

Workshop on Welfare

June 22 – 23, 2022
Saarland University

Thursday, June 23

- 09:00 – 09:30 *Welcome and Introduction*
- 09:30 – 10:15 Jonas Harney (Saarland University)
[*How Narrow Person-Affecting Considerations Influence the Goodness Of Outcomes*](#)
- 10:45 – 11:30 Kacper Kowalczyk (University College London)
[*A New Argument for Fanaticism*](#)
- 12:00 – 12:45 Vuko Andrić (Bayreuth University/Institute for Futures Studies Stockholm) & Anders Herlitz (Institute for Futures Studies Stockholm)
[*Prudential Prioritarianism and Intrapersonal Trade-Offs*](#)
- 12:45 – 14:30 *Lunch break*
- 14:30 – 15:15 James Brown (University of Sheffield)
[*Additive Aggregation and Whole-Life Welfare*](#)
- 15:45 – 17:15 Nils Holtug (University of Copenhagen)
[*Prioritarianism, Risk and the Gap Between Prudence and Morality*](#)
- 19:00 *Dinner*

Friday, June 24

- 09:30 – 10:15 Joseph Van Weelden (Ahmedabad University)
[*Putting the Desire Back in Subjectivism*](#)
- 10:45 – 11:30 Shu Ishida (Osaka University)
[*Pluralist Welfare-Subjectivism and Procedural Perfectionism*](#)
- 12:00 – 12:45 Ana Gavran Miloš (University of Rijeka)
[*Mid-level theories of wellbeing: capabilitarian account*](#)
- 12:45 – 14:30 *Lunch break*
- 14:30 – 15:15 Peter Zuk (Harvard University)
[*Welfare and Other-Involving Goods*](#)
- 15:45 – 17:15 Guy Fletcher (University of Edinburgh)
[*All's Well that Ends Well?*](#)
- 19:00 *Dinner*

Jonas Harney (Saarland University)

How Narrow Person-Affecting Considerations Influence the Goodness of Outcomes

The Narrow Person-Affecting View is deeply plausible: the extent to which particular people are better or worse off than they could have been makes outcomes better or worse. It arguably underlies the moral significance of counterfactual harms and benefits, gives rise to individual claims and complaints, and motivates widely accepted Pareto principles. But the view has been accused of implying intransitive rankings of outcomes and, assuming transitivity of better-than, contradictions (Parfit 1984: 395–396; Broome 2004: ch. 10). The objection can be avoided only by rejecting the transitivity of better-than. In this talk, I explore whether the Narrow Person-Affecting View has the resources to overcome the costs of intransitivity.

I distinguish three different understandings of the Narrow Person-Affecting View: (1) narrow person-affecting considerations relative to two outcomes *A* and *B* are essentially comparative; or (2) they are features of *A* and *B* that make these outcomes better or worse even relative to a third outcome *C* in a given set of outcomes $\{A, B, C\}$; or (3) they are such features but make *A* and *B* better or worse only relative to those outcomes in which the individual exist whose welfare gain or loss gives rise to the consideration.

I argue that (1) and (3) imply cyclical betterness assessments and, assuming that one ought to bring about the better of two outcomes, lead to cyclical obligations and the threat of agents being money pumped. (3), by contrast, allows for a method to rank outcomes within fixed sets of compared outcomes that secures set-relative transitivity. This avoids the threat of cyclical assessments, cyclical obligations, and money pumps. However, (3) may be rejected, because it contradicts the spirit of the Narrow Person-Affecting Views. It implies that, given a set of outcomes, one outcome can be better or worse than another even though it is better or worse for no one. I present a case inspired by Jacob Ross (2015) that suggests that (3) has counterintuitive implications. I question the intuitions behind such an objection and argue that it relies on an unwarranted distinction between the moral relevance of narrow person-affecting considerations related to individuals that exist, on the one hand, independently and, on the other hand, dependently on the set of compared outcomes. Finally, I offer three explanations for why we should accept the implications of (3).

