A Publicly Informed Objective Theory of Illness Severity

Odd Borgar Jølstad
Department of Medicine
University of Oslo
August 2023
PhD Dissertation

Series of dissertations submitted to the
Faculty of Medicine, University of Oslo

ISBN 978-82-348-0374-1

Cover: UiO.
Print production: Graphic center, University of Oslo.

Table of contents

Acknowledgments

My time as a PhD-student at the Severity and Priority Setting in Health Care project has been varied and interesting. I have had the privilege of taking part in interviews exploring public views on illness severity, but have also spent long days reading philosophy, discussing normative questions with colleagues, and exploring theoretical issues related to priority setting. I have been able to primarily work on the theoretical issues that interest me. For this I am thankful to Mathias Barra, my main supervisor and PI on the project. Mathias has also been very available and supportive as a supervisor; the time he has spent with me has surely surpassed what any student can expect from their supervisor. Without his keen eye for logical structure and precision, my work would surely be worse than it is.

Without Niklas Juth, my second supervisor, navigating the theoretical literature would have been much more difficult. Our regular supervisory sessions have been instrumental for the development of my work. The time I spent with Niklas at the Centre for Research Ethics and Bioethics at the University of Uppsala was perhaps the most productive part of my time as a PhD-student.

I am also grateful for the supervision of Carl Tollef Solberg. Having a supervisor who bridges the gap between the theoretical and the practical has been very helpful. He has provided valuable feedback on papers and ideas, and he has been a dedicated co-author. He has also been a source of excellent career advice and advice on how to navigate the world of academia.

I have been lucky to spend this time with other PhD-students. Sharing an office with Mille Sofie Stenmarck and Marius Torjussen has been a pleasure. I am particularly grateful that Mille and I have shared an office thorough my entire time as a PhD-student. Being able to share the daily ups and downs with someone involved in the same project makes all the difference. While Sindre Horn has not been based in Oslo, I feel that we have gone through this process together. I am grateful for all the time we have spent discussing our work.

During my time as a PhD-student I have been given feedback on my work from many people and presented it at various seminars and conferences. It would not be possible to list everyone who has made an impact, but I would like to mention Adam Ehlert, Erik Gustavsson, Aksel Sterri, Sadie Regmi, Peder Skjelbred, Ole Fritjof Norheim, Bjørn Hol, Torbjörn Tännsjö, Kim Rand, Rachel Baker, Bjørn Hoffman, Greg Bognar, Anand Bhopal, Ben Davies, Preben Sørheim, Lars Sandman, and Erik Kjos Fonn. My work would not have been the same without your input.

I have enjoyed working at the Health Care Research Unit at Akershus University Hospital. Working at such a diverse institution has helped situate my work in the larger sphere of health care research. I wish to thank Tone Breines Simonsen and Nina Viksløkken Ødegård for administrative support. I am very grateful to the leader of the department, Hilde Lurås, for deciding to provide the additional funding that made my employment possible. I would also like to thank Jorun Rugkåsa for her work as group leader. I am also grateful for the feedback I have received on my work during internal seminars, and to all my coworkers for providing a supportive work environment.

Finally, I wish to thank my friends and family, and in particular my parents and my brother. Your support has contributed greatly to my well-being during the past three years, as it always does.

Abstract

When allocating scarce medical resources, we strive to do the best we can with the available resources. However, while we aspire to provide as "much health" as possible, we also care about who the beneficiaries are. Severity criteria are operative in Norway, Sweden, The Netherlands, and the UK. These criteria, while differing in interesting ways, all dictate that we should be willing to accept a higher cost of helping those who are severely ill. But what does it mean to be severely ill? Severity is a thick and heterogeneous concept. It is, at least partly, a measure of health-related worse off-ness. Severity also seems to have a prescriptive aspect: we believe that we should do something when an illness is severe.

In this dissertation, I attempt to answer two questions that should be provided by a theory of illness severity: (1) what is the relevance of subjective and objective factors of well-being to illness severity, and (2) to what extent is severity an aggregative concept. I argue for a fully aggregative objective list-theory of illness severity, informed by theoretical arguments and popular views. The dissertation builds on and extends arguments from the four included papers.

In paper 1, I examine the relationship between age and illness severity. More specifically, I examine a limited aggregation approach and find the approach problematic. A limited aggregation approach to the relationship between age and severity leads to difficulties in providing a coherent and privileged priority ordering. I argue that we should reject a limited aggregation approach to the relationship between age and illness severity. In the dissertation, I extend this argument to well-being; that is, we should similarly reject a limited aggregation approach to the relationship between health-related well-being and severity. I also argue that we should reject a non-aggregative conception of severity. In sum, we should accept an unlimited aggregation approach to illness severity.

In paper 2, I argue that we should count the effects of adaptation when assessing illness severity. Adaptation reduces suffering and suffering is relevant to any plausible theory of well-being relevant to health. Still, counting the effects of adaptation may be problematic if our theory of severity is a strictly subjective one. To account for our intuitions when subjective and objective well-being diverge, I argue that we should accept an objective list-theory of illness severity with a prominent subjective feature. This theory allows us to both accept the effects of adaptation and claim that people can be worse off due to illness even when their subjective well-being is not reduced.

In paper 3, I argue that popular views should be entered into reflective equilibrium processes if, and only if, they approximate considered judgments. This argument is in line with the method of reflective equilibrium as envisioned by John Rawls and Norman Daniels. I argue that views that have survived deliberation and reflection and views on familiar topics are more likely to approximate the standard of considered judgments. Some means of eliciting popular views, such as focus groups, citizen panels, and in-depth interviews, are more likely to yield views approximating the standard than others, such as polls and experiments.

In paper 4, I argue that we can progress towards a publicly informed reflective equilibrium even when our data do not meet the standard proposed in paper 3. We can bolster popular views by linking them with theoretical proposals that echo similar underlying intuitions. I identify some popular views on features of illness severity from an empirical study and attempt to bolster them by drawing on theories of well-being. In the dissertation, I link the discussion of these popular views with the argument in paper 2. The popular views I examine are heterogeneous, suggesting that an objective list-theory of severity is necessary to avoid discarding considered judgments in the ensuing reflective equilibrium. Regarding the necessity of an objective list-conception of severity, the popular considered judgments and my theoretical arguments are in alignment.

In the dissertation, I also examine a set of popular views on the relationship between age and illness severity and interpret them as views on aggregation. Since theories of aggregation are mutually exclusive, we cannot keep all considered judgments in the ensuing equilibrium. I argue that we have good theoretical reasons, discussed in paper 1 and extended in the dissertation, to reject most of the popular views in favor of an unlimited aggregation approach to severity.

I also make some suggestions regarding how to decide which goods and bads to include in the objective theory off illness severity and how to measure and weight these goods and bads. I argue that we should be interested in relevant facts and the considered judgments of the public and propose that the best way to progress on these questions is by going through a publicly informed reflective equilibrium process.

In conclusion, I make the case for a fully aggregative objective theory of illness severity. In such a theory, some features of illness are subject-independent, while others are subject-dependent. While the theory will necessarily be complex, this is necessary if we are to account for the complexity of our judgments on illness severity.

