

# Workshop on Welfare

July 9 – 11 2021

*Time specifications in CET.*

## Friday, July 9

- 14:00            *Introduction*
- 14:45–15:30    Thorsten Helfer (Saarland University)  
*[The Vagueness of Desires](#)*
- 15:45–16:30    Catherine Robb (Tilburg University)  
*[Prudential Particularism](#)*
- 17:00–18:30    Dale Dorsey (University of Kansas)  
*[The Aesthetic Life](#)*
- 19:00            *Social Event*

## Saturday, July 10

- 12:00            *Introduction*
- 12:15–13:00    Joseph Van Weelden (Ahmedabad University)  
*[The Single Life Repugnant Conclusion](#)*
- 13:15–14:00    Stephane Lemaire (Rennes University)  
*[Why would fulfilling a pro-attitude contribute to well-being?](#)*
- 14:30–16:00    Eden Lin (Ohio State University)  
*[Pleasure, Pain, and Pluralism about Well-Being](#)*
- 16:45–17:30    Charlotte Unruh (Technical University of Munich)  
*[The problem with measuring non-comparative harm](#)*
- 17:45–18:30    Jonas Harney (Saarland University)  
*[The Pseudo Person-Affectingness of the Interpersonal Comparative View](#)*
- 19:00            *Main Social Event*

## Sunday, July 11

- 12:00            *Introduction*
- 12:15–13:00    Luca Stroppa (University of Turin / University of St. Andrew)  
***The Monstrous Conclusion: a Welfare Limit in Population Ethics?***
- 13:15–14:00    Nicholas Makins (London School of Economics)  
***Rational Choice Under Uncertainty About Population Axiology***
- 14:30–15:15    Gaia Belardinelli (University of Copenhagen)  
***Comparative and non-comparative aspects of desert-adjusted axiologies***
- 15:45–17:15    Ralf Bader (University of Fribourg)  
***Person-affecting population ethics***
- 17:45            *Social Event*

*The Vagueness of Desires*

Desire-Satisfactionism claims that all and only A's episodes of desire satisfaction are constitutive for A's well-being. It is astonishing to see that most proponents of Desire-Satisfactionism of well-being say very little about the concept of desire. One promising concept is the pleasure-based concepts of desire.

*Pleasure-based concept of desire:* A desires  $p$  iff if A was in conditions C, A would feel pleasure.

I will argue that the pleasure-based concept (and other similar accounts) faces a problem because of the underspecified conditions C, and I will present a solution for this problem. For good reasons, these conditions C often include A representing  $p$  to herself. Some philosophers have proposed to include some strong idealisations for the representation of  $p$ . If we use such idealisations, we will end up with an alienating concept of well-being. It might turn out that my idealised self would find some represented states of affairs pleasurable that I would find utterly repugnant if actually realised.

On the other hand, if there is no idealisation, there still has to be some condition C under which A has to represent  $p$  to herself. Now, it seems that A could represent  $p$  to herself more or less detailed or vivid and these differences in representation might make a difference in whether she feels pleasure representing  $p$  to herself. To go to an extreme, I do not see a reason why it should not be the case that for all conditions C it is true that the representation of  $p$  is pleasurable or unpleasurable for A under conditions C but just the opposite, so unpleasurable or pleasurable, respectively, under the slightly more or less idealised conditions C'. In other words, it seems unacceptably ad hoc to settle for one specific condition or a specific range of conditions. So, the concept of desire is either unacceptably ad hoc or leads to an alienating concept of well-being.

I claim that a relatively unidealised concept of desire that explicitly allows vagueness can solve this problem:

*Vague pleasure-based concept of desire:* A desires  $p$  with strength  $s$  iff for all worlds of a vague range of worlds R in which A represents  $p$  to herself, A feels an average amount of pleasure  $s$ .

I will argue that the vague pleasure-based concept of desire is neither unacceptably ad hoc nor leads to an alienating concept of well-being.