June 23, 10:45 – 11:30

Kacper Kowalczyk (University College London)

A New Argument for Fanaticism

It can feel uneasy to sacrifice a sure good for the sake of a very small probability of a much larger good. It can even feel *fanatical*. Indeed, *fanaticism* has become a label for the view that, for any nonzero probability p and any finite good x , there is a much larger finite good y such that a p probability of y is better than x . Bostrom, for example, argues that, since humanity could thrive on Earth for another billion years, it is better to reduce the risk of extinction by one millionth of one percent than to directly save one hundred million human lives. However, Bostrom's fanatical argument (and, to a lesser degree, other recent arguments for fanaticism, due to Beckstead and Thomas, and Wilkinson) rely on a number of controversial premises, such as total utilitarianism. In this paper, I describe a different – and more compelling – argument for a fanatical conclusion in the more circumscribed context of rescue cases, that is, cases which present a choice between an action that is certain to save a number of people from some harm and an action that has an arbitrarily small probability to save a potentially astronomical number of other people from the same harm. My argument appeals to opaque versions of these cases, that is, versions where every person involved is equally likely to be in either group. I show that my argument begs no questions against proponents of risk aversion, small probability discounting, and aggregation sceptics of the Taurek type, and that it is compatible with some salient instances of imprecise probabilism. I then consider whether and how my argument can be generalized to cases where the harms involved are not equal and where some of the people involved do not yet exist and might never exist. I also respond to recent objections to fanaticism, due to Beckstead and Thomas, and Russell. These objections are about cases that involve infinitely many possible people. In response, I develop a novel approach to infinite ethics which diagnoses an important ambiguity in our concept of betterness. This new approach is inspired by recent historical work on the mathematics of the infinite which diagnoses an important ambiguity between an Aristotelian and a Cantorian concept of size.

References

- Nick Bostrom (2003). 'Astronomical Waste: The Opportunity Cost of Delayed Technological Development'. *Utilitas*, vol. 15, no. 3, pp. 308-314.
- Nick Beckstead & Teruji Thomas (2021). 'A Paradox for Tiny Probabilities and Enormous Values'. GPI Working Paper No. 7-2021.
- Hayden Wilkinson (2022). 'In Defense of Fanaticism'. *Ethics*, vol. 132, no. 2, pp. 445-477.
- Jeffrey Sanford Russell (2021). 'On Two Arguments for Fanaticism'. GPI Working Paper No. 17-2021.

June 23, 12:00 – 12:45

Vuko Andrić (Bayreuth University/Institute for Futures Studies Stockholm)

& Anders Herlitz (Institute for Futures Studies Stockholm)

Prudential Prioritarianism and Intrapersonal Trade-Offs

When are trade-offs in well-being permissible? John Rawls (1971: 27) argued that there is a significant difference between inter- and intrapersonal trade-offs in well-being, and that a cardinal flaw of utilitarianism is that it cannot recognize this. Utilitarianism fails to recognize the moral difference between, on the one hand, decreasing an individual's welfare at one time to increase her welfare at another time and, on the other hand, decreasing one individual's welfare to increase another individual's welfare. Utilitarianism is insensitive to how welfare is distributed among several individuals and fails to accommodate the “separateness of persons”. But is the difference between inter- and intrapersonal trade-offs in well-being really morally significant? Derek Parfit (1984: 334–45) has suggested that the separateness of persons objection has less force if reductionism about personal identity is true. For on reductionism about personal identity, we may consider it less important that the individual whose welfare is decreased, rather than some other individual, is benefitted by the decrease.

Nils Holtug (2010) accepts a reductionist account of personal identity which he combines with both a theory of prudence and a distribution-sensitive moral theory, viz. a version of prioritarianism. According to Holtug (2010: 299), Relation M (roughly: the continuous physical realization of a psychology) provides lives with the sort of unity that is necessary for prudence, despite reductionism about personal identity. On Holtug's Prudential View, present self-interests are a function of discounted benefits, where benefits are discounted to the extent that Relation M holds between the bearer of the selfinterest and the recipient. On Holtug's moral theory, Prudential Prioritarianism, the moral value of a benefit to an individual at one time depends on both the size of the benefit and that individual's self-interest, at that time, in the other benefits that accrue to the individual at this and other times.