Sammendrag

Når vi prioriterer ressurser i helsevesenet bryr vi oss både om å få så mye helse som mulig og om å hjelpe de verst stilte. I Norge, Sverige, Nederland, og Storbritannia, opereres det med alvorlighetskriterier. Disse kriteriene tilsier at vi skal være villige til å bruke mer ressurser på å hjelpe de som er alvorlig syke. Men hva vil det si å være alvorlig syk? Alvorlighet har både deskriptive og preskriptive aspekter; vi bruker alvorlighet for å beskrive hvor ille en sykdom er, men vi føler også typisk at alvorlige tilstander er noe vi må gjøre noe med.

I denne avhandlingen forsøker jeg å svare på to spørsmål om alvorlighet. Det første er hvilke faktorer som bidrar til en sykdoms alvorlighetsgrad. Det andre er hvordan man skal aggregere (legge sammen) disse faktorene for å gjøre en helhetsvurdering. Jeg argumenterer for en objektiv liste-teori om alvorlighet og ubegrenset aggregering, basert på befolkningsperspektiver og teoretiske argumenter. Kappen bygger på og videreutvikler argumenter fra de fire artiklene som er inkludert i avhandlingen.

I artikkel 1 undersøker jeg sammenhengen mellom alder og alvorlighet. Mer spesifikt undersøker jeg en begrenset aggregerings-tilnærming, og viser at den ikke er tilfredsstillende. En slik tilnærming gjør det vanskelig å fremstille en systematisk og privilegert prioriteringsordning. I kappen utvider jeg denne argumentasjonen til spørsmål om helserelatert velferd. Jeg argumenterer også for at en tilfredsstillende teori om alvorlighet må akseptere aggregering. Samlet sett bør vi akseptere ubegrenset aggregering i vår teori om alvorlighet.

I artikkel 2 argumenterer jeg for at effektene av tilpasning bør medberegnes når vi vurderer alvorlighetsgrad. Tilpasning reduserer lidelse, og lidelse er relevant på alle plausible teorier om velferd som er relevante for helse. Samtidig er det er problematisk å medberegne virkningene av tilpasning dersom vi opererer med en fullstendig subjektiv teori om velferd. For å håndtere våre intuisjoner om tilfeller der subjektiv og objektiv velferd skiller lag foreslår jeg en objektiv liste-teori om alvorlighet som inkluderer subjektive faktorer. En slik teori vil la oss medberegne virkningene av tilpasning, samtidig som vi kan gjøre rede for at sykdom kan være negativ for en person selv dersom deres subjektive velferd er upåvirket.

I artikkel 3 argumenterer jeg for at befolkningens perspektiver bør inkluderes i refleksiv likevekt-prosesser hvis, og bare hvis, de nærmer seg veloverveide synspunkter (considered judgments). Dette er i tråd med refleksiv likevekt som metode, slik den er beskrevet av John Rawls og Norman Daniels. Jeg argumenterer for at synspunkter som har overlevd deliberasjon og refleksjon, samt synspunkter om emner som er velkjente for deltakere, i større grad nærmer

seg denne standarden. Noen måter å samle inn befolkningens synspunkter på, som fokusgrupper, borgerpaneler, og dybdeintervjuer, vil i større grad samle perspektiver som nærmer seg standarden. Andre måter, som spørreundersøkelser og eksperimentelle studier, vil være mindre egnet.

I artikkel 4 argumenterer jeg for at vi kan gjøre fremskritt på veien mot en refleksiv likevekt-prosess informert av befolkningens perspektiver, selv hvis vi ikke har data som tilfredsstiller gullstandarden jeg foreslo i artikkel 3. Vi kan styrke befolkningsperspektiver ved å knytte dem til teori med liknende underliggende intuisjoner. I artikkelen identifiserer jeg perspektiver på hvilke faktorer som er relevante for alvorlighetsgrad i en empirisk studie, og forsøker å styrke dem ved å trekke på filosofiske teorier om velferd. I kappen knytter jeg denne diskusjonen til argumentasjonen i artikkel 2. Befolkningsperspektivene jeg diskuterer i artikkel 4 er svært heterogene. En objektiv liste-teori om alvorlighet vil i mindre grad enn andre teoretiske alternativer kreve at vi forkaster befolkningens synspunkter refleksiv likevekt-prosessen. En objektiv liste-teori om alvorlighet vil derfor kunne støttes både av befolkningens perspektiver og teoretiske argumenter.

I kappen undersøker jeg også synspunkter om sammenhengen mellom alder og alvorlighet fra den samme studien, og tolker dem som synspunkter på aggregering. Fordi teorier om aggregering er gjensidig utelukkende, kan ikke alle befolkningsperspektivene inkluderes i den endelige teorien om alvorlighet. Jeg argumenter for at argumentene som fremmes i artikkel 1 og videreutvikles i kappen tilsier at vi bør forkaste flere av befolkningssynspunktene og akseptere ubegrenset aggregering.

Mot slutten av kappen diskuterer jeg kort hvordan vi bør bestemme hvilke faktorer som skal være på den objektive listen og hvordan disse faktorene bør måles og veies opp mot hverandre. Jeg argumenterer for at det vi bør være interesserte i er relevante fakta om sykdom og befolkningens veloverveide synspunkter, og foreslår at disse spørsmålene bør håndteres gjennom en refleksiv likevekt-prosess informert av befolkningens perspektiver.

Samlet sett fremmer jeg en sak for en objektiv liste-teori om alvorlighet og ubegrenset aggregering. Noen av faktorene på listen bør være objektive og andre subjektive. Denne teorien vil være kompleks, men dette er nødvendig for å fange opp kompleksiteten i fenomenet alvorlighet.

List of Papers

Paper 1: Jølstad, B., & Juth, N. (2022). Age and Illness Severity: A Case of Irrelevant Utilities?. *Utilitas*, 1-16.

Paper 2: Jølstad, B. (2023). Adaptation and illness severity: the significance of suffering. *Medicine, Health Care and Philosophy*, 1-11.

Paper 3: Jølstad, B., Solberg, C. T., Juth, N. & Barra, M. (under review in *Erkenntniss*). When Should Popular Views be Included in a Reflective Equilibrium?

Paper 4: Jølstad, B., Stenmarck, M. S., & Barra, M. (submitted to *Social Science and Medicine*). Preparing Popular Views for Inclusion in a Reflective Equilibrium: A Case Study on Illness Severity

Introduction

> "The majority has might on its side—unfortunately; but right it has not. I am in the right—I and a few other scattered individuals."
>
> –Ibsen, An Enemy of the People.
>
> "[E]ven if ethicists were to agree on how autonomous vehicles should solve moral dilemmas, their work would be useless if citizens were to disagree with their solution, and thus opt out of the future that autonomous vehicles promise in lieu of the status quo. Any attempt to devise artificial intelligence ethics must be at least cognizant of public morality."
>
> (Awad et al., 2018)
>
> "I believe […] that the hedonistic conception of happiness and especially of suffering, has enormous moral significance in its own right. It is what constitutes much if not most of the moral horror in experiences such as torture, starvation, military combat, disease, humiliation, clinical depression, and psychosis."
>
> (Mayerfeld, 1996)

This dissertation brings together two strands of inquiry that I have been pursuing as a PhD student. One is the question of how to think about public perspectives on moral issues. This question became salient for me while we examined the views of the Norwegian public on the meaning of *illness severity*. While it is interesting to know what the public thinks about morally charged topics, it is unclear why or how it matters. It is something close to philosophical orthodoxy to assume that moral questions cannot be settled by majority vote, but it is also the case that arguing for a moral position involves justifying the position to others. This justification must clearly involve engaging with the considered judgments of these relevant others. To make headway on this question, I explore the option of including popular views in reflective equilibrium processes. Several authors have recently written on this option, but I believe that they do not take seriously enough that what we are interested in are considered judgments, rather than "bare" intuitions.