Catherine Robb (Tilburg University)

July 9, 15:45–16:30

***Prudential Particularism***

There are three main types of theory that make claims about the nature of prudential value. Subjectivists claim that what is prudentially valuable for a person must fit with their own pro-attitudes, whilst objectivists claim that prudential value is determined by what has objective value, irrespective of a person's pro-attitude towards it. As a third option, 'hybrid' theories aim to incorporate elements of both subjective and objective accounts, generally claiming that something is prudentially good for a person if it has objective value *and* the person has a pro-attitude of some kind towards it.

In this paper I offer a novel response to the debate between subjectivism and objectivism, by rejecting them both, and putting forward a new 'particularist' theory about the nature of prudential value: *Prudential Particularism*. This account draws from the already established theory of 'moral particularism', which makes certain claims about the particularist nature of morality and the way in which we determine what counts as morally right. This paper argues that the moral variety of particularism can be successfully adapted to claims about the particularist nature of well-being and the way in which we determine what counts as prudentially good. Prudential Particularism holds that the reasons we have for explaining why something promotes prudential value are not captured by general principles, but are instead sensitive and vary according to the evaluative features that are particular to the specific context. What counts in favour of an action in one circumstance may count against it in another, and in another it may be altogether irrelevant. As a result, there is no such thing as a universal and unconditional fact about what counts as prudentially valuable.

I argue that both subjective and objective accounts of well-being are in fact generalist theories about prudential value. A prudential generalist claims that the reasons we have for explaining why something contributes to our well-being are captured by general principles; the reasons we give in support of a claim about well-being will always carry the same and invariable evaluative significance in every situation. For example, objective theories hold the generalist principle that what is good for a person is the fulfilment or acquisition of certain objective values, and this is *always* good for that person. Subjective theories state the generalist principle that what is good for a person is that which they have a certain (perhaps modified) pro-attitude towards it, and from this follows the prescription that if a person has that pro- attitude

towards something it will *always* have prudential value.

Prudential Particularism holds that there are no true generalist claims about what counts as prudentially valuable. As a result, both subjective and objective theories of well-being, and the hybrid accounts that originate from them, do not successfully capture what it is for something to count as prudentially valuable. In this paper I lay out and assess the prospects of Prudential Particularism, and respond to initial objections that may be raised against it.

Dale Dorsey (University of Kansas)

July 9, 17:00–18:30

***The Aesthetic Life***

Fans of welfare narrativism (the view according to which a life is improved by living in accordance with a valuable narrative) are plagued by concerns about mere aesthetic value: why should my life go better just because I live in accordance with a valuable narrative, given that valuable narratives can be full of suffering, tragedy, and so forth? While narrativists have argued that a good narrative need not be aesthetically valuable, in this paper I explore an alternative response: could the instantiation of aesthetic value in a life improve its quality? I suggest a qualified “yes”.

Joseph Van Weelden (Ahmedabad University)

July 10, 12:15–13:00

***The Single Life Repugnant Conclusion***

The contemporary debate about whether the prudential value of a life should be understood as a simple function of the prudential value of the moments within that life, or as bearing some more complicated relationship to momentary well-being (call these views atomism and holism, respectively) has largely focused on the question of whether the shape of a life matters.[1]

In this paper I want to consider a different kind of challenge to atomism, one which has been comparatively underdiscussed by theorists of prudential value. This is the problem that I will call, following Larry Temkin, the Single Life Repugnant Conclusion (SLRC). Temkin states it as follows:

For any possible life, no matter how long or how high the quality of that life

might be, there must be some much longer imaginable life whose existence, if other things are equal, would be better, even though each period of that life would be barely worth living.[2]

While the challenge posed by Parfit's repugnant conclusion has set the agenda in population ethics for decades, its intrapersonal analogue has been largely ignored.[3] However, Richard Kraut's latest work is a welcome exception to this pattern of neglect. [4] Kraut takes a version of the SLRC (due to J.M.E McTaggart), as his starting point in formulating an argument for the existence of two distinct and incommensurable orders of prudential value.