In this paper, we argue that Prudential Prioritarianism is overly restrictive. Consider cases in which early investments in terms of well-being will pay off much later in life. Examples include infant vaccinations, pension funds, and redeeming a mortgage for a private home. Such cases involve Pareto improvements in the form of intrapersonal tradeoffs that lead to later net benefits for the same individuals. However, since we are considering long-term investments, Holtug's Prudential View implies that in some such cases, it would be irrational for individuals to make these investments. For the later benefits need to be discounted, because Relation M

will be very weak between an individual at young age and the same individual at old age. Accordingly, Prudential Prioritarianism implausibly implies negative evaluations of the occurrence of some such seemingly rational long-term investments.

What is the upshot if our critique of Prudential Prioritarianism is correct? Our finding suggests that, if a view like Holtug's avoids some intuitively problematic interpersonal trade-offs in well-being despite accepting reductionism about personal identity, then the view is likely to have problematic implications regarding intuitively desirable intrapersonal trade-offs in well-being.

References

Holtug, Nils (2010): *Persons, Interests, and Justice*. New York: Oxford University Press.

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

Rawls, John (1971): *A Theory of Justice*. Oxford: Oxford University Press.

June 23, 14:30 – 15:15

James Brown (University of Sheffield)

Additive Aggregation and Whole-Life Welfare

Having a headache is a bad thing. But it would be worse if we both had a headache than if only one of us did. It would be even worse if all of Saarland had a headache, worse still if all of Germany had one, and so on. It seems, then, that the more of us have a headache, the worse things are. A premature death is also a bad thing. It is a much worse than having a headache. But is there any number people having headaches that would be worse than a premature death? It strikes many that no number of people having headaches would be worse than a single premature death. However, if welfare can be aggregated such that the goodness of an outcome can be determined by the amount of welfare it contains, then it seems inevitable that, eventually, if enough people have a headache, this will be worse than a premature death. This is what Temkin (2011) calls *the problem of additive aggregation*.

The problem of additive aggregation is most obviously a problem for utilitarians. But really it is everyone's problem. For even if we deny that what we ought to do is determined solely by the goodness of outcomes, and even if we deny that goodness is exhausted by aggregated welfare, the problem remains. As long as the goodness of an outcome is sometimes determined by the amount of welfare it contains, and as long as we think that greater aggregated welfare corresponds to better outcomes, then we seem committed to the conclusion that a great many people having a headache would be worse than a premature death. As in any case, one might simply bite the bullet and accept the unwanted conclusion (e.g. Broome 2004). However,

to many this will come at a high price. So it is worth examining whether we should accept the conclusion in the first place.

Dorsey (2009) and Bramble (2014) argue that the problem can be avoided by taking a lifetime perspective of welfare. Roughly, they argue that while a single headache may affect our *momentary* welfare, it does not affect our *whole-life* welfare, which is what ultimately matters. I argue, however, that neither approach works. First, Dorsey's approach fails because it fails to explain why a sufficient amount of momentary welfare cannot outweigh whole-life welfare in evaluating the overall goodness of an outcome. Second, Bramble's approach fails because it requires denying that small pains or pleasures matter *at all*, meaning that, contrary to appearances, headaches are not in fact bad for us. Nonetheless, I argue that the whole-life approach can avoid these problems when we focus on the *holistic* nature of the determination of whole-life welfare. Thus, it can be true that a headache can be a bad-making feature of a life without it being true that an otherwise identical life without that headache is worse for the person who lives it. I conclude that lifetime approaches to welfare can solve the problem of additive aggregation.

References

- Bramble, Ben (2014). Whole-Life Welfarism. *American Philosophical Quarterly* 51 (1):63-74.
- Broome, John (2004). *Weighing Lives*. Oxford University Press.
- Dorsey, Dale (2009). Headaches, Lives and Value. *Utilitas* 21 (1):36.
- Temkin, Larry S. (2012). *Rethinking the Good: Moral Ideals and the Nature of Practical Reasoning*. Oxford University Press.