The second strand of inquiry centers on some theoretical questions regarding illness severity as a priority setting criterion. I have explored two issues of relevance to the question of how to conceptualize severity. First, what is the relevance of subjective and objective factors of well-being to illness severity? In one of the papers of this dissertation, I argue that both objective and subjective factors are necessary to account for our intuitions on severity. Secondly, to what extent is severity an aggregative concept? That is, how should we add up goods and bads to arrive at an assessment of severity? In paper 1, I argue that any amount of time spent in states of illness should contribute to our assessment of severity.

It is my belief that these two strands of inquiry can profitably be brought together. In the process of examining popular views in general, I have examined views specific to illness severity. This examination has led me to a set of popular views on illness severity that approximate considered judgments, on a charitable interpretation. These considered judgments merit consideration during the philosophical task of conceptualizing severity as a priority setting criterion. But importantly, the fact that a view merits consideration does not imply that the view should be accepted upon due consideration. As with other views, the popular views should be revised or rejected if opposed by weighty reasons. Paper 1 in this dissertation presents arguments that are, in my opinion, weighty enough to merit discarding a popular view. Paper 2 gives a theoretical reason for endorsing an objective conceptualization of the worse off-ness of severe illness, in line with the plurality of popular views.

Bringing these lines of research together allows for saying some important things about illness severity as a priority setting criterion. First, severity should be the measure of a form of worse off-ness that accounts for both objective and subjective factors of well-being. Secondly, severity should be an aggregative concept, where any difference in worse off-ness should contribute to illness severity. Accepting aggregation is in line with seeing illness severity as being morally relevant due to consequentialist reasons; we care about and should care about, the severity of illness because severe illness is bad for people. The strength of our reasons for aiding those with severe illness is a function of this badness.

I will also briefly discuss two issues that arise when conceptualizing severity in this manner: which factors should be on our objective list, and how to measure and assess the relative weight of the different factor. I have unfortunately not solved these issues, but I believe that I have pointed in the right direction. If we have reasons to care about how both objective and subjective factors of well-being affect those with illness, we must find ways to measure and weigh these factors to improve our systems of priority setting. The most plausible way to do this is via a publicly informed reflective equilibrium process.

The dissertation builds on and further develops arguments in line with the four papers. In paper 1 I argue that we should reject limited aggregation when it comes to the relationship between illness severity and age. In the dissertation, I generalize the argument to health-related well-being and also argue further for the conclusion that we should accept unlimited aggregation as a part of our theory of illness severity. In paper 2, I argue that we should accept an objective list theory of illness severity to account for our intuitions on adaptation. In the dissertation, I argue that this conclusion aligns with the considered judgments discussed in paper 4. In paper 3, I argue that popular views are relevant for normative theory when they approximate considered judgments, and that we should be sensitive to whether views are deliberated and about familiar topics. In the dissertation, I argue further that such a publicly informed reflective equilibrium process can allow us to make progress on the questions of which factors to include in our objective list theory of severity and how to weigh these factors.

Ethical Considerations

This dissertation has been written as part of the research project Severity and Priority Setting in Health Care (SEVPRI: RCN grant 303724), at HØKH, Akershus University Hospital. I have been enrolled in the PhD program at the Medical Faculty at the University of Oslo.

The Regional Committee for Medical and Health Research Ethics (Regional Ethics Committee South-East B) advised that SEVPRI falls outside their mandate (i.e., the aims and objectives are not regulated by the Health Research Act (Helseforskningsloven, 2008)).

The work presented in this dissertation is part of SEVPRIs WP3, which analyzes ethical and economic accounts of severity. While the work presented in this dissertation is theoretical, I have taken part in interviewing the public about their views on severity and Paper 4 in this dissertation normative analysis of the results from one of the project’s empirical studies, published in Stenmark et al. (2023).

For the empirical work I analyze in paper 4, the data protection representative at Akershus University Hospital evaluated the protocol for the data collection and advised that it could be conducted (*PVO. Nos 20_200 and 21_200*). Akershus University Hospital and the Principal Investigator (MB) were responsible for project oversight, including all aspects of ethical research conduct and data privacy. The SEVPRI team has strived to adhere to all relevant ethical guidelines and frameworks, such as the ethical guidelines for the social sciences (Forskningsetiske retningslinjer for samfunnsvitenskap, humaniora, juss og teologi) and the ethical guidelines for qualitative research in medicine (Veiledning for forskningsetisk og vitenskapelig vurdering av kvalitative forskningsprosjekt innen medisin og helsefag). The SEVPRI team has also had meetings discussing ethical concerns when necessary.

Still, the part of my work that analyzes public views (paper 4) does not build directly on the data gathered, but rather proceeds from the results published in Stenmarck et al. (2023). It is important to distinguish between the data collection of which I took part, and which was approved by the appropriate body, and the work in this dissertation which builds on publicly available data and which was exempted from ethical approval by Regional Committee for Medical and Health Research Ethics. The work in this dissertation still adheres to the ethical guidelines governing all of SEVPRIs work.

I wish to address an ethical concern that is specific to the work in this dissertation. In papers 3 and 4, I discuss a potential problem when analyzing popular views: there is a risk of misinterpreting the views and therefore misrepresenting the ethical judgments of the participants. I have attempted to remain both vigilant to this risk and open when it comes to how much interpretation is involved in this work. In addition, I believe that while a risk of misinterpreting the popular views remains, there is a greater risk involved in not taking popular views seriously. When analyzing questions of healthcare priority setting, we are handling issues that are relevant to most members of our community. We should take the views of the community seriously, and this will inevitably involve a measure of interpretation.

Background

My work builds on several theoretical perspectives, which I introduce in this section. I will discuss illness severity, the relationship between severity and distribution, theories of well-being, problems with measuring the value of health, deprivationism, aggregation, public perspectives, reflective equilibrium and considered judgments.

Illness Severity

When distributing scarce medical resources, we want to do the best with what we have. That is, we want to maximize the benefits of our healthcare expenditures. However, we also care about who the recipients of benefits are; in particular, we care about the worse off. That we have reasons for aiding the worse off is asserted by prioritarian, egalitarian, and sufficientarian theories of distributive fairness (Hirose, 2014). In several countries, a severity criterion is operative, dictating that we should accept a higher cost-effectiveness ratio when treating severe illness (Barra et al., 2020). According to such a criterion, we should be willing to pay more for a benefit for a patient with schizophrenia than for the same benefit for a patient with a mild anxiety disorder, all else equal. The recent introduction of severity as a priority setting criterion in the UK (National Institute for Health and Care Excellence, 2022) has increased the number of people living in a region where a severity criterion is operative. In addition, large-scale efforts are made to introduce priority setting systems in low- and middle-income countries, sometimes including measures of severity or other criteria emphasizing the importance of giving priority to the worse- or badly off. Getting severity right thus matters a great deal.

