

Lund University

Department of Philosophy

FPRK01:2 Practical Philosophy: Level 3

Bachelor Thesis (15 credits)

2011-01-07

Author: Simon Rosenqvist

Supervisor: Dan Egonsson

A Utilitarian Dilemma On Diachronic Preferences

Abstract

When preferences from the past or future disagrees with present preferences about what to do now, it seems counterintuitive for preference utilitarianism to take these past and future diachronic preferences into account. In this thesis I discuss two different solutions to this problem. First, preference utilitarianism could solve the problem by becoming more sophisticated - by suggesting that we should give the diachronic preferences equal weight to other preferences, but that we in order to fulfil as many preferences as possible at all times should give the diachronic preferences less consideration in our daily lives. Second, preference utilitarianism could exclude or downgrade the weight of certain diachronic preferences. I argue that it is not clear how the second solution could solve the problem of counterintuitiveness without becoming sophisticated to at least some degree. I conclude that, since we anyway need to become sophisticated, we could just as well accept the first view from the start – treating the problem of diachronic preferences no different than other similar allegations of utilitarianism being counterintuitive.

Keywords: Diachronic, preferences, sophisticated, utilitarianism.

Contents

1. Introduction.....	1
2. Preferences.....	2
3. Sophisticated Utilitarianism.....	3
4. The Problem of Diachronic Preferences.....	7
5. Defending the Simple View: Past and Future Preferences.....	8
6. Defending the Simple View: Changing Preferences.....	14
7. Criticising the Complex View: Excluding Preferences.....	19
8. Criticising the Complex View: Downgrading Preferences.....	21
9. Conclusion.....	23

1. Introduction

When preferences from the past or future disagrees with present preferences about what to do now, it seems counterintuitive for preference utilitarianism to take these past and future diachronic preferences into account.¹ This suggests a problem for a certain normative theory – preference utilitarianism – which holds that we should maximize the fulfilment of all preferences.

This thesis discusses a dilemma between two different solutions to this problem. First, a *simple view* which holds that all preferences should be given equal weight. Second, a *complex view* which holds that we can exclude or downgrade the weight of some preferences. I argue that, while the complex view provides the perhaps most obvious route of dealing with the problematic diachronic preferences, the preference utilitarian should instead opt for the simple view, thus giving equal weight to all preferences. My argument for this is twofold:

First, (1) I argue that the simple view can be defended through utilitarianism becoming more *sophisticated* – by pointing to how we, if we give the preferences of past and future persons full consideration in our daily lives, would lower total preference fulfilment. This is similar to how the utilitarian would explain why he or she, even when calculating that it would lead to greater overall preference fulfilment, should not consider murder someone. Instead, we should point to the indirect negative effects of people going around trying to calculate when to murder people.

Second, (2) I argue that it is not clear how the complex view could be appropriately justified and formulated to solve the problem of counterintuitive diachronic preferences, at least not without becoming sophisticated to at least some degree. And if we anyway are to become sophisticated, I argue, we could just as well accept the simple view from the start.

The disposition of the thesis goes as follows: First, I present a general view on preference and combine it with a rough account of sophisticated utilitarianism (Section 2 and 3). In then examine the problem of diachronic preferences and single out two different groups of preferences for further study (Section 4). The rest of the thesis is concerned with defending the simple view (Section 5 and 6), and criticising the complex view (Section 7 and 8). In the end I sum up the main points, concluding that both the simple and complex views are possible as solutions to the diachronic preferences, but that the simple view is the preferable solution (Section 9).

¹ That past preferences are problematic in this way have been suggested by Brandt (1979), p 249-251, Bykvist (2003), p 115, and Rønnow-Rasmussen (1998), section 1, et al. Note that, since the paper by Rønnow-Rasmussen is not paginated because of it only being available online, I instead refer to the different sections in his text.

2. Preferences

What is a preference? On the account assumed here, a preference first of all requires a preferrer – someone that prefers something. It would seem strange to claim the existence of a preference without being able to pinpoint who has the preference in question. I do not take a stand on issues on what kinds of creatures can have preferences or not. It is enough for my purposes here to simply denote such preferrers, when necessary, as *persons*, without giving any more specific description of what it means to be a person.

Further, there must be the preferred object – the thing which the person prefers. I here follow Krister Bykvist in his identification of "preference objects with bare state of affairs".² Once again, this seems to be required by the notion of a preference. If I tell you "I have a preference" and you ask "for what?", you would not accept "nothing, I just have a preference" as an answer. There needs to be some state of affairs in the world that we prefer to be realised.

Both preferences and their preferred objects are located in time – they have a temporal location. Assuming this is true, we have two different kinds of preferences to take into consideration. When the preference and the preference object are located in the same temporal interval the preference is *synchronic*. When the preference and the preference object are located in different temporal intervals they are *diachronic*.

What is a temporal interval? I do not believe it to be relevant for this project to decide on a fixed length for each preference. What is important are the differences between groups of diachronic and synchronic preferences, not the exact duration of each preference. Similar to Bykvist I therefore adapt the minimal view that time "can be split up into an exhaustive and exclusive set of intervals of minimal duration [...] minimal in the sense of being just long enough that it makes sense to say that someone wants something during one such interval".³ If this is true, we could place preference objects and preferrers into both the same, and different, intervals, thus differing between the two kinds of preferences.

Finally, I distinguish between preferences being fulfilled or not fulfilled (or "frustrated"). A fulfilled preference is one which has its object realised. For example, your preference for ice cream in the present situation is fulfilled if the object "I have an ice cream in the present situation" is realised.

2 Bykvist (1998), p 15.

3 Bykvist (1998), p 16.

3. Sophisticated Utilitarianism

Preference utilitarianism, on the view taken here, is described as the combined views (1) that it is the consequences that matters when we act morally, and (2) that it is the fulfilment of preferences that makes consequences valuable.

Regarding (1) utilitarianism is *consequentialist*. Only the expected consequences of our actions count when determining whether an action is right or wrong. I here take “consequences” to mean certain state of affairs. A world with more good consequences - in our case a world with more fulfilled preferences, is better than one with less good consequences – a world with less fulfilled preferences. Also, utilitarianism is here considered a *maximizing* theory: We should attempt to create the best possible world, as opposed to being satisfied with a world that is “good enough”.

Regarding (2), on what is valuable, I assume that utilitarianism is *universalizable* – that all preferences by all persons, at all times, are equally valuable and included in morality. This includes *external* preferences – “a preference for a state of affairs other than the experiences of the preferer”.⁴ For example, we should value fulfilling someone's preference for the preservation of the rain forest even if the person will not know whether the rain forest is actually cut down or not.⁵

The view including all of the above preferences, and giving them equal weight, will here be referred to as a *simple view*. If we, on the contrary, have reservations against including all kinds of preferences, or if we wish to give some of them different weight, we could favour a less inclusive view. I refer to such a view as a *complex view*.