In this paper, I will begin by examining Kraut's argument. I argue that Kraut's strategy for resisting the SLRC requires him to posit certain discontinuities in value that are difficult to accept. Moreover, it leaves other versions of the SLRC untouched. It is possible to construct cases where exactly the same prudential goods are available at each time in a very long (but intuitively mediocre) and a much shorter (but intuitively wonderful) life. I present several such cases. If I am correct that, in these cases, preferring the longer life is scarcely less repugnant than in McTaggart's original thought experiment, this suggests that a quite different response is needed if we are to avoid succumbing to a version of the SLRC.

We might instead endorse something like the following principle:

A sufficiently long life that is sufficiently high in average well-being would always be better (for the person living it) than any amount of life that is sufficiently low in average well-being.[5]

Whereas Kraut relies on the claim that certain kinds of experience, when enjoyed for sufficiently long, are better than any amount of other kinds of experience, the proposed principle makes no reference to the kinds of prudential goods involved. It can therefore be generalized to any version of the SLRC. This proposal, however, is not without its problems. I close by considering what I deem to be the most damaging objection to our alternative principle. This is that it appears vulnerable to an intrapersonal version of Gustav Arrhenius's Sadistic Conclusion.[6]

[1] This terminology is adopted from Raibley 2015. Probably the most influential discussion of the 'shape of a life' issue is Velleman 2000.

[2] Temkin 2012 p.119. Although this formulation leaves it ambiguous as to whether it is the welfare

value of the life that is Temkin's subject, at other points this seems to be what he has in mind. It is in any case an explicitly prudential version of the SLRC that is my target in this paper.

[3] Indeed, Parfit's conclusion (and the devising of ever more clever strategies for resisting it) has been so central to this literature as to prompt a group of almost 30 philosophers to coauthor a piece in *Utilitas* just this year, urging population ethicists to move on to other problems.

[4] Kraut 2018

[5] Cf. Temkin's Fourth Standard View (Temkin 2012 pp.123-124)

[6] Arrhenius

Stephane Lemaire (Rennes University)

July 10, 13:15–14:00

### ***Why would fulfilling a pro-attitude contribute to well-being?***

Subjectivism about well-being claims that an event contributes to the well-being of an individual to the extent that this individual has or would have a certain type of pro-attitude towards this event under certain conditions. One question that is rarely answered fully explicitly is why the fulfillment of the type of pro-attitude that the subjectivist has picked up *contributes* to well-being. This is what I call the explanatory challenge.

One partial response to the explanatory challenge is that if someone is fully indifferent to an event, then it is hard to make sense of the idea that it nevertheless affects this person in a way that is relevant to her well-being. This argument explains why a resonance constraint has seemed to many to be a necessary condition for well-being. However, this remains a partial response. It remains to be explained why the fulfillment of the type of pro-attitude that the subjectivist has picked up is not only *required* but *contributes* to well-being.

In the present paper, I will argue that in order to meet the explanatory challenge, the subjectivist cannot merely rely on desires, or at least on a thin understanding of them. In order to overcome the explanatory challenge, subjectivism must opt for a richer understanding of desire or another richer pro-attitude. As far as I can see, there is only two paths to enrich the relevant pro-attitude: either one claims that certain pro-attitudes and maybe desires construe their objects as good or good for one, or one points to the link that desire have to other mental states such as emotions and pleasures. The claim of the paper is that the former is unable to meet the explanatory challenge while the second succeeds.

More precisely, I argue that subjectivism meets the explanatory challenge only if it appeals at least to a minimal account of the attitude of valuing. The claim that such a pro-attitude exists and that subjectivism should appeal to it have been recently made by several authors although there are disagreements about the specific nature of this valuing attitude. My aim here is not here to resolve these divergences. It is rather to show that the core features that friends of a pro-attitude of valuing all acknowledge are sufficient to meet the explanatory challenge. My aim is thus to bring a newer argument in favor of such a valuing attitude as the one that a fully satisfying subjectivist account of well-being should adopt.