As of 2023, a severity criterion is operative in Norway, Sweden, The Netherlands, and the UK (Barra et al., 2020; National Institute for Health and Care Excellence, 2022). Notably, these countries have operationalized severity in different ways. In Norway, severity is operationalized as *absolute (QALY) shortfall:* the severity of an illness is a measure of how many good life years are lost to illness compared to a reference life. In the Netherlands, a *relative (QALY) shortfall* operationalization is used: severity is a measure of how many good life years are lost proportional to the patient’s expectations barring illness. In the UK. both operationalizations are in play. In Sweden a more qualitative definition of severity is used in priority setting (Barra et al., 2020).

Conceptually, severity is a thick concept: a concept with both a descriptive and a prescriptive aspect (Väyrynen, 2021). Severity is, at least partially, a measure of the *badness*[1] of a given illness. That is, an illness is more severe than another illness if it is worse for a person in question, all else equal. Severity also has a prescriptive aspect; we believe that we should do something when an illness is severe. As with other thick concepts, concepts with both descriptive and prescriptive aspects, the measure of badness and the action-guiding nature of severity are intertwined. It is due to the badness of illness that severity has an action-guiding nature. As such, it is natural to consider severity from the perspective of consequentialist ethics, particularly consequentialist distributive theory. It is a concept of a particular kind of badness, giving reasons for action, that severity will be analyzed in this dissertation. Severity might also have deontological or virtue ethical aspects, but I will not concern myself with these topics. I will restrict myself to questions of what kinds of badness are relevant for severity, and how these kinds should be aggregated.

A satisfactory theory of illness severity should fulfill several criteria. First, it should be able to account for our considered judgments on which illnesses are severe. If a theory claims that mild anxiety is a more severe illness than full blown psychosis or late-stage cancer, this should make us reconsider the theory. Secondly, a theory of severity should provide reasons for why an illness is severe or not, or for why an illness is more severe than another. We should for example be able to explain why schizophrenia is more severe than mild anxiety.[2] Third, the theory of illness severity should allow us to use severity as a priority setting criterion. This means that we should be able to operationalize severity in a way that allows us to construct a coherent priority ordering.

Severity and Distribution

Severity is a measure of a certain kind of badness that is relevant to any consequentialist moral theory. On a purely utilitarian theory, severity would simply be a shorthand for relevant badness without having any particular distributive implications (Jølstad & Juth, 2022). That is, a more severe illness would be worse, but we would have no specific moral reasons to be concerned

---

[1] I use the terms ‘bad’, ‘good’, ‘badness’, and ‘goodness’ to refer to features of something that impacts someone’s well-being or a state of affairs in a negative (bads) or positive (goods) way. Pain is for example a bad, whereas achieving one’s goal is (presumably) a good. The badness of an illness would be composed of all relevant bads.

[2] These criteria are inspired by how justification is envisioned on a reflective equilibrium approach to moral theory, which I draw heavily on in my work. A theory should not only be coherent with our considered judgments, but also provide explanations for them (Rechnitzer & Schmidt, 2022).

with this illness beyond maximizing utility. On such an account severity seems largely superfluous. It is within theories of distribution that emphasize our responsibility for the disadvantaged that severity finds its natural place. The three main contenders are sufficientarianism, egalitarianism, and prioritarianism.

On sufficientarian theories of distribution, it has moral significance whether people are above or below a certain threshold of the relevant good (Crisp, 2003; Hirose, 2014). What matters is whether people are *well enough off*. Sufficientarianism does not seem to align with how we use the term severity. A natural way of conceptualizing severity in a sufficientarian manner would be to claim that anyone below a certain threshold of health or health-related well-being are severely ill. Perhaps, for example, everyone whose quality of life is below 3 on a scale from 1 to 10 are severely ill. This would mean that people are either severely ill or not, and this is not how we descriptively use severity as a concept. We naturally talk about *A* being more severely ill than *B*, while *B* is more severely ill than *C;* severity of illness is clearly a continuum of sorts. This feature of severity is not easily accommodated within a sufficientarian theory of severity. While we may have moral reasons rooted in sufficientarian theory, these reasons are not meaningfully elaborated in terms of illness severity. We should therefore reject sufficientarianism as a relevant theory for conceptualizing severity and examine egalitarianism and prioritarianism.[3]

Egalitarian theories state that there is a particular badness associated with some people being worse off than others (Hirose, 2014). Some varieties of egalitarianism are luck egalitarianism (see for example Arneson, 1989 and Cohen, 1989), telic egalitarianism (as discussed by Parfit, 1991, and defended by Segall, 2016)), and relational egalitarianism (see for example Anderson, 1999). Accepting egalitarianism involves a willingness to trade total welfare to achieve a more equal distribution. It is thus a theory that attributes value, either intrinsically or extrinsically, to the higher-level properties of the distribution of resources.

---

[3] There is an argument to be made for the inclusion of a threshold at the upper end of the quality-of-life scale where it does no longer make sense to speak of severity, as Juth et al. consider in a forthcoming article. Nordheim et al. (Omsorgsdepartementet, 2014) have argued for lifetime QALY prioritarianism where the severity of illness is measured against an upper threshold of 80 QALYs. There will then be no severity above 80 QALYs. These proposals have much in common with the idea of a 'fair innings': that everyone is entitled to a normal span of health, as discussed by John Harris (1985) and Alan Williams (1997). I will not discuss these possibilities in the dissertation, but will say this: while it seems weird to call an illness such as a mild flu severe, it does not seem strange to say that a mild flu is more severe than a common cold. This resembles how it is weird to call me tall, even though I am taller than some of my friends.

Prioritarianism states that providing benefits to those who are badly off matter more, in an absolute sense (Hirose, 2014; Parfit, 1991). On a prioritarian theory, what matters is not whether someone is worse off than others, but how badly off they are simpliciter. Well-being, or improving well-being, has decreasing marginal moral importance. Prioritarianism does not give absolute priority to the worst off, like maximin- or leximin egalitarianism, and does not attach importance to distribution. Like on a utilitarian theory, a large number of small improvements can outweigh large improvements. As such, prioritarianism is essentially a form of utilitarianism that attributes more weight to low levels of well-being (or whatever matters). (Prioritarianism has for example been called utilitarianism with a prioritarian addendum by Torbjörn Tännsjö (2019).)

I will not spend a lot of time on the issue of whether severity is an egalitarian or prioritarian concept, because I believe that Juth et al. (forthcoming) have argued convincingly for prioritarianism as the best explanation. The essential question, as I see Juth et al.'s discussion, when deciding between an egalitarian and prioritarian basis for severity is whether severity is a relational concept. That is, does it matter for the severity of my condition whether your illness is better or worse? Would my illness become more severe if you became better of health-wise? Would no one have a severe illness if everyone were equally badly off in terms of health? I share Juth et al.'s intuition that this is clearly not the case and will therefore assume that illness severity can be understood as a non-relational property of an illness or state of illness. That is, a severe illness is one that leaves a person badly off in an absolute sense.