The two different views can be stated in the following way:

The Simple View (SV): All preferences should count equally, regardless of whether they are diachronic or synchronic.⁶

The Complex View (CV): Not all preferences should count equally. Some preferences should either be excluded, or be given less weight.⁷

4 Rønnow-Rasmussen (1998), section 3.

5 Two different views on the value of preferences are sometimes mentioned – one *object* and one *satisfaction* variant, Rabinowicz (1996), p 1-27. I do not believe this distinction to be relevant for the questions considered here. However, if questioned on this topic I would subscribe to the satisfaction variant, for reasons given by Rønnow-Rasmussen (1998), section 2.

6 The distinction between a simple and a complex view can be compared to the discussion on “purity” in utilitarianism. For a discussion on such a pure model, see for example Bykvist (1998), p 37-42, and Egonsson (1990), p 32-60.

7 The possibility of solving the problem of past preferences through giving different weight to them, as compared to excluding them, have been discussed by Rønnow-Rasmussen (1998), section 3. It should be noted that I here take the exclusion of irrational preferences to be a version of the complex view.

What does it mean that a view of utilitarianism is “sophisticated”? Let us consider an example by William Shaw:⁸ Imagine that we are to set up a system of traffic management according to a utilitarian normative theory – in this thesis a variant aimed at maximising preference fulfilment. Should we, in such a system, have only a single rule demanding of us that we “drive as to maximize preference fulfilment”? It seems likely that such a principle would lead to immense problems. For instance - how do we know what side of the road to drive on - or when to stop for pedestrians? Shaw writes:

Sophisticated utilitarians know that the motivations and dispositions they have and the rules they follow are for the best; they know that the most sensible way for them (and others) to live precludes their trying to maximize utility with each and every action they perform.⁹

The sophisticated utilitarian realises that to maximise the total amount of good consequences, in this case the fulfilment of preferences, we cannot go around calculating what it would be best to do in every given situation. Doing so would lead to frequent mistakes and to less good total consequences. Therefore the wise preference utilitarian does not, in most situations, think in terms of preference fulfilment, but rather use a larger number of simple principles to guide his or her actions.

In the following pages I borrow a certain terminology from Richard M. Hare to describe such a position. Hare differs between two levels in moral thinking, one critical and one intuitive. On the critical level we decide and evaluate the principles that guide us in our daily lives. These principles make up the intuitive level and consist of simple rules, such as “never kill” or “do no steal”. Hare compares the two levels with a game of chess. Since it is impossible to calculate every possible move in every given situation, we would do better to follow simple principles, as, for example, “always sacrifice your peasant for a queen”. These principles are then evaluated and justified on the critical level - in this case, on the basis of whether they are likely to help us win the game, or “maximize the amount of chess-victory-state-of-affairs”.¹⁰

In the case of traffic management the utilitarian can reply that we, instead of following only a single overarching principle, should adapt a number of principles on the intuitive level, indirectly aimed at maximizing the total amount of preference fulfilment on the critical level. We would in such a system have several different rules guiding traffic management, for example stating that we

⁸ Shaw (2006), p 211-212.

⁹ Shaw (2006), p 214-215.

¹⁰ Hare (1981), p 1-64.

should always drive on one side of the road, that we should signal to other drivers when turning, as well as rules stating in a clear and unambiguous manner when we are allowed to reconsider such rules, and when we are not.

Let us now turn to the problem of counterintuitiveness. Within the study of normative ethics we are often presented with different examples that, when we apply our moral theory to them, render results that we find repulsive and that revolt our “moral intuitions”. There are in general two ways to treat such problems: Either we “bite the bullet” and accept the counterintuitive conclusions, or we try to revise our theory. Both the views discussed here are of the latter kind - they admit that some sort of response is needed to the counterintuitive examples.

How far should we go when adapting the theory to our moral intuitions? Hare states in a clear manner that “[t]he appeal to moral intuitions will never do as a basis for a moral system”.¹¹ Hare instead suggests that we base morality on linguistic rather than moral intuitions. If our moral intuitions are designed to help us follow particular principles, how can they at the same time decide the critical level?

Many are, however, sceptical of such a way of discarding our moral intuitions. I therefore opt for an alternative route here, adapting a “reflective equilibrium”, such as suggested by John Rawls, where we continuously revise both our opinions about particular examples and our more general moral theory, bringing both of them into coherence with each other.¹² I therefore take it to be possible for us to revise the critical level to escape the problematic intuitions. The complex view will do so through the exclusion or downgrading of certain preferences.

What could such a sophisticated response look like? Let us return to the example of setting up a traffic system. Imagine that you are out driving with a friend. Suddenly you arrive at a crossover. You know that the traffic rules forbid you to drive against red lights – however, you also know that the area is empty. Your friend, who knows that you are a utilitarian, is quick to remind you that you ought to maximize preference fulfilment and drive against the red lights. Yet you might find this counterintuitive. Is this a non-utilitarian feeling that we would best get rid of?

The sophisticated utilitarian would reply that our intuitions are not always utilitarian, even if we justify them with a utilitarian theory, and that we would never be able to increase preference fulfilment if all of our intuitions were utilitarian in this way. It is not strange that some actions, such as driving against red lights, seems to us counterintuitive, because if they did not we would consider breaking these rules on a regular basis. To maximize preference fulfilment on the whole, and on the long run, some principles should very seldom, if at all, be contemplated to break. This provides a

11 Hare (1981), p 12.

12 Rawls (1999), p 18, 42-45.

ready answer to accusations of utilitarianism being counterintuitive – we could point to how not trying to calculate how to maximize preference fulfilment in the particular situation indirectly leads to more total preference fulfilment.

Two objections to this strategy need to be considered. First, how do we know that a certain principle leads to more preference fulfilment than another? This seems to be a query more apt for empirical studies rather than philosophical thought experiments. Second, the defence might be too strong – could any of our moral intuitions, indeed, almost any conceivable action, be justified through pointing to how they somehow would maximize preference fulfilment? How can we ever show such a position to be false – and would not this lack of falsifiability reveal a serious flaw in the sophisticated utilitarian position?

Answering (a), I agree that empirical studies would be preferable. Until such are available, however, we are forced to take the more uncertain route of thought experiments and guesswork. This is unsatisfying. However, the sophisticated utilitarian is here forced follow his or her opponent in the tracks. At least two such empirically testable claims are already assumed by the opponent. First, that certain cases actually are counterintuitive, and that some are not. As we shall see this is far from clear in all cases. Second, mirroring the utilitarian claim, it claims that certain principles would not lead to more preference fulfilment. Both the sophisticated utilitarian and his or her opponent are here on muddy and uncertain, though even, ground.