Eden Lin (Ohio State University)

July 10, 14:30–16:00

***Pleasure, Pain, and Pluralism about Well-Being***

Many philosophers are inclined to accept a pluralistic theory of well-being—one on which there are either a plurality of basic goods or a plurality of basic bads. Such theories can, and often do, deem pleasure a basic good and pain a basic bad. But it might seem that they cannot accommodate just how important pleasure and pain are to well-being. Intuitively, if your life goes badly enough with respect to pleasure and pain—if, for example, you never feel any pleasure and you spend two years in unrelenting agony—then your life is not very good for you, no matter what else it contains. But if there is some basic good distinct from pleasure, as any pluralistic theory must claim, then it seems that a sufficiently large quantity of that good could more than make up for the hedonic badness of your life and give you a high lifetime welfare score. Indeed, pluralistic theories appear to imply that, by extending your life and adding more of this good to it, we could in principle make you arbitrarily well off even though your life is devoid of pleasure and full of pain. My aim in this paper is to defend pluralistic theories against this objection. After responding to the simplest version of the objection, I will present and answer a more sophisticated version of it that has recently been advanced by Theron Pummer. One lesson of this discussion will be that there are importantly different ways in which pluralistic theories of well-being can be structured. Another will be that even the standard version of hedonism may be unable to accommodate all of the ways in which pain detracts from well-being.

***The problem with measuring non-comparative harm***

Prohibitions against harming others are ubiquitous. However, the concept of ‘harm’ is notoriously difficult to define. According to the non-comparative account of harm, someone is harmed if and only if they are made badly off on some measure of welfare. The virtue of this account, or so it is thought, is that it does not require any comparisons with baselines. It can thus avoid worries that plague rival accounts.

In this paper, I argue that this apparent virtue of the non-comparative account does not withstand closer scrutiny. In order to measure harm, that is, to say just how badly off that person is, the non-comparative account needs to provide an account of what it means for a state to be intrinsically better or worse for a person according to some measure of welfare. ‘Better’ and ‘worse’ are clearly comparative notions. In order to measure harm, the non-comparative account of harm therefore needs to specify appropriate comparison baselines, contrary to what its name suggests.

I then discuss different options that are open to non-comparativists for specifying such baselines, and argue that all of them faces significant difficulties. First, defenders of the non-comparative account could define a non-comparative baseline as a state that is absolutely free from negative welfare. However, this baseline would be overly sensitive, because not all instances of negative welfare (e.g. insignificant instances of pain or desire frustrations) count as harm. Second, non-comparativists could determine baselines for harm counterfactually, thus combining the non-comparative account with a counterfactual comparative element. However, this would re-introduce problems arising in non-identity and pre-emption cases in the non-comparative account. Third, non-comparativists could argue that baselines should be determined according to some empirical or normative threshold. However, future generations lack a current level of wellbeing that could serve as an intuitively plausible empirical threshold. Moreover, if non-comparativists use a normative threshold to measure harm, they need to reject the common belief that harm is a non-moral concept.

If my argument is correct, then the question of how to measure harm presents a major challenge to defenders of the non-comparative account. What makes this challenge especially pertinent is that it arises specifically in cases involving future generations. These are the very cases that motivated the non-comparative account of harm in the first place.

***The Pseudo Person-Affectingness of the Interpersonal Comparative View***

According to the person-affecting view, the part of ethics that is concerned with the welfare of individuals should be cashed out in terms of how the individuals are affected. While the narrow version, which considers intrapersonal comparative value to be morally significant, fails to solve the Non-Identity Problem, the wide version, which focuses on absolute personal value, is subject to the Repugnant Conclusion, though. Recently, a middle view has been proposed that promises to capture our moral concerns towards future individuals but avoid their flaws. It modifies the narrow view by abstracting away from individuals' identities in order to account for *interpersonal*, not just intrapersonal, comparative value. I call this position the Interpersonal Comparative View (ICV).

In my talk, I will argue that ICV is flawed. By striking the middle, it cuts to many things along the way: it abandons intuitions that underlie the narrow view, obstructs the advantage of the wide view to account for the welfare of all future individuals, and violates its own presuppositions. I will concentrate on the third point. The value ICV considers to be morally significant is supposed to be personal (that is, value *for individuals*) and comparative (that is, the extent to which something is *better* or *worse*). I will argue that, in different people comparisons, it cannot be both. If the value is personal, it fails to be comparative. If it is comparative, it cannot be personal. I discuss several possibilities how proponents of ICV could circumvent that result all of which run into similar problems: ICV either deviates from the person-affecting view by constructing comparative value interpersonally or it fails to provide comparative value at all. Therefore, I conclude, ICV turns out to be merely *pseudo person-affecting*.