June 23, 15:45–17:15

Nils Holtug (University of Copenhagen)

Prioritarianism, Risk and the Gap Between Prudence and Morality

According to a prominent objection to prioritarianism, it inappropriately implies a gap between prudence and morality, even in single-person cases. Thus, according to prioritarianism, we should sometimes sacrifice an individual's expected welfare in order to protect her from the risk of a worse outcome. I present a critical discussion of this objection. I first provide a more precise account of axiological prioritarianism and what it implies for the relation between prudence and morality. Then I provide an account of four prioritarian theories that (unlike axiological prioritarianism) have implications for risky choices, namely *ex ante* prioritarianism, *ex post* prioritarianism, pluralist prioritarianism, and factualist prioritarianism. I then present

the objection that prioritarianism implies a gap between prudence and morality in single-person cases in greater detail, which includes explaining the extent to which this objection applies to the four different versions of prioritarianism mentioned above. Finally, I defend the view that the prioritarian gap between prudence and morality is unproblematic, even in single-person cases.

June 24, 09:30 – 10:15

Joseph Van Weelden (Ahmedabad University)

Putting the Desire Back in Subjectivism

The desire-satisfaction theory of well-being has gotten a bad rap lately. While the objections are familiar, the direction from which much recent criticism has come is striking. Self-proclaimed *subjectivists* are now among this theory's chief critics, despite its pedigree as the subjective account par excellence.

Thus, in perhaps the most important statement of well-being subjectivism in decades, Dale Dorsey's 2021 book *A Theory of Prudence*, the heart of subjectivism is identified with the following principle:

“*Good-Value Link*: for any object, event, state, etc., ϕ and agent x , ϕ is good for x only if, and (at least in part) because ϕ is valued, under conditions c , by x ”.¹

Good-Value Link doesn't explicitly rule out views according to which valuing something, in the prudentially relevant sense, equates to desiring it. However, Dorsey utilizes the proposed tight connection between well-being and *valuing* to argue against a desiderative subjectivism. Other contemporary subjectivists similarly recommend divorcing *desiring* from *valuing* and redirecting philosophical attention to the latter.²

This paper offers a partial defense of the desire-satisfaction theory. I argue that the most plausible value-oriented subjectivism about well-being *just is* the most plausible version of the desire-satisfaction account. The advantages claimed for the newfangled nondesiderative subjective theories are either skin deep, or purchased at too great a cost. My argument takes the form of a dilemma for such approaches.

Value-based subjectivists might adopt a *permissive* stance according to which, even though desiring is not *always* sufficient for (prudentially relevant) valuing, some subset of a welfare-subject's desires do constitute such valuing. Since most desire-satisfaction theorists already hold that *some* desires are prudentially irrelevant, this makes the difference essentially terminological. But the chief selling point of value-focused accounts was their supposed ability

to disarm the notorious *scope problem*, a thorn in the side of desire-satisfaction theorists at least since Sidgwick. If the desire-satisfaction theory is vulnerable to the charge of over-inclusiveness, *simply changing our terminology can't solve this*, as any proposed method for picking out only those desires that 'constitute valuings' (and so are prudentially relevant) is equally available to the old school desire-satisfaction theorist.

The alternative is to adopt a *restrictive* stance which builds into the idea of valuing emotional/cognitive/motivational demands which are robust enough to exclude mere desires, no matter how *strong, stable, deeply held, and affectively laden* these may be. However, this incurs the reverse charge of *under-inclusiveness*. These nominally value-based theories now appear to preclude much of what welfare-subjects *actually care about* from counting as good for these same subjects. It strains credulity to deny that I *care a great deal* about the satisfaction of at least some of my desires. Since valuing something and caring about it are (as a matter of linguistic intuition) more or less equivalent, the upshot is that these theories are betraying exactly that tight connection between my good and what I value that they were designed to capture.

The real challenge for the modern subjectivist is not new. It is precisely the challenge desire-satisfaction theorists have been grappling with for generations, that of finding a principled and plausible way of cordoning off those pro-attitudes that are relevant to wellbeing from those that are not. Here shifting the target from desires to values does not help.

1 Dorsey 2021 p.80

2 See Raibley (2010 and 2013) and Tiberius (2018).

June 24, 10:45 – 11:30

Shu Ishida (Osaka University)

Pluralist Welfare-Subjectivism and Procedural Perfectionism

Welfare subjectivists face a dilemma. If we are to provide a subjectivist theory of well-being that applies to typical mature human beings, we must appeal to something "sophisticated" such as whole-life satisfaction or value fulfilment. However, that is too demanding for welfare agents without the requisite cognitive capacities, including newborns, those with profound cognitive impairments, or non-human animals. On the other hand, a less "sophisticated" theory of well-being, such as a simple desire-fulfilment theory, cannot account for the well-being of typical mature human beings, as most of today's philosophers admit.