An upshoot of understanding severity as prioritarian in nature is that we can analyze it as a feature of illness for a person in isolation from others. We avoid having to consider distributional concerns when analyzing severity as a descriptor. Distributional considerations will then have to *follow from* the relevance of the badness of severe illness; that is, the severity of an illness is prior to our distributional concerns.

Well-Being and Health

Why should we care about the severity of illness, or about health whatsoever? We care about health because of why health matters: because illness has bad effects on the lives of individuals and for society. That is, we care about the value of health (Hausman, 2010). If illness was not bad in some sense, we would not be concerned with it. I will, in this dissertation, focus only on the value of health for individuals experiencing ill health. I will not concern myself with questions of productivity or the effects of illness on national finances (or whatever). I will also not consider the question of whether an illness can be severe due to its effect on people other than the afflicted person. This is not to say that this is an implausible suggestion. Some illnesses, such as Alzheimer's disease or severe addiction, have a considerable negative effect on kin and friends. I will stick to examining to what extent an illness makes a life worse for the person experiencing the illness. To answer this question, I turn to two sources: the philosophical literature on well-being and popular views on which bads are relevant for the severity of illness.

Theories of Well-Being

Theories of well-being are theories of "what is non-instrumentally or ultimately good *for* a person." (Crisp, 2017). Theories of well-being are highly heterogenous, but they all attempt to answer the question of what we have reason to want for ourselves, or about what makes life go well. As the discussion of different theories will make clear, there are many potential answers to this question in the philosophical literature. Before delving into these theories, I should make two things clear: how I use the term well-being, and why I focus on theories of well-being.

For the purposes of this dissertation, I use the terms 'well-being', 'theories of well-being', and similar terms in a wide sense. I for example include perfectionism as a theory of well-being, contrary to how the term is used by Thomas Hurka (1993) or Daniel Hausman (2015). Some of the features of people's life that I discuss as potentially important for well-being would by others be considered goods that are independent of well-being.

Why do I turn to theories of well-being? There are four primary reasons. First, theories of well-being are about what matters for people; about what makes life go well. As such, the theories I will discuss attempt to provide answers to what makes a life good or bad for the person living it. This makes them relevant to the analysis of severity, since I wish to analyze in what ways severe illness is *bad for a person*. Theories of well-being attempt to answer this question in a fundamental manner that should be relevant for a variety of lives and throughout life. Even if it turns out that a theory of well-being is not forthcoming or that well-being is hard or impossible to measure, we can still learn much from existing theory. Secondly, the philosophical literature on well-being is highly heterogeneous and diverse and can help us understand the diverse aspects of illness and its effects on people's lives. This heterogeneity and diversity should be helpful when elucidating popular views on severity. Third, there is the fact that theories of well-being often play two roles: they should tell a person what they should aim for in life, and they make claims as to which goods should be distributed in consequentialist distributive ethics. Severity also prima facie has the same structure: it describes something that is bad for someone, and it (presumably) provides us with reasons for action or motivates distributional concerns. Fourth, there already is an established tradition of understanding the value of health in terms of well-being (see for example Brock, 2002; Broome, 2002).

In the philosophical literature, it is standard to categorize theories of well-being into three categories: desire theories, hedonism, and objective list theories (Crisp, 2017; Parfit, 1984). I will, in addition, discuss perfectionism and the capability approach.

Desire Satisfaction Theory

Desire theories, or preference satisfaction theories, claim that what contributes to well-being is the satisfaction of desires or preferences; that is, what makes life go well for a person is getting what they want (Heathwood, 2015). If I desire to become a philosopher, becoming a philosopher will contribute to my well-being. Conversely, becoming a philosopher is not good for my well-being if I have no desire for being one. Desire theories derive their plausibility from the suggestion that for *X* to matter for *A's* well-being, *X* should matter *for A*. Desire theory is *subject dependent*: whether something is good for a person depends on the attitudes of the person (Hall & Tiberius, 2016). Desire theory, in its simplest forms, is vulnerable to a range of objections. Some objections begin by describing *defective desires*: cases where it turns out that getting what one wants does not make one happier (Heathwood, 2005). Say that *A* desperately wants to be immensely rich and spends their early life pursuing this desire. It then turns out, after *A* has achieved their objective, that money does not, in fact, make them happy but in fact makes them paranoid and miserable. This seems to be a case demonstrating that what we desire does not necessarily contribute to our well-being. Desire theorists have attempted to meet this challenge by positing various information- or idealization requirements (Heathwood, 2015; Lemaire, 2021). One could, for example, posit that what makes life go well is satisfying desires that one would have given enough information (Sidgwick, 1981), that survive cognitive psychotherapy (Brandt, 1979), or what your ideal self would want (Railton, 1986). Chris Heathwood (2005) has argued that defective desires are either not defective or otherwise compatible with an actual (non-idealized) desire theory, so the demand for an information- or idealization requirement is not universal.

Highly idealized desire satisfaction theories suffer from the problem that they are no longer robustly subject dependent. Does it really matter what I would have preferred given perfect information when this is not the state I am in? Imagine someone who prefers sweet white wine. Perhaps they would, given perfect information and knowledge of wine, prefer dry Riesling over sweet Riesling. Perhaps it would even be good for them to cultivate this preference. But still, if given a glass of Riesling now, they would be happier with the sweet one. They would enjoy it more. Idealized desire satisfaction theory – in attempting to overcome some of the inherent problems of non-idealized versions – risk losing contact with the very reason why we were tempted to accept desire satisfaction theory in the first place. We seem to be at an impasse: idealizing removes us from the actual experiences of the person whom we are trying to assess the well-being of, whereas not idealizing leaves the problem of defective desires.

A separate strand of objections stems from cases of adaptive preferences and their importance for distributive ethics. Adaptive preferences are shaped by (in the literature typically adverse) circumstances (Elster, 1982). Consider the *happy slave*: a slave that, even after procedures of information and idealization, prefers to be a slave. Perhaps they have, after due reflection, concluded that slavery is a correct part of the order of the universe (perhaps they endorse a "great chain of being" type belief where one's position in the moral order is ordained by God (as described by Dreyfus & Kelly, 2011)). If we were to distribute according to well-being it might be the case that this person is not badly off (on a subjective theory). The slave's desires are fulfilled. This seems wrong. This person lacks autonomy and presumably lives in (objectively) bad circumstances. What we want to say is that these preferences stem from their circumstances, that this person has reasoned in a way that makes them able to accept their (objectively) terrible situation. Intuitively, it may not seem right to allow these types of preferences to influence judgments of distributional fairness or our reasons for action (Sen, 2009).

Desire satisfaction theory captures a seemingly important aspect of illness: it prevents us from living life in the manner of our choosing. Having a chronic bowel disorder might, for example, prevent someone from attending desirable social events, dating, or from eating certain foods. This disruptive quality may not, due to hedonic adaption, be adequately captured by measures of hedonic well-being. You may, for example, have a high level of hedonic well-being while not being able to do the things that matter to you.