Answering (b), I do not think that all cases of counterintuitiveness could be explained through referring to indirectly beneficial rules. For example, we could not justify killing people randomly through suggesting that such principles somehow are conducive to more preference fulfilment. The reasonableness of a certain behaviour actually increasing preference fulfilment is a matter of degree. What the sophisticated utilitarian can do is, through careful reasoning, show how it seems plausible that such principles are actually increasing preference fulfilment.

Having outlined a response to the two worries discussed above, let me suggest a more modest ambition: The sophisticated utilitarian should aim at no more than to hold the ground against its opponent - to point to how easy our intuitions change when we change certain features in examples, and to draw parallels between cases of diachronic preferences and less controversial non-diachronic preferences, resolving part of, or all of, the experienced counterintuitiveness. What could not be aimed at is a conclusive proof of how utilitarianism could give a final answer to the different problematic cases. It is of small comfort that the opponent to utilitarianism is in no better position.

I soon return to the simple and complex view. First, however, I should conduct a closer examination of the diachronic preferences. What kinds of diachronic preferences are problematic, i.e. what groups of preferences leads to counterintuitive conclusions for preference utilitarianism?

4. The Problem of Diachronic Preferences

Given our previous distinction between synchronic and diachronic preferences, and given that we differ between three temporal positions – past, present and future – we can illustrate the taxonomy as follows:

Preference / Preference object	Past	Present	Future
Past	Synchronic/ (1) Diachronic	Diachronic (2)	Diachronic (3)
Present	Diachronic (4)	Synchronic (5)	Diachronic (6)
Future	Diachronic (7)	Diachronic (8)	Synchronic/ (9) Diachronic

Figure 1, Preferences and preference objects.

Of these nine groups of preferences, eight concern diachronic preferences. Two of these groups are, I believe, problematic in that they could give counterintuitive results for the utilitarian. These are (2) past preferences for present objects, and (8) future preferences for present objects. More specifically, the problem concerns future and past preferences that conflict with (5) present preferences for present objects, for instance if a large number of past or future people preferred us to, in the present, build a monument to their glory – and if we would dislike to undertake such a project on their behalf.

Before moving on, I should briefly mention why I consider the other six groups to be less problematic:

(1), (4) and (7): The diachronic preferences with a past object are excluded on the basis of their objects being impossible to realise. For example, suppose that you wanted to eat an ice cream ten years ago at that particular moment. Regardless of whether we wish to fulfil these preferences, it seems that we are not obliged to do anything about them now – simply because we cannot change the past. These preferences, whether included in preference utilitarianism or not, are unproblematic.

It should be noted that, in rare cases, preferences with a past object could be fulfilled through their being “transported” to the present by a present preference, such as when you prefer for a past person's preference for a present object to be fulfilled. These cases, however, are rare exceptions and in a clear minority, and should regardless not cause any instances of counterintuitiveness.

(3), (9): The cases in which the preference is past and the object is future could be problematic, but the discussion on (2) past preferences for present object cover the relevant parts of these cases. The same goes for cases of future preferences for a future non-synchronic situation – these are

covered by either (2) or (8).

(6): Finally, I will not consider cases with a present preference for a future object. For instance, if someone in the present wants humanity to colonise the stars in two hundred years. These preferences appear to me at least to some extent intuitively acceptable. We often reason in terms that suggest that our preferences for future objects beyond the scope of our lives should be realised. However, I admit that I am troubled by the possibility of present preferences for future objects conflicting with future synchronic preferences, such as if the people living in the future would prefer to go on living on earth rather than fulfilling my preference for humanity exploring the stars. I am for this reason tempted to include these preferences in the subsequent discussion, but since they seem to be of a different kind, and less problematic than the other two groups, I choose to exclude them.¹³

I now move on to discuss different counterintuitive examples of (2) and (8). In the following two sections I frequently use a certain type of thought experiments, or “intuition pumps” - examples tailored specifically to bring out our moral intuitions in different cases. For clarity I have named the examples, formatting them in italic and with indents, as well as pointing out which group of preferences they concern. When the examples are borrowed from another author I refer to him or her in a footnote.

5. Defending the Simple View: Past and Future Preferences

First of all, both the simple view (from now on labelled as SV) and the complex view (CV) would agree that some preferences in (2) and (8) are not problematic because they are *conditional on their own persistence* - they include a condition that the preference persists in the present situation. That some preferences are conditional in this way has been suggested by Derek Parfit.¹⁴ For instance, imagine that a friend of yours in the past wanted you to eat ice cream in the present, but only on the condition that you now still like ice cream. If you presently have lost your taste for ice cream, you can obviously ignore this past preference. The same goes for future preferences: For example, future persons could prefer us to eat ice cream now only on the condition that we enjoy it.

Since these preferences will not clash with the present preferences they avoid the objections of counterintuitiveness. And just as synchronic preferences often appears to be tied to the experience of well-being of people, so diachronic preferences often seems to be tied to these same synchronic preferences. Imagine, for instance, that you have a preference for the preservation of the rainforest

¹³ Hare discuss the problems concerning such “now-for-then”-preferences in Hare (1981), p 104-106. Also, see Rabinowicz (1989), p 145-151, and Hare (1989), p 152-158.

¹⁴ Parfit (1984), p 151.

in the future. However, you are now informed that preserving it will lead to a large amount of synchronic preference frustration at the time. Confronted with these facts, most of us would probably at least briefly reconsider our attitude to the rain forest. If SV can claim that all diachronic preferences are conditionally tied to synchronic preferences in this way, the view is easily defended.

However, such a claim would be too strong. We may, even after having given these facts due consideration, persist in our preference that the rainforest should be preserved. Perhaps we are egoists, or we simply dislike other people. The remaining non-conditional preferences therefore continue to challenge PV. More specifically, the cases that should worry us the most are those where the past and future diachronic preferences clashes with the present synchronic ones.

Krister Bykvist discuss one such example where the diachronic preferences of future people (8) decides what we ought to do now:¹⁵

(8) The Future Religious Fanatics: In the future almost everyone will be religious fanatics. They will feel sorrow over the heretic past when many people did not believe in the right religion. Therefore they will prefer that everyone was a believer in the past.

Let us assume that the number of future people greatly outnumber the present or past generations. Further, let us assume that their preferences are in no way conditional, for example on the actual existence of a god, or us enjoying our new faith. Would not this example show it intuitively right to exclude the diachronic preferences?