This paper illustrates a subjectivist view of well-being, which belongs to the wider class of “subjective list” or “subjective–subjective hybrid” theories of well-being. I suggest the following view.

Highest-Attitude View (HA):

An object or state φ is good (resp. bad) for a welfare agent x to the extent that x has a pro-attitude (resp. con-attitude) towards φ and the attitude in question is in x 's highest kind.

There are various kinds of valenced attitude (pro- or con-attitude), such as being pleased with φ , being attached to φ , desiring φ , and judging φ as valuable (I use “valenced attitude” in its broadest sense). They differ in the requisite levels of cognitive capacity. For instance, judging something as valuable requires the capacity to judge, while being attached to something does not. I stipulate that a kind of valenced attitude, A_1 , is higher than another kind, A_2 , if and only if A_1 is cognitively more demanding than A_2 . The advantage of HA is as follows: (1) it is sensitive to the fact that different welfare agents have different levels of cognitive capacity, (2) it admits the possibility that (say) a person with a severe cognitive disability can be as happy as a typical human being, and (3) it provides a unifying, invariabilist explanation of why that is so.

Note that HA is perfectionist in a sense, in that it assesses x 's well-being in terms of x 's *highest* kind of valenced attitude. That is, x 's suboptimal kinds of valenced attitude, however strong, are irrelevant to x 's well-being. HA is *procedurally* perfectionist since it denotes the kind(s) of valenced attitude relevant to x 's well-being; however, it is not *substantially* perfectionist since it does not denote any object/state that contributes to x 's well-being.

HA might seem too high-minded in that any suboptimal valenced attitude is irrelevant to one's well-being. In response, I provide two variations of HA that admit that suboptimal kinds of valenced attitudes may be relevant, prudentially speaking. The **Highest-Attitude-Based Rank-Discounting View** holds that A_1 is prudentially more relevant than A_2 if and only if A_1 is higher than A_2 . According to the **Highest-Attitude-Based Lexical View**, x 's valenced attitude, A_i , towards φ is prudentially relevant only if x is indifferent to φ in terms of any kind of valenced attitude, A_j , such that A_j is higher than A_i . I finally reply to Eden Lin's argument against subjectivism (Lin, “Against Welfare Subjectivism”).

Ana Gavran Miloš (University of Rijeka)

Mid-level theory of well-being: capabilitarian account

Philosophers debating over well-being start from the big questions: What makes a good life for a person living it? What is in the best interest of a person? What is good for a person? In answering these questions, they typically try to figure out which thing(s) are ultimately good or bad for us, that is, what are the essential constituents of well-being. Standard philosophical debate distinguishes between three main positions: hedonism, desire-fulfilment theory and objective list theories (Parfit 1984). Although, philosophers so far developed different and to some extent more accurate classifications, common characteristic of philosophical theories of well-being is that they aim to provide a single, monistic, highly abstract and general account of well-being.

It has been argued recently that the problem with such attempts is that they are not useful for any other concerns besides strictly philosophical competition over the correctness of one of high well-being theories (Alexandrova 2017, Robeyns 2017, 2020). As Alexandrova argues, the main job of these philosophical theories is to provide an explanation and justification of the axiological nature of well-being. In doing that, they ask fundamental questions and offer substantive theories that rise above particular contexts searching for the common ground to all of them. In order to achieve the common ground, usually described in terms of necessary and sufficient condition, philosophical theories get highly abstract. The result is that such theories become practically useless: they cannot guide scientists since they are too general and it is impossible to measure highly abstract concepts. In order to be more useful for interdisciplinary research, it is argued, philosophers should give up on high-level theories and turn to contextualist understanding of well-being (Alexandrova 2017, Robeyns 2020).