Hedonism

Hedonism claims that a person's well-being is constituted by the experience of pleasure and pain, and thus captures an important intuition: how you feel matters for your well-being (Gregory, 2015; Tännsjö, 1998). A life of pleasure is obviously preferable to a life of pain, all else equal. L. W. Sumner goes further, in stating that "…nothing can make our lives go better or worse unless it somehow affects the quality of our experience." (Sumner, 1999, p. 112). Hedonism also, in contrast to other theories of well-being, respects what James Griffin (1986) calls the experience requirement: for *X* to be good or bad for an individual, the person must be aware of *X*. Hedonism undoubtedly captures some of the most essential elements of well-being. Whether you are in pain, suffering from anxiety or depression, or leading a life filled with pleasure surely matters for your well-being. Hedonism is part of a broader category of theories of well-being we can call *mental state theories of well-being*: theories that claim that well-being is determined by mental states more broadly defined, such as happiness, emotional well-being or life satisfaction (Haybron, 2016).

The perhaps most intuitively compelling argument against hedonism (and other mental state theories) is that our experiences are, in fact, *not* all we care about. Robert Nozick (1974) makes this case against hedonism by constructing a thought experiment labeled the Experience Machine. Imagine that neuropsychologist could give you the experience of anything you would like. They could, for instance, give you the experience of writing a successful book, being an important member of your nation's football-team, or falling in love. You are given the option of entering an experience machine: a machine where you will, for some number of years, have the experiences of doing everything you would like to do and will not remember that this is indeed an experience curated by these (extremely capable) neuropsychologists. Would you enter the machine? Nozick argues that you would, and should, not do so. Presumably this has to do with the fact that we do not only want the experience of doing something, but want to do and be the things that we want. The experience machine argument is not universally accepted. (See for example Tännsjö (2007) for a critical discussion of the experience machine argument). A more down-to-earth example, which is also clearly relevant in the context of priority setting in health, comes from Daniel Hausman (2015). A person who suffers from depression and is told that a lobotomy could reduce their suffering while rendering them cognitively impaired is clearly not irrational if they refuse the treatment and instead decide to keep on living in their current state. This argument also points to hedonic quality not being all that matters for our well-being.

The controversial aspect of hedonism is thus that it claims that experience is *all there is* to well-being. A theory of well-being that disregarded the value of pleasure and pain entirely would be highly implausible. Bradford (2017) has for example argued that a lack of regard for the importance of pleasure makes perfectionism a less plausible theory of well-being.

Hedonism points us towards an important aspect of illness: illness is associated with, and sometimes even constituted by, different forms of suffering. Undoubtedly, the major badness of illnesses such as depression, anxiety, or chronic pain stems from their hedonic quality. While these illnesses also have objective consequences, such as unemployment, loneliness, and lack of achievement, these are not the cores of these conditions, so to speak. It is in no way incoherent, and it is probably common, to have a substantially reduced quality of life due to these illnesses without any of these "objective" criteria being present.

Objective List Theories

Objective list theories of well-being are lists of goods and bads that contribute to well-being. These theories have been defended by, among others, Richard Arneson (1999), Christopher Rice (2013), and arguably Timothy Scanlon (2000).[4] For a theory to qualify as an objective list theory, at least one of the goods or bads on the list must be *subject-independent* (Fletcher, 2015). That is, at least some part of what it means to have a certain level of well-being on the theory does not depend on the attitude of the person in question. Typically, objective list theories are also pluralistic: they claim that there are multiple components to well-being. This is not always the case though. Roger Crisp (2017) has, for instance, argued that hedonism can be seen as an objective list theory of well-being with only one item on the list. Objective lists are very flexible when it comes to understanding health and well-being, because more or less anything we consider important can be on the list. A point in favor of objective list theories is that they seem to respect people considered judgments on well-being to a greater extent than competing theories (Rice, 2013). Importantly, objective does not mean non-individual. Arneson (1999) has argued that it makes perfect sense to allow for the possibility that different people have different lists of goods that contribute to their well-being.

In contrast to perfectionism, objective list theories do not provide a coherent theory as to why the goods and bads are on the list. They thus do not give a theory on why the things on

---

[4] Scanlon describes a list of basic goods but remains skeptical of the possibility of constructing an adequate theory of well-being.

the list matter for well-being. This has been considered an objection to objective list-theories. Rice (2013) has argued against this objection, writing that "Well-being is not necessarily a natural kind, but is an evaluative concept that is used to evaluate lives and to identify certain goals as worth pursuing.". It is not obvious that well-being has a monistic nature in the sense that well-being *can* be reduced to one concept. Objective list theories face the problem that they are not subject-dependent: something can, on these theories, impact the well-being of a person without the person caring about the something. Another objection is precisely that it seems implausible with a theory of well-being that is not subject dependent. As Peter Railton (1986) puts it: "It would be an intolerably alienated conception of someone's good to imagine that it might fail in any such way to engage him." Objective list theories also suffer from a certain arbitrariness: it is not obvious how to argue for the inclusion of a set of items, and only that set of items, on the list (Fletcher, 2015). Rice (2013) argues that we could make progress towards a non-arbitrary list of objective goods by consulting peoples considered judgments on well-being. Additionally, there is the problem of how to weigh the items on the list against each other (Fletcher, 2015). How do we, for example, weight knowledge against pleasure? While it is not obvious whether other theories of well-being avoid this question (how do we for example weight different forms of subjective experience?), the problem seems particularly difficult for objective list theories with heterogenous factors all contributing to well-being.

Perfectionism

Perfectionism postulates that our good is related to the kinds of creatures that we are. It argues that there is an ideal, and that our good depends on how closely we approximate the ideal or develop ideal qualities (Bradford, 2015; Hurka, 1993). As Thomas Hurka [23, p3] puts it: "The good life […] develops[human] properties to [a] high degree or realizes what is central to human nature." Perfectionism is an objective theory of well-being, in the sense that it argues that some aspects of human well-being are not subject-dependent: some features of a life can be good or bad for an individual even when the person does not care about the feature.[5] On a perfectionist theory of well-being, having a reduced cognitive capacity due to an accident could for example reduce your level of well-being even if you are (subjectively) happy and content. This is in contrast with hedonism, where a person in a severely reduced cognitive state but with a high

---

[5] It is worth mentioning here that Hurka (1993) does not think that perfectionism is a theory of well-being as such, but rather a theory of what makes life go best. This clearly shows that there are large conceptual differences in how well-being is seen.

level of subjective well-being must be considered to be well off. Perfectionism points us in the direction of ideal functioning. If the ideal human life includes normal human physical, emotional, social, and cognitive abilities, any illness that impedes these functions will count towards a lower well-being, even in cases where subjective well-being is unaffected.

We have already seen that subjective theories of well-being are vulnerable to adaptive preference-objections. A perfectionist theory avoids this objection and offers a reason for distributing according to our intuitions on what matters, but suffers from the same problems as idealized preference theories and objective list theories. Why should it matter to me what ideal human functioning is if I am happy without it? Additionally, perfectionist theories clearly require an account of human nature or essence, and providing one is difficult (Kitcher, 1999).

Interestingly, some psychiatric diagnoses described in the DSM and ICD follow the pattern of perfectionist theory: personality disorders explicitly require non-normal functioning in emotional, social, or cognitive functioning. On the other hand, this does not necessarily speak to the reason for being concerned with these traits. Personality disorders involve a lot of suffering, and as such may be bad due to their hedonic properties.