Before drawing any hasty conclusions, however, we should follow Bykvist in moving the religious believers to the present (the below example is mine):

(5) The Present Religious Fanatics: Almost everyone are religious fanatics, with the exception of yourself. They feel sorry for the fact that you do not believe in the right religion. They therefore prefer that you are a believer.

As Bykvist notes, the example is still counterintuitive.¹⁶ And even if we switch the religious fanatics to something more sympathetic, for instance political views, the counterintuitiveness remains. Also, note that when varying the example we cannot go as far as saying that the future or past persons prefer something in the present on the condition that the present person appreciate it, for example

¹⁵ The example is borrowed from Bykvist (1998), p 82-83.

¹⁶ Bykvist (1998), p 82-83.

some kind of utopia. If we do, the preference is conditional in the way described earlier. To fill the purpose of the example, the preferences of the past and future persons need to be independent of the present preferences.

What could the utilitarian say about the present fanatics? Bykvist suggests that we could ban preferences that are about other people's lives and private concerns.¹⁷ In a later article he develops these thoughts and suggests that we should exclude those preferences that are not personal. A personal preference, according to Bykvist, is one “about a state of affairs that concerns the preferrer's life without entailing anything about other people”.¹⁸ This is a defence of a modified version of preference utilitarianism where some preferences have been excluded – that is, a defence of CV.

I discuss the prospect of Bykvist's variant of CV in Section 8. For now I believe SV can retain much of Bykvist's suggestion by, rather than accepting CV, suggesting a principle on the intuitive level giving less weight to preferences that are not personal in the way described above, and claim that this principle is actually maximizing preference fulfilment on the critical level. To put it simply, SV could suggest that we ignore the fanatics with good moral consciousness, because if we give such preferences about the lives of other persons equal consideration to the personal preferences on the intuitive level, we would lower total preferences fulfilment indirectly. We could therefore keep our moral intuitions in the case while at the same time defending SV.

This argument would do equally well for different kinds of similar diachronic preferences, for instance if the religious fanatics were past and preferred for us or future people to become believers.

What if we think that the preference being about other people's lives is not enough to explain the counterintuitiveness in the example? SV could respond by adding more principles. I will briefly consider two such possible suggestions by SV:

First, SV could argue that the intuitive level should include some kind of minority bias, both with regard to groups and individuals. This is how the utilitarian can reply to the allegations that it would promote the enslavement of a small group or a single individual for the benefit of the many – we should not adapt such oppressive principles, because they would lead to a society where total preference fulfilment is substantially lowered. Providing a certain amount of freedom to individuals and smaller groups, SV could argue, will maximize preference fulfilment on the whole.

Second, SV could suggest that we should include a bias towards our friends, relatives, children and others that are close to us. For instance, if we tried to give equal care to all children and not only our own, or started to give away all of our money to poor people, then the result would be a

17 Bykvist (1998), p 82-83.

18 Bykvist (2003), p 122.

lowered total amount of preference fulfilment at the critical level. In addition, we often have much less information on events far from us than those occurring in the immediate vicinity, thus actions aimed at changing what is close to us often, it could be argued, have a higher chance of success.

Both these principles could be used in the case of the "Future Religious Fanatics" to weaken their claims on us. Because they are a clear majority trying to decide over a smaller group, and because they are not close to us, our moral intuitions are correct in giving little weight to them. We might add further principles to strengthen our cause, for example that irrational beliefs should not be encouraged etc. Such principles are part of many utilitarians' standard repertoire against their opponents.

Let me, to illustrate the above argument, suggest an example of diachronic interpersonal preferences modified according to the three above principles suggested by SV. Such an example should consider (1) a close relative or friend, (2) a personal preference ranging over facts about the preferer and his or her life and (3) only the preference of a single person:¹⁹

(2) The Old Heirloom: Your mother died a few years ago. Before she died she told you she very much would like a certain heirloom to be buried at a certain place. One day you find the heirloom. You realise you could fulfil her wish without any greater burden or cost on yourself.

Note how quickly our intuitions change, and how they do so while the preference remains diachronic. It seems, at least to me, that fulfilling this preference would be quite important, perhaps even as much as the preference for a task of similar proportion belonging to another present relative. And, I suggest, it would be important not solely because you preferred her wish to be fulfilled, but because you valued her preference even after she was dead.

So far so good. Can we do the same with (8)? Consider:

(8), The Blessed Son: A couple is going to have a son. However, because of poor income and living conditions they are forced to give away the child for adoption. The son will never know his parents, but both of them are sure that he in the future would prefer that he at the time of his birth was blessed by the local religious leader. He will never know if he was blessed in this way or not.

¹⁹ The standard example in the literature of a valuable diachronic preference is a dead father's wish to be buried at a certain place, see for example Egonsson (1990), p 43, and Bykvist (1998), p 83.

It appears to me that our moral intuitions are favourable to giving weight to this diachronic preference and to ensure that the son's preference is fulfilled. And that the parents should do so not just to fulfil a synchronic preference for them to feel happy about themselves in the present moment.

Let us move on to a trickier case. Toni Rønnow-Rasmussen have suggested the following example:²⁰ Imagine two worlds. In the first world there are beings that have their synchronic preferences for pleasurable experiences fulfilled. In the second world all beings have died, but the present objects of the past preferences have been realised. When the amount of fulfilled preferences in the second world is greater in number than the amount of preferences fulfilled in the first world, it is possible that we would choose the second world – that we would choose an empty world with fulfilled past preferences. But such a position seems counterintuitive.

Could SV create a more plausible variant of the two worlds, where we would think it reasonable to choose the second world?

First we need to be assured that there are no additional reasons to choose the first world. One such reason would be that there are future synchronic or diachronic preferences to take into consideration. For instance, the second world would be considered bad by many utilitarians because it holds no hope for any future generations. To clarify the example, we should therefore remove all future preferences in both the possible worlds. Consider then:

The Last Days of Earth: For several years, the earth's population has been dedicated to gathering massive amounts of scientific data. However, because of an impending solar catastrophe the world will end in a little more than five days, and in only two days massive bursts of solar energy will kill everyone not sheltered. The limited resources available can go to either preserving the lives of almost the whole population for the next five days, until they then die from the freezing cold, or to, in seven day's time, place a monolith on the moon with the complete gathered knowledge of humanity. Right now everyone wants to finish the project and place the monolith on the moon, but everyone alive knows that if they are alive after the two days have passed they will prefer to continue their lives. The decision is up to a single individual, the president, controlling the resource allocation, and who in an instant can choose to realise either world.

20 Rønnow-Rasmussen (1998), section 4.

No one will be alive when the rocket is launched and lands on the moon, or know if they are going to die or not in two days.