In this paper I focus on one of the most promising contextualist well-being accounts developed by Robeyns within capabilitarian framework. Robeyns argues that the advantages of the capabilitarian account of well-being are that it is sensitive to different contexts, allows for pluralism and multidimensionality, but also provides useful guidance for social scientists and policy makers. In spite of its attractiveness, I discuss certain problems with Robeyns' account, namely that as a philosophical theory it lacks normative and descriptive adequacy. Therefore, I propose a different, more substantive capabilitarian account of well-being, based on Nussbaum's list of capabilities that keeps advantages of Robeyns' account but is also normatively and descriptively adequate. The result, in my view, is a hybrid theory of wellbeing

that consists of objective value (achievement of the central capabilities that constitute human flourishing) and subjective engagement with that value as a result of agent's deliberative desires and choices that shapes an agent's own conception of well-functioning. Therefore, I argue, such theory is both normatively and descriptively adequate. Finally, I argue that such capability account of well-being, by providing a promising framework for understanding the value element in the notion of human well-being, opens a possibility for further interdisciplinary and contextual investigation of well-being and as such is also empirically adequate.

June 24, 14:30 – 15:15

Peter Zuk (Harvard University)

Welfare and Other-Involving Goods

Other-involving goods such as romantic love, friendship, and virtue are underexplored in the contemporary philosophical literature on welfare. I develop three main claims about these goods.

The first is that each constitutively involves, as part of the good, a pro-attitude toward another's welfare.¹ While the goods mentioned involve pro-attitudes toward various features of the other, or even (in love) toward the other *in toto*, concern for the other's welfare is an essential feature of each. A relationship that included all of romantic love's other characteristic features but lacked concern for the (supposed) beloved's welfare would not be romantic love. A relationship that included friendship's other characteristic features without concern for the supposed friend's welfare would not be friendship. And a state of character that included virtue's dispositions to right action without concern for the welfare of those whose interests are thereby promoted would not be virtue.

My second main claim is that this has an important, underappreciated implication for subjective theories of welfare. Subjective theories say that welfare consists in the satisfaction of pro-attitudes such as desire, valuing, or (my preferred candidate) affective attraction. In conjunction with my first main claim, these theories imply that other-involving goods blend the welfare of self and other. In having the attitude that constitutes part of an other-involving good, we make the welfare of the other part of our own. Because my preferred candidate mental state is affective attraction, I call this *affective entanglement*.

The fact that other-involving goods extend the scope of our welfare in this way explains their importance in our lives, setting them apart as a special class rather than merely

one sort of object about which we might simply happen to subjectively care. Subjective theories, it turns out, have a ready account of why we have strong reason to seek other-involving goods, an account that need make no appeal to objective (mind-independent) value.

My third main claim is that this is a deeper and more principled explanation of the goodness of these goods than rival objective theories of welfare provide. While rejecting the idea that our attitudes in any way *explain* the goodness of welfare goods, some contemporary objectivists (e.g., Guy Fletcher) nonetheless agree that what is good for an individual must resonate with (engage, attract, or appeal to) her. This leads Fletcher to say that every (objectively valuable) welfare good constitutively includes a (subjective) pro-attitude as part of the bearer of objective value. However, this leaves unexplained why we should accept the resonance requirement in the first place. For the Fletcherian objectivist, the resonance requirement is inexplicable. All welfare goods involve positive resonance, but this tells us nothing about the *nature* of welfare. For the subjectivist, by contrast, positive resonance is an integral component of what explains the goodness of welfare goods. Subjectivism is in this respect a more coherent (and not just merely consistent) view of welfare, and to that extent preferable.

¹ This account is indebted to Guy Fletcher's more general proposal discussed below.

June 24, 15:45 – 17:15

Guy Fletcher (University of Edinburgh)

All's Well that Ends Well?

Distinguish how well someone's life is going at a particular point -- their momentary well-being -- from how well their life went as a whole, their lifetime well-being. How are these related?

The simplest answer is that lifetime well-being is just aggregate momentary well-being. Theories that deny this are forms of holism about lifetime well-being. Recent discussions of holism, inspired by David Velleman, have focused heavily on one particular species of it, the so-called 'shape of a life' hypothesis. This is the claim that having an "uphill" distribution of momentary well-being contributes to lifetime well-being and does so over and above the instrumental effects that such a distribution might have upon momentary well-being.

In this talk I outline a set of problems for the shape of a life hypothesis. I then introduce an alternative view which avoids those problems and argue for its independent plausibility.