Capability Approach

The capability approach, championed by Amartya Sen (2009) and Martha Nussbaum (2011), is a third theory of what matters that includes objective factors. It entails both that it matters that people have the opportunity to achieve well-being, and that rather than focus on well-being as it is understood by the previously discussed theories we should be concerned with capabilities for functioning and being (Robeyns & Byskov, 2023). The capability approach captures the fact that people's ability to convert resources into well-being differs, and that this difference is relevant for distributional fairness (Sen, 1980). In contrast to the other theories of well-being discussed in this section, the capability approach focuses on capabilities rather than realized functioning. It is thus not strictly a theory of well-being, but rather of what we can do and be. Martha Nussbaum's list of central capabilities include life, bodily health, bodily integrity, senses, imagination and thoughts, emotion, practical reason, affiliation, other species, play, and control of one's environment (Nussbaum, 2011). The version of the capability approach championed by Nussbaum (1988) has its roots in Aristotelian theory of well-being and as such has a lot in common with perfectionism

Capability theory has some interesting features that make it relevant to the question of severity. First, it emphasizes that we are interested not only in people's actualized well-being, but also in their opportunities. It seems plausible that the importance of several features of illness is that they remove opportunities for functioning (Hausman, 2015). Secondly, capability theory was explicitly created in the context of answering questions of distributional fairness. It seems intuitive that it matters whether people have the opportunity to fashion a life of their choosing. Unfortunately, capability theory suffers from the same problems as the objective theories and idealized preference theory, in the sense that it can fail to respect both subject dependence and the experience requirement. If I do not care about, for example, affiliation, it is not obvious that it matters for my well-being. Capability theory also shares the problem with objective list-theories that it is not obvious how to argue for *a specific list* of capabilities. Lastly, it has been argued that the capability approach as argued for by Nussbaum fails to show people what is called *secondary recognition respect*: it fails to respect people as authorities on their own good (Terlazzo, 2014). Showing proper secondary recognition respect seems difficult for all theories of well-being that are at least partly subject-independent.

The Subjective and the Objective

These various theories of well-being point to different aspects that we associate with life going well or not. In the context of health, they also all pick out aspects of the effects of illness on well-being. An important issue when discussing these theories of well-being and severity is to what extent the severity of an illness is subject-dependent or subject-independent, and to what extent a theory must respect the experience requirement. Another way of delineating the difference between objective theories is by defining objective theories of well-being as "views which hold that claims about what is good can be correct or incorrect and that the correctness of a claim about a person's good is determined independently of that person's volition, attitudes, and opinions" and subjective theories as those that deny this (Arneson, 1999). It is thus a feature of objective theories of well-being that a person can be *wrong* when making a judgment about their own well-being. If we are hedonists or desire satisfaction theorists, we will argue that the value of health is subject-dependent: whether an illness is severe depends on the state of mind of the person with illness, or on whether the illness prevents the satisfaction of important desires. If we endorse one of the theories with objective factors, we will argue that the value of illness is, at least partially, subject-independent: there are factors that contribute to the severity of illness that have nothing to do with a person's state of mind or preferences. How we decide on this issue is not only important for giving a satisfactory theoretical account of the descriptive aspects of severity, but also for the more practical question of how we should measure the severity of an illness.

Measuring the Value of Health

Typically, health states are valued by some metric of good life years: metrics that combine quantity and quality of life. I will describe the QALY, as it is the one used in regions where a severity-criterion is operative. I will then describe some challenges that present themselves when measuring the value of health: the question of whether to measure experienced or hypothetical health and the problem of adaptation to illness.

HRQoL estimates

The most used method of valuing health states, and the one which is used in all the countries which currently uses severity as a priority setting criterion, is the Quality-Adjusted Life Year (QALY). The QALY combines number of years with the health-related quality of life (HRQoL) experienced during those years, in the sense that duration of life comes in different qualities (Weinstein et al., 2009).

HRQoL is typically estimated by eliciting answers from the public on exercises like the time trade-off or the standard gamble. During a time trade-off exercise, participants are presented with a hypothetical health state and asked to indicate how many years of life they would give up in order to be healthy instead. If, for example, participants are indifferent between ten years in health state *X* and five years in 'perfect health', the health-related quality of life given health state x is 0.5 (relative to 'perfect health'). The standard gamble is calculated differently but is assumed to give somewhat the same estimates of the quality of different health states as the time trade-off. The (HRQoL) values of different health states are then used in health economic evaluations.

The value estimates elicited using QALY-methodology are based on descriptions of health states, often in the form of the EQ-5D. The EQ-5D, the preferred instrument in Norway, comprises a descriptive system of health states along five dimensions: mobility, self-care, usual activities, pain and discomfort, and anxiety and depression (Dolan, 1997). People suffering from various illnesses report their level on each of these dimensions, ranging from no problems to extreme problems. It is these descriptions of health states that participants express preferences over when the value of health states are estimated. Presented like this, the EQ-5D looks like an objective list of the value of health that includes both objective and subjective factors. Presumably mobility, self-care, and usual activities are something like objective (subject independent) factors, whereas pain and discomfort and anxiety and depression are subjective (subject dependent) factors.[6]

Still, given how the EQ-5D-health states are subsequently valued, it seems more plausible to argue that the QALY relates to a desire satisfaction (preferentialist) theory of well-being. People implicitly value health states subjectively by making a choice between different futures of different duration and in different health states. The QALY is thus a sort of hybrid: researchers have produced a list of relevant factors, which participants subsequently express preferences for. Individual preferences are then aggregated, providing numbers for cost-utility calculations.

Given that cost-utility analysis is based on aggregated preferences from participants, a debate has ensued as to whom should express preferences over health states. The most common method is to measure the preferences of the general population, who then express preferences for *hypothetical* health states. Alternatively, one could measure *experienced* health*:* the preferences of people who have the illness we seek to value (Versteegh & Brouwer, 2016). As will become clear in the next section, the question of whether to measure hypothetical or experienced health is far from straightforward.

Hypothetical and Experienced Health

It is central to the question of whether to measure experienced or hypothetical health that the methods differ in their valuation of health states. Notably, people typically value their experienced health states higher than others value them hypothetically (Damschroder et al., 2005). While this is the typical pattern, there are exceptions. People with depression for example report lower values than those generated by hypothetical health measurement (Pyne et al., 2009).

---

[6] This is not to say that mobility, self-care, and usual activities cannot be interpreted in any number of ways. One participant may, for example, report that he is highly mobile in his wheelchair (Michel et al., 2016). Similarly, what "usual activities" entails will differ between participants. It is thus possible to argue that these factors are subjective in the sense that reporting will vary due to the participants interpretations. By objective I mean subject-independent, in that when we query about these factors we are not interested in the state om mind of the participant.

There are several reasons for this discrepancy between hypothetical and experienced health. One reason is the focusing illusion: when asked to value a health state that one is not familiar with, we are likely to focus solely on limitations associated with the illness (Ubel et al., 2005). If one is asked to value a disability that means having to use a wheelchair it is easy to focus on limitations such as not being able to use stairs, needing help to board the bus or tram, and perhaps not being able to participate in familiar sports. What one does not think about is the fact that people with wheelchairs have fulfilling jobs, fulfilling social lives, and participate in (other kinds of) sports. Most people have a fairly restricted view of what life with an illness involves. When people value their experienced health, they are aware of what a life living with the illness is like, rather than only the salient aspects. Other reasons for the discrepancy include health state descriptions being incomplete, response shifts in the sense that a good life may look different from different vantage points, and a lack of accounting for adaptation (Ubel et al., 2005).