The choice is between the first world where we fulfil the three additional days of synchronic experienced preferences, or the second world where we fulfil the somewhat larger amount of diachronic preferences for placing the monolith on the moon. I believe our intuitions could differ here, but I do not think it would be completely unreasonable if someone claimed that we should actually place the monolith on the moon. Of course, what is not ruled out is a variant giving some, but less, weight to the diachronic preferences. I remain hesitant of such an approach, however, for reasons that will be fully discussed in Section 8.

More could be said along these lines, and a more detailed and convincing argument could probably be provided by SV if required. I only suggest that SV can defend itself against these kinds of cases in the same way as it could defend itself against other objections against utilitarianism.

CV may concede that some examples of diachronic preferences in (2) and (8) should remain part of our morality, but still claim that we cannot account for all the counterintuitiveness of the cases. The diachronic preferences should at least be given less weight, or we should exclude some preferences. Imagine the enormous amount of preferences involved when we include all the generations now dead and those yet to be alive - surely we cannot think that all these preferences should be given equal weight to the present ones at the critical level?

SV can respond in the following way: Assume that we admit that some features tend to increase the counterintuitiveness of examples: Such as if they involve people far away, that they are not personal, or that they involve a majority deciding over a minority. Is it not also reasonable to believe that each added counterintuitive feature will increase the total amount of counterintuitiveness? Let us assume that, in addition to the three suggestions above, the preferences are irrational and about breaking the law. Surely it would not be strange if we put an even lower value on the preference with five counterintuitive features, than on the preference with three?

If counterintuitiveness can be aggregated in this way, it would not be strange that preferences far away in time appear to us as less important. Add to this that these preferences are often conditional in the ways described earlier, that we have less information about them than we have on present preferences, that they are more often irrational etc., and our reluctance to consider them should seem more and more plausible. This, SV could claim, is a reason to not try to surgically remove them or give them less weight with CV, but rather to keep them within a framework of sophisticated utilitarianism.

I am going to wait with discussing the prospect of CV until Section 7 and 8. First I will move

on to discussing a more specific, and much more troubling, case of diachronic preferences – preferences changing over the course of someone's life.

6. Defending the Simple View: Changing Preferences

What about the cases of changing preferences? For example, you may yesterday have preferred to eat ice cream today, but today you loathe it. Here we have two different preferences to take into consideration, in this case past-you and present-you. Should you give equal weight to both of the preferences?

There is, you might claim, something strange in splitting up our lives in this way. However, I doubt that this impression endures after a closer examination: How often have we not disagreed with what we did in the past, or solemnly promised to do differently in the future? Even if the stages of our lives are connected, we are still to some degree different from what we are in the past, or what we will be in the future.

Regardless, SV holds that the preferences belonging to a future or past stage of a person should be given equal weight to his or her present preference. What is required for a preference, according to SV, is simply a person preferring something and some preferred object that should be realised, nothing else. And since past-you and present-you occupy different temporal positions, you also give rise to different preferences. SV will therefore have to deal with examples that are intrapersonal in this way.

First of all it should be mentioned that these intrapersonal preferences are likely to, at least to some degree, confound our intuitions. This is expected, since we are seldom used to think of ourselves as having moral obligations to future or past versions of ourselves, regardless of whether we are preference or hedonist utilitarians, or not utilitarians at all. To clear up amongst our intuitions I here differ between two kinds of examples: First, (a) those where we are considering conflicting preferences within the lives of another person. Second, (b) those where we consider conflicting preferences within our own lives. The strategy for SV will be to argue that we can make sense of (a) in terms of a special relationship with our past, and to show that (b) is much less counterintuitive than it appears on first sight.

Beginning with (a), let us look closer at an example suggested by Rønnow-Rasmussen:²¹

(2), The Old House: Tom is considering whether or not to sell a house that he has inherited from his parents. Having given the matter much attention he reaches the conclusion that he is actually indifferent to

²¹ The example is borrowed from Rønnow-Rasmussen (1998), section 4.

the question. However, he knows that his sister once strongly opposed the idea of selling the house of their childhood. Recently she changed her mind – she now shows some interest in selling the house. It is clear to Tom that her original preference is stronger than her recent one.

Should Tom give priority to his sister's past preference? Two additional assumptions need to be made: First, that her previous preference was not conditional on its persistence in the future situation. Let us say she would not have wanted to sell the house even if her preference changed - she might even, at the time, have feared that she would later change her mind. Secondly, we need to make sure that there are no future preferences influencing the case. Let us assume that she in the future would be indifferent to the outcome of the affair.

Even with these qualifications, however, it does seem to me that we should give no weight to her past preference.

That the present or future preferences of others can override their past preferences in this way appears to be a problem exclusive for past preferences. We can imagine cases where we go against the synchronic wishes of a present person in order to ensure that he or she does not regret not doing a certain action in the future.²² But when the intrapersonal present or future preferences change into the past, they seem to be almost immediately discarded – not slowly losing their relevance as they move further away into the past.

Let me first note that there is nothing suspicious about our moral intuitions ignoring a certain preference, taken as a principle on the intuitive level. For example, we usually think we could ignore someone's preference to have us killed - we would not deliberate between fulfilling the preference to have us killed against fulfilling our own preference for staying alive. Similarly, a number of preferences are discarded on a daily basis in society – for example, the desire of a co-worker to be promoted on the basis of her friendship to the employer, rather than on the basis of her performance. There is a kind of special relationship between the employer and the co-worker which allows for some preferences to be ignored on the intuitive level. And, the utilitarian could claim, this is not necessarily a problem, provided that this relationship maximize preference fulfilment on the critical level.

Could SV suggest a similar principle for (a) in the case of changing preferences? One way of

²² SV would of course claim that we are biased to our present friend because of the distance in time, similarly as to when we are biased to people close to us in space. So it would not be surprising if the preference of our present-friend would take precedence over our future-friend. What is important is that the future preference of our present friend would count when we make our decision.

doing so would be to claim that we have a special relationship to our past, and that such a relationship is indirectly maximizing preference fulfilment. How could SV motivate such a principle?

First, SV could argue that it would be difficult to cooperate if we were not certain that other people would give priority to the preferences of our present-selves. And if we cannot cooperate with the people we share a temporal location with, total preference fulfilment will suffer.

Second, we tend to be better informed than our past selves as we grow up and accumulate wisdom and experience over the years. For example, just as we often consider adults to have “better” preferences than children because of the latter's inexperience and lack of knowledge, so the present persons could be said to have better preferences, on the intuitive level, than their past selves.