***The Monstrous Conclusion: a Welfare Limit in Population Ethics?***

Population Ethics is puzzled by the Repugnant Conclusion, an implication of some theories for population axiology according to which for any population *A*, no matter how high its welfare level, there is a better population *Z* consisting of arbitrarily many people with low welfare levels. In this paper I examine the implication of a structurally

analogous Conclusion, that I call the Monstrous Conclusion, according to which for any population, no matter how high its welfare and how large its population, there is a better population consisting of only one individual of much higher welfare.

The argument for the Monstrous Conclusion relies on two premises: first, that there is no upper limit for welfare levels; second, what I call *Welfare Dominance*: for any population  $P$ , if everyone fares much better in  $Q$  than in  $P$ ,  $Q$  has higher total welfare than  $P$  and  $Q$  is as equal as  $P$ , then  $Q$  is better than  $P$ . *Welfare Dominance* is extremely attractive, since it encompasses all criteria for which a population can be better than another given the parameters considered in Population Ethics. I show that revising *Welfare Dominance* in order to avoid the Monstrous Conclusion leads to even more disturbing implications. Since we seem forced to accept *Welfare Dominance*, the only way to avoid the Monstrous Conclusion is to accept an upper limit for welfare levels.

Who rejects the Repugnant Conclusion cannot be inclined to accept the Monstrous Conclusion, since the two Conclusions are equivalent but with terms inverted, and has thus to accept an upper limit for welfare levels. This acceptance comes at a cost. In fact, accepting the upper limit for welfare levels eases the defence of the Repugnant Conclusion: the limit implies that the individually better off population  $A$  in the Repugnant Conclusion cannot be better than a certain threshold, and who accepts the Repugnant Conclusion needs only to defend that the individually worse off people  $Z$  in the Repugnant Conclusion do not fare much worse than the individually better off people in population  $A$ . I however suggest that such defences have not been fully satisfactory so far.

The paper would not be completed without an attempt to identify the upper limit for welfare levels. I suggest that the limit may be how good of a life we can empathise with: if we can't empathise with the welfare of a life, this life is above the limit and we should not consider it while doing Population Ethics. However, this is problematic. First, I show that, if the correct wellbeing theory is non-experiential, there is no welfare level with which we can't empathise, and I express my doubts that any useful limit can be found for non-experiential wellbeing theories. Second, there seems to be no reasons to think that the limit is metaphysical, but only epistemological: since this is so, accepting the upper limit leads to hard choices concerning how to think about the metaphysically possible lives above it. Both concerns are worrying, but the alternative of accepting the Monstrous Conclusion does not seem particularly more attractive.

***Rational Choice Under Uncertainty About Population Axiology***

Many decisions we face will determine the number of people who come to exist and their welfare, from choices about having children to governmental policy on climate change. As such, any complete moral theory requires a way of assessing the value of the possible outcomes of these choices: a population axiology. Although numerous population axiologies have been proposed, all seem entail some deeply counterintuitive conclusions that many people are unwilling to accept and there is no consensus on the correct theory of population axiology. We must therefore make decisions under uncertainty about population axiology.

Some philosophers have advocated for the adoption of an analogue of expected value theory to guide decisions under such normative uncertainty, but this proposal is subject to the Problem of Intertheoretic Comparisons. That is, calculating expected value requires us to make comparisons of value across different normative theories, but we do not seem to be able to make sense of such comparisons. This paper assesses the specific application of this problem to normative uncertainty about population axiology. I aim to show that the problem is not applicable to this particular variety of normative uncertainty.

I will show there are two distinct sources of intertheoretic incomparability. The first form of incomparability, which I will call the Reference Problem, stems from the idea that some different moral theories refer to fundamentally different conceptions of moral value. The second form of incomparability, which I will call the Scale Problem, is due to the fact that moral theories may have no single, unique numerical representation and there is no clear way of representing different moral theories on the same scale. I will argue that neither of these problems applies to comparisons of value between different population axiologies, thereby saving the use of expected value theory under normative uncertainty about population axiology from its most substantial challenge.

