOLUME 47. C.M. BEKKER, *Rechtvaardigheid, Onpartijdigheid, Gender en Sociale Diversiteit; Feministische filosofen over recht doen aan vrouwen en hun onderlinge verschillen* (dissertation), 2004.
- VOLUME 48. A.A. LONG, *Epictetus on understanding and managing emotions* (public lecture), 2004.
- VOLUME 49. J.J. JOOSTEN, *Interpretability formalized* (dissertation), 2004.
- VOLUME 50. J.G. SIJMONS, *Phänomenologie und Idealismus: Analyse der Struktur und Methode der Philosophie Rudolf Steiners* (dissertation), 2005.
- VOLUME 51. J.H. HOOGSTAD, *Time tracks* (dissertation), 2005.
- VOLUME 52. M.A. VAN DEN HOVEN, *A Claim for Reasonable Morality* (dissertation), 2006.
- VOLUME 53. C. VERMEULEN, *René Descartes, Specimina philosophiae: Introduction and Critical Edition* (dissertation), 2007.
- VOLUME 54. R.G. MILLIKAN, *Learning Language without having a theory of mind* (inaugural lecture), 2007.
- VOLUME 55. R.J.G. CLAASSEN, *The Market's Place in the Provision of Goods* (dissertation), 2008.
- VOLUME 56. H.J.S. BRUGGINK, *Equivalence of Reductions in Higher-Order Rewriting* (dissertation), 2008.
- VOLUME 57. A. KALIS, *Failures of agency* (dissertation), 2009.

- VOLUME 58. S. GRAUMANN, *Assistierte Freiheit* (dissertation), 2009.
- VOLUME 59. M. AALDERINK, *Philosophy, Scientific Knowledge, and Concept Formation in Geulincx and Descartes* (dissertation), 2010.
- VOLUME 60. I.M. CONRADIE, *Seneca in his cultural and literary context: Selected moral letters on the body* (dissertation), 2010.
- VOLUME 61. C. VAN SIJL, *Stoic Philosophy and the Exegesis of Myth* (dissertation), 2010.
- VOLUME 62. J.M.I.M. LEO, *The Logical Structure of Relations* (dissertation), 2010.
- VOLUME 63. M.S.A. VAN HOUTE, *Seneca's theology in its philosophical context* (dissertation), 2010.
- VOLUME 64. F.A. BAKKER, *Three Studies in Epicurean Cosmology* (dissertation), 2010.
- VOLUME 65. T. FOSSEN, *Political legitimacy and the pragmatic turn* (dissertation), 2011.
- VOLUME 66. T. VISAK, *Killing happy animals. Explorations in utilitarian ethics.* (dissertation), 2011.