What could such a special relationship look like? Let me illustrate with an analogy to a non-temporal case. Imagine a group of people cooperating to help each other. Sometimes people leave the group, never to come back. These are replaced by a continuous stream of new people. The group gives priority to the current members, tries to anticipate the preferences of the future members and gives little or no weight to the preferences of those who have left. Being utilitarians, they all consider themselves maximizing preference fulfilment – we should aid people in the same groups as others because such cooperation is better on the whole. If we do not, cooperation breaks down, or is at any rate less efficient.

Similarly, we could view the past stages of other persons as having left a group cooperating through time – the past stages can no longer influence our current or future situation, and should therefore be given less or no weight on the intuitive level.

Moving on, how could SV defend against (b), considering cases where our future or past selves conflict with our own present preferences? To some extent we may recycle the explanation of (a) here, but I also think that we generally find our own past preferences more plausible to give weight. This is not unexpected, since we generally are biased in favour of ourselves, when compared to others.

Let us look closer at an example of a diachronic preference suggested by Richard Brandt to be counterintuitive:²³

(2), The Roller-Coaster Ride: Suppose your six-year-old son has decided he would like to celebrate his fiftieth birthday by taking a roller-coaster ride. Are we to take into account only the desires we think my son will have at the time his desire would be 'satisfied', here

23 The example is from Brandt (1979), 249-250. Similar examples can be found in Bykvist (1998), p 86-88.

at the age of fifty?

Imagine that we are planning our birthday at the age of fifty and are considering whether to take our past preference into account. Should we do it? Even if we assume that our preference at the age of six was not conditional on its persistence, it does seem counterintuitive to take this preference into consideration.

However, clearly not all past preferences appear to us as counterintuitive as the one for taking a roller-coaster ride. Examples of such past preferences are decisions to follow through with a certain plan, or to fulfil a certain goal. For example, imagine that you for several weeks have been secretly in love with someone. You have repeatedly promised yourself to act and ask him or her out for a date, but your synchronic preferences always let you down - fear of rejection is ever present. Surely there would be nothing strange in asking him or her out - not because present-you would like to do so, but because past-you would want you to, or because your future self would regret if you did not. These kind of situations happen to us all the time – there is, I would claim, nothing strange about giving weight to our past and future preferences in this way.

Or consider a different example: We may choose to not drink excessive amounts of beverages today and in a drunken rage insult a friend, because we know that tomorrow we would prefer that the object “I should not have insulted my friend while drunk” should have obtained in the present situation. It seems reasonable to count this preference even if we tomorrow would not actually know whether the preference was fulfilled or not - perhaps we would wake up with a perplexing case of memory loss, and our friend would never speak about what happened. That is, even if the frustration of the preference was external rather than experienced.

This does not change the fact that some preferences, such as the one for taking a roller-coaster ride at the age of fifty, may appear counterintuitive to us. But, SV could claim, this is not because it is strange to include our own past and future diachronic preferences in utilitarianism, but rather because of other features in the examples. In the case of “The Roller-Coaster Ride” there are two such features that, I suggest, if modified appropriately would change our intuitions considerably.

First, we should move the past-us closer to present-us, from childhood to just a few years back. If there is a distance to our past selves similar to the distance we have to other people through space and time, it is reasonable that moving past-us closer to present-us will decrease the level of counterintuitiveness.

Second, it is natural that preferences beneficial for society are more appreciated by our moral intuitions. This is because, SV could claim, giving these preferences priority on the intuitive level will raise overall preference fulfilment. Our example should therefore concern something more

valuable for society than “simple entertainment”.

Let us therefore consider:

(2), The Business Manager: Five years ago, at the age of forty-five, you decided to work hard for five years of your life in order to become a successful business manager. You did so to earn a certain amount of money to give to charity. Are you to take into account only the desires you have at the time your desire would be 'satisfied', here at the age of fifty, or should you also count the desires of the past five years?

When the different features of the examples are changed in this way, it appears much more reasonable to take the past preference into account.

To sum up the previous discussion: In the past two sections I have defended SV. The strategy has been to suggest that the counterintuitiveness stems from other features than the preferences being past or future. Rather, they can be explained through pointing to different principles on the intuitive level increasing or decreasing the intuitiveness of the examples – the principles evaluated on the basis of whether they lead to greater preference fulfilment.

Of the preferences considered, the changing preferences are by far the trickiest for SV to respond to. But this is expected, given our reluctance to view ourselves as split up over time. In this regard, preference utilitarianism is no worse off than any other normative theory.

I now move on to consider CV, to evaluate to what extent it provides an alternative to SV, and specifically with regard to the counterintuitive diachronic preferences. I discuss two variants of CV – the exclusion and the downgrading of preferences.

7. Criticising the Complex View: Excluding Preferences

Could CV solve the problem of diachronic preferences, and how compelling is CV compared to SV? In this and the next section I discuss two possible variants of CV: One excluding, and one downgrading the weight of, certain preferences. I treat each of these in turn, beginning with the option of excluding preferences.

There are at least two problems facing any complex view. First, there is the problem of justification. How do we justify excluding certain preferences, without such a justification being *ad hoc*? If we accept the reflective equilibrium, when do we know that a preference is counterintuitive enough to be excluded – and when can we simply accept the counterintuitive conclusions? Second,

there is the problem of arbitrariness - how do we decide on which preferences to exclude? It is possible to imagine certain examples, such as “The Future Religious Fanatics”, that could be made plausible enough, but could be done so through several different exclusions of preferences. Should we exclude irrational preferences? Or preferences that are not personal? If we cannot produce a way of deciding between these exclusions, we might end up with several equally valid theories of utilitarianism. This would be problematic if we are drawn to utilitarianism in hope of achieving some common method of decision making.

Regardless, assuming these problems can be responded to, and that we can give a properly justified non-arbitrary exclusion of certain preferences, there remains a crucial problem – would such a exclusion solve the problem of diachronic preferences? Let us look at a few different alternatives:

(a) *Excluding Diachronic Preferences*: The most obvious way of solving the problem would be to exclude all of the diachronic preferences. But there are cases of present preferences for future situations, such as a preference for the preservation of the rain forest in a future situation, which it seems plausible to give weight.

(b) *Excluding Non-present Preferences*: We might try to remedy this problem through only excluding non-present diachronic preferences. But as I have tried to show in the cases of “The Old Heirloom” and “The Beloved Son” it is possible to create examples where we would want to give weight to at least some past and future preferences.

(d) *Excluding Irrational Preferences*: Excluding irrational preferences would help us get rid of several past preferences, for example those based on certain faiths etc. But would it not be reasonable to believe that future generations will be more rational than us, for example because of higher education etc.? If such rational people replaced “The Future Fanatics” and directed our lives from the future, that would not, it seems to me, make it reasonable to give these preferences full weight.