Two main conclusions will be drawn from this result. Firstly, in the absence of any other significant problems, expected value theory may be adopted for decision making under normative uncertainty about population axiology. Secondly, there is a broader conclusion regarding the Problem of Intertheoretic Comparisons itself: it should no longer be thought of as a single problem for choice in all cases of normative uncertainty. Instead, the Reference Problem and the Scale problem should be treated

separately and assessed in their application to different types of normative uncertainty. This may lead to uncovering other applications of expected value theory that are free of the Problem of Intertheoretic Comparisons. For example, when one's credence is divided between variations on theories of the same class, such as consequentialism, deontology or virtue ethics, there may be ways around the Reference and Scale Problems. Some instances of normative uncertainty may be amenable to expected value theory, while others may not.

Gaia Belardinelli (University of Copenhagen)

July 11, 14:30–15:15

***Comparative and non-comparative aspects of desert-adjusted axiologies***

A consequentialist theory defines the rightness of an action in terms of the goodness of its consequences and, typically, one assumes that goodness resides in individual well-being alone. For instance, according to utilitarianism, an action is right if it results in the maximization of the sum-total of individual well-being. A consequentialism that understands well-being as the sole good is forced to accept intuitively implausible conclusions: it has to accept as morally right certain distributions of goods in which wicked people are benefited over virtuous ones. Some scholars have thus argued that consequentialism *as such* is to be given up, as it cannot properly incorporate considerations of justice [1].

However, Fred Feldman suggests a way to overcome such *objections from justice* within a consequentialist framework [2]. His idea is to define an action as right if it maximizes the sum total of *good*, where the good is now dependent not only on individual well-being, *but also on the fit between received and deserved well-being*. This so-called *desert-adjusted consequentialism* (DAC) is still an individualistic ethical theory, since it is defined only in terms of individual features: individual welfare, individual desert and the fit between these two values. Importantly, it does not involve information about comparisons between different individuals and their individual features. Thomas Hurka has then questioned whether a DAC can account for comparative considerations as well [3]. Since "distributive justice is essentially holistic", he submits that disregarding such comparisons would amount to allow for injustice, as for example situations in which equally deserving agents are treated unequally. On whether an individualistically defined DAC can fully account for

comparative perspectives, Hurka expressed severe doubts.

In our work, we focus on comparative and non-comparative aspects of desert-adjusted axiologies and aim to address Hurka's doubts. We back up our analysis with a formal framework and follow the literature in considering the desert-adjusted axiology as based on an individualistic intrinsic value function, i.e. a function that only depends on individual welfare and desert, and the fit between the values [4]. In order to account also for comparative aspects of justice, we identify a set of axioms that captures what behaviour this function should have. We conclude that it is possible for an individualistically defined axiology to be comparatively correct as well.

#### **References**

- [1] W. D. Ross. *The Right and the Good. Some Problems in Ethics*. Clarendon Press, 1930.
- [2] Fred Feldman. Adjusting utility for justice: A consequentialist reply to the objection from justice. *Philosophy and Phenomenological Research*, 55(3):567–585, 1995.
- [3] Thomas Hurka. Desert: Individualistic and holistic. In Serena Olsaretti, editor, *Desert and Justice*, pages 45–45. Oxford University Press, 2003.
- [4] Gustaf Arrhenius. Meritarian axiologies and distributive justice. 05 2012.
- [5] Erik Carlson. Consequentialism, distribution and desert. *Utilitas*, 9(3):307, 1997.

Ralf Bader (University of Fribourg)

July 11, 15:45–17:15

#### ***Person-affecting population ethics***

This paper argues that a person-affecting approach should reject impersonal good and reduce considerations of general good to considerations of personal good. It will be argued that the reducibility of general good when combined with non-comparativism about personal good, i.e. the view that existence and non-existence are not comparable in terms of personal good, implies that populations of different sizes are non-comparable.