(e) *Excluding Non-Personal Unsustained Preferences*: Perhaps we should return to the solution offered by Bykvist, excluding preferences ranging over facts about a person's life? I previously discussed this solution as a principle on the intuitive level, but Bykvist suggest that we rather modify our theory on what I here take to be the critical level. Such a modification would be part of a complex view, but would it solve the problem of diachronic preferences?

First of all, and as mentioned earlier, Bykvist excludes the preferences that are not personal – that do not range over facts about the person and her life.²⁴ This helps us get rid of “The Future Fanatics” and their rational counterparts. Second, he adds that the past preference should only count

²⁴ Bykvist (2003), p 121-124.

if the person's priorities are sustained in the present situation, or more specifically that “a person's personal preference at [t1] for something to happen at time [t2] counts iff the person's priorities at [t1] for [t2] are sustained at or throughout [t2]”.²⁵ For example, our past wish for eating an ice cream now should only count if our past priority for eating ice cream is sustained in the present situation. This helps us get rid of problematic examples where changing preferences cause a conflict between our past and future diachronic and our present synchronic preferences, such as in “The Roller-Coaster Ride” and “The Old House”. Bykvist calls the combination of these two exclusions “The Harmony View”.²⁶

However, Bykvist's model will not give the right results in the cases of “The Last Days of Earth” and “The Business Manager”. In the former case, the diachronic preferences will not count because they are not about persons, but about placing a monolith on the moon.²⁷ In the latter case, the diachronic preferences will not count if the fifty-year-old would like to spend his or her money on something else than charity – and the past preference would not just be downgraded, it would be given no weight at all.²⁸ This seems, at least to me, counterintuitive.

(f) Other Exclusions: We might consider other exclusions of preferences, for instance those where a majority tries to decide over a minority, or preferences that belong to people that we have no relationship to. However, not only would, I believe, most utilitarians balk at the prospect of such exclusions, I also think that each of them fails to exclude one or more intuitively important preferences. In the case of excluding the preferences of the many deciding over the few we would fail to account for the single intrapersonal past preference. And in the case of excluding the preferences of people that we do not know, such a move would force us to let the diachronic wishes of past friends override present synchronic preferences of people that we do not know – even if these synchronic preferences are many times more numerous.

(g) Combined Exclusions: One way of escaping the above problems would be to combine

25 Bykvist (2003), p 124.

26 Bykvist (2003), p 121-128. For a previous account of Bykvist's “Harmony View”, see Bykvist (1998), p 207-213 and Bykvist (2001), p 17-44.

27 Bykvist makes a similar comparison to preferences for the survival of the rain forest and argues that “The moral relevance for the survival of the rain forest should be determined by considering how their destruction would affect personal preferences”, Bykvist (2003), p 123.

28 Note that I previously chose to discuss “The Roller-Coaster Ride” with the six-year-old in the past and the fifty-year-old in the present. Bykvist instead considers the six-year-old in the present. This does not affect the point of neither it nor “The Business Manager” because both Bykvist and the simple view are time-neutral – they claim that there is no normative difference in the temporal position of the diachronic preference. Bykvist argues for time-neutrality in Bykvist (1998), p 67- 92 and Bykvist (2003), p 116-117.

several of the above exclusions, carefully “tailoring” our theory to avoid counterintuitiveness. However, we now run into the same problem as discussed earlier – it is both probable and possible that there are several combinations of exclusions that could lower the counterintuitiveness to the desired degree. And how do we choose between these combinations?

(h) Sophisticated Exclusions: Finally, we could follow SV and become more sophisticated – while at the same time embracing one or more exclusions. For example, we could exclude all non-present preferences, and then explain why some past and future preferences should be given weight on the intuitive level, suggesting that doing so will maximise preference fulfilment on the critical level. For example, CV could claim that respecting the wishes of the dead is beneficial for society, even if we on a critical level do not value fulfilling these preferences.

Such an approach might be workable, or at least the best alternative, for an exclusionary complex view. Especially since such a view probably will have become sophisticated to some extent already in order to avoid accusations that it would promote murder etc.

I will return to the problems of a sophisticated CV, but first I will discuss the second variant of the complex view. When defending SV it was clear that the cases discussed were compatible with a variant giving less weight to the preferences at the critical level. Could CV be defended through accepting all preferences, but downgrading the weight some preferences should be given at the critical level?

8. Criticising the Complex View: Downgrading Preferences

Downgrading the weight of certain preferences makes for a more subtle approach than simply excluding them. As such it might be more appealing to someone defending a complex view.

However, there are also problems with downgrading preferences. First of all, and similar to the exclusion of preferences, downgrading face a problem of justification: Why should some preferences be given different weight than others? The problem of arbitrariness also remains. How do we decide which preferences should be downgraded?

It is also doubtful to what extent downgrading the weight of preferences actually solves the problem of diachronic preferences being counterintuitive. I can think of at least three different ways of downgrading preferences, all of which seems to me problematic:

(a) Overriding: Certain groups of preferences could “override” others. For example, we might give weight to diachronic preferences, but only as long as they do not conflict with synchronic preferences. Such a solution, however, would not allow us to send a monolith to the moon in the case of “The Last Days of Earth”. It also suffers from the same problems as the exclusionary view,

it is too strongly formulated – we only require a single synchronic preference to overturn all the diachronic preferences in the world.

(b) *Different weight*: A less extreme solution would be to give certain preferences less weight than others, but still some weight. The obvious problem is to decide how much certain preferences should be downgraded with. Should diachronic preferences be worth half the value of synchronic ones, or less? Also, diachronic preferences close to us in time appears more valuable than those from two thousand years back. How do we account for this difference?

(c) *Variable weight*: To handle this problem, diachronic preferences could be given different weight at different times. At some times they could be given full weight, at others less weight or no weight at all. But is there not something troublesome in giving different weight to preferences depending on where we are located in time? I am worried that by accepting such a model we will lose much of what makes utilitarianism appealing - the non-discriminating way of giving all persons their equal value on the critical level. And it would be difficult to explain why we, if we consider distance in time, should not at the same time consider distance in space.

Another way of accounting for the temporal differences in (b) is, CV could suggest, becoming more sophisticated. However, I believe we should be vary of such a sophisticated complex view, both with regard to the exclusion and downgrading of preferences. Why should CV wait until until now with excluding or downgrading problematic preferences? Why not do so when first facing the anti-utilitarian examples, for example excluding or downgrading all preferences about killing other people? Perhaps there is a certain barrier where the amount of counterintuitiveness simply cannot be tolerated, even when we are explaining the examples adequately. But how do we know when we have reached this line? If it is arbitrary we face the problem of how we choose between competing theories, if it is not, the problem remains on how we find such a line – and how we agree about whether we have found it.

And – if CV is anyway going to become more sophisticated, why not become so from the start? Why not reply to the different problems in the same way, either explaining or revising our moral positions when appropriate? Why stop half-ways? There is a certain value in simplicity, a certain theoretical elegance that speaks in favour of SV here. In this regard the complex view, while not by necessity incoherent or incomprehensible, does strike me as redundant.

9. Conclusion

Two groups of diachronic preferences have been in the focus of this study. These are preferences from the past or future that disagrees with present preferences about what to do now. Giving these

future and past preferences equal weight to the present preferences seems to lead to counterintuitive results for a variant of preference utilitarianism.

This thesis discussed a dilemma between two solutions for the preference utilitarian: A simple view giving equal weight to all preferences, and a complex view excluding or downgrading the weight of some preferences.

First, I defended a simple view of utilitarianism. Such a view would solve the problem through becoming more sophisticated - by suggesting that giving some diachronic preferences less weight at the intuitive level would maximize preference fulfilment on a critical level where we give equal weight to all preferences. I tried to show how the counterintuitiveness in the diachronic cases could be explained in a similar way to synchronic cases. Such features affecting the intuitiveness of examples were if the preferences (a) concern the lives of other persons, (b) are far away in time, or (c) do not concern close friends or family. When these features were modified in the diachronic examples, they became much more plausible. I also treated the instances of preferences changing over the course of a person's life. In the case of other people's lives, I argued that much of our intuitions can be explained with us having a special relationship to the past stages of other person's lives. In the case of our own lives, I argued that the diachronic preferences are less counterintuitive than is apparent at first sight.

Second, I criticised the complex view. I first pointed to difficulties with both formulating and justifying such a view, and then considered a number of possible exclusions and different ways to downgrade the weight of preferences. None of the exclusions could solve the problem on their own, and a combination of them would risk becoming arbitrary. Downgrading the weight of preferences faced the problem of intuitiveness changing with the distance in time. Letting the weight of preferences change with their distance from us in time might be possible. However, a complex view would need to argue why it should not treat distance in space in a similar way.

An alternative solution for the complex view was to, similarly to the simple view, become more sophisticated to make the exclusions or downgrading of preferences seem more plausible. But it is, I argued, not clear why we should wait until encountering the diachronic preferences before we move from sophisticated solutions to the exclusion or downgrading of preferences. And why we should not, as the simple view, stop the anti-utilitarian "at the door".

Stopping the anti-utilitarian at the door might seem a thankless task. There is an almost limitless supply of counterintuitive examples to throw against the utilitarian position. And it is, when facing the more problematic examples such as the diachronic preferences, always tempting to defend utilitarianism through cutting of the "dead weight" of what is considered less important preferences, rather than carefully explaining them within a more sophisticated framework.

However, though such a complex view is possible to hold in the case of the diachronic preferences, it would need to become sophisticated to at least some degree in order to remain intuitively plausible. I have argued that, when becoming sophisticated, we could just as well accept a simple view from the start – treating the diachronic preferences no different than other similar allegations of utilitarianism being counterintuitive.

References

- Brandt, Richard B. (1979) *A Theory of the Good and the Right* (Oxford: Oxford University Press), ISBN: 0-19-824550-5.
- Bykvist, Krister (1998) *Changing Preferences. A Study in Preferentialism* (Uppsala: Uppsala University).
- Bykvist, Krister (2003) “The Moral Relevance of Past Preferences” in Dyke (2003) *Time and Ethics: Essays at the Intersection*, p 115-136.
- Bykvist, Krister (2001) “What is Wrong with Past Preferences” in Rabinowicz (2001) *Value and Choice, volume 2: Some Common Themes In Decision Theory and Moral Philosophy*, p 17-44.
- Dyke, Heather (2003) *Time and Ethics: Essays at the Intersection* (Dordrecht: Kluwer Academic Publishers), ISBN: 978-1-4020-1312-6.
- Egonsson, Dan (1990) *Interests, Utilitarianism and Moral Standing* (Lund: Lund University Press), ISBN: 91-7966-114-9.
- Hare, Richard M. (1981) *Moral Thinking: Its Levels, Method and Point* (Oxford: Clarendon Press), ISBN: 978-0-19-824660-2.
- Hare, Richard M. (1989) “Prudence and past preferences: Reply to Włodzimierz Rabinowicz”, *Theoria*, vol LV, Part 3, 1989, p 152-158.
- Parfit, Derek (1984) *Reasons and Persons* (Oxford: Oxford University Press), ISBN: 978-0-19-824908-5.
- Rabinowicz, Wlodek (2001) *Value and Choice, volume 2: Some Common Themes In Decision Theory and Moral Philosophy* (Lund: Department of Philosophy, Lund's University).
- Rabinowicz, Wlodek., and Österberg, Jan (1996) “Value Based on Preferences: On Two Interpretations of Preference Utilitarianism”, *Economics and Philosophy*, 12, 1996, p 1-27.
- Rabinowicz, Wlodek (1989) “Hare on Prudence”, *Theoria*, vol LV, Part 3, 1989, p 145-151.

Rawls, John (1999) *A Theory of Justice* (Oxford, Oxford University Press), ISBN: 0-19-825055.

Rønnow-Rasmussen, Toni (1998) *Preference-utilitarianism and Past Preferences* (Paper given at the Twentieth World Congress of Philosophy in Boston, 1998), direct link: <http://www.bu.edu/wcp/Papers/Valu/ValuRonn.htm>, accessed from server 2011-01-06, 09:12 am.

Shaw, William H. (2006) "Contemporary Critics of Utilitarianism: a Response" in West (2006) *The Blackwell Guide to Mill's Utilitarianism*, p 201-216.

West, Henry R. (2006) *The Blackwell Guide to Mill's Utilitarianism* (Malden, Oxford, Victoria: Blackwell Publishing Ltd), direct link: [http://uwch-4.humanities.washington.edu/Texts/Philosophy%20Guides.%20Analysis%20and%20Resources%20\(ver.2\)/Blackwell/The%20Blackwell%20Guide%20to%20Mill's%20Utilitarianism%20\(Blackwell%20Guides%20to%20Great%20Works\).pdf#page=203](http://uwch-4.humanities.washington.edu/Texts/Philosophy%20Guides.%20Analysis%20and%20Resources%20(ver.2)/Blackwell/The%20Blackwell%20Guide%20to%20Mill's%20Utilitarianism%20(Blackwell%20Guides%20to%20Great%20Works).pdf#page=203), accessed from server 2011-01-06, 09:22 am.