The discrepancy between hypothetical and experienced health matters for several reasons. One reason has to do with distributional consequences. Lower health state valuations means that the curing of an illness is more valuable. On the other hand, it also means that extending the life of someone with the same illness is less valuable than if the valuation was higher.[7] Another reason has to do with respect: if someone claims that their quality of life is good, who are we to say otherwise?

Adaptation to illness

One of the reasons for the discrepancy between hypothetical and experienced health valuations is that people who value health hypothetically do not account for *adaptation* to illness (Ubel et al., 2005). To adapt is to change in response to new circumstances (Menzel et al., 2002). Menzel et al. discuss the following forms of adaptation: cognitive denial of dysfunction, suppressed recognition of full health, skill enhancement, activity adjustment, substantive goal adjustment, altered conception of health, lowered expectations, and heightened stoicism. For the purposes of my dissertation, two points are of particular importance. First, that these forms of adaptation seem to be morally problematic to differing extents. While activity adjustment and goal

---

[7] This is assuming that our goal is to maximize health within our budget constraints. It is possible to get around this, for example by arguing that any life extension should be valued at perfect health. This presents its own range of problems.

adjustments seem laudable, lowered expectations seems reason for concern. If we are to prioritize based on adaptation-sensitive measures, we may want some way to acknowledge that not all forms of adaptation are morally relevant for distributional concerns. Secondly, people do not adapt equally to all illnesses. Chronic pain leads to long-term reductions in well-being (Kahneman & Krueger, 2006), and people with depression value their health lower than corresponding hypothetical health valuations (Pyne et al., 2009). Part of the badness associated with these illnesses seem to be precisely that people do not adapt to them.

Adaptive preferences

Adaptation as a phenomenon, and the associated normative literature, resembles the phenomenon and literature on adaptive preferences. Adaptive preferences are "…shaped … by facts of, or perceptions of, availability or possibility." (Dorsey, 2017). People have argued that adaptive preferences are problematic due to their being irrational (Eftekhari, 2021), because they undermine autonomy (Colburn, 2011; Elster, 1982), and because of the distributional consequences of accounting for them (Sen, 2009). An important distinction in the literature on adaptive preferences is between adaptive preferences as a problem for theories of well-being, and, as a problem for theories of distributional justice (Terlazzo, 2015). Jessica Begon (2020) uses the terms "well-being adaptive preferences" and "justice adaptive preferences" to disambiguate two uses of adaptive preferences. Well-being adaptive preferences follow the pattern referred to by Dorsey, whereas justice adaptive preferences are '… preferences that are a poor guide to individuals' entitlements.' (Begon, 2020). When discussing well-being, the subjective aspects that are affected by adaptation seem impossible to ignore. How can you disagree with someone when they tell you that they are faring well subjectively? Disagreeing may seem to imply a lack of respect, in the sense of failing to recognize a person as an authority on their own good (Terlazzo, 2014). When discussing justice on the other hand, adaptive preferences may seem deeply problematic. This is clearly shown by cases such as the "happy slave" example used above.

While it is possible to disambiguate the uses of 'adaptive preference' like Begon does, this is made more complicated when working from the perspective of consequentialist distribute ethics. Even if we argue that adaptation to illness is not a case of adaptive preferences, like Polly Mitchel (2018) does, I believe that we should still acknowledge that patients who adapt to their conditions intuitively still have a claim to health care resources. Since theories of well-being often play the part as a theory of the relevant currency in consequentialist distributive theory, we must often wrangle with the weighing of recognition respect and seemingly unacceptable distributive consequences. This is at the heart of the discussion of adaptation in the priority setting literature: how do we weigh our competing reasons having to do with respecting peoples experience on the one hand, and, ensuring a morally acceptable distribution of resources on the other? I will also argue that this weighing is relevant for how to conceptualize illness severity.

The Badness of Death and the Badness of Illness

How can death harm you? How do different goods and bads combine to leave you better or worse off? These are questions that are essential for making headway on a theory of severity. I will discuss the deprivation account of the badness of death, and competing theories of aggregation.

Epicurean riddles

Intuitively, one of the ways in which an illness can be severe is by shortening life; that is, by killing a person. When considering how death is bad for a person, one must confront the Epicurean argument against the badness of death: there is no time at which death can be bad for a person. Being dead (rather than alive) cannot be harmful, Epicurus asserts, because there is no one who *can be* harmed (Solberg, 2019). The standard response to the epicurean argument is that death is a comparative harm: one is deprived of the life one would otherwise have had but for dying. This is known as the *deprivation account of the badness of death*, first proposed by Thomas Nagel (1970). In this way, we can not only explain why death is a harm, but we can also measure how substantial the harm is relevant to a counterfactual. Being fatally run over by a bus at age 20 is intuitively a greater misfortune than being similarly run over at age 80, and the deprivation account claims that this is because the 20-year-old loses out on more life due to the accident.

While the deprivation account can explain our intuitive responses to these cases, it still leaves difficult questions. First off, it is not trivial to specify what the relevant counterfactual is. Should terminal illness at 20 years old be measured against a hypothetical life span of 80 years, or against some other counterfactual? Secondly, it is not clear whether the comparative harm accounts for all of death’s badness. Francis Kamm (Kamm, 1993) has for example suggested that death involves an *insult factor*: death not only deprives us of what could have been, but also takes away that which is ours, exposes our vulnerability, and can be seen as a decline. If Kamm is right, then the deprivation account does not account for all the badness associated with death.

Different answers can be proposed as to what the exact relationship is between time lost and relevant badness. Perhaps early years carry more weight? The Global Burden of Disease Study, for example, previously weighed the years in young adulthood more heavily (Norheim, 2019). Or perhaps the earliest years of life carry less weight? The time-relative interest account, proposed by Jeff McMahan (2002) proposes that the badness of death is moderated by the time-relative interest a person has in their future. A very young child, who presumably is not very connected to their future due to a lack of psychological capacities, has a weaker interest in their future than an adult. This implies that the death of a very young child can be less bad than the death of a young adult, even considering that the very young child has more years to lose.

Theories of Aggregation

Unlimited aggregation is a position on the question of what we can, or should, aggregate when assessing the goodness and badness of states of affairs. On an unlimited aggregation account, all differences in the relevant desiderata should be aggregated. Importantly, we can accept unlimited aggregation within a specific domain, for example when it comes to the relevance of life years or quality of life lost for assessing severity. On an unlimited aggregation account of the severity of illness, we would for example argue that small differences in quality of life or small differences in life years lost matter for the relative severity of different conditions. This is a position I will defend in this dissertation, by arguing that neither limited aggregation nor rejecting aggregation is plausible.

On what I will call *complex aggregation theories*, which can again be limited to specific domains, the badness of states of affairs are not straightforwardly the sum of all its component goodness and badness. David Velleman (1991) has for instance argued that to assess the goodness or badness of a life we must consider the whole of life: its structure and narration. On such an account the value of my life is not simply the aggregate of my well-being at specific times. Aristotle (2012) also seem to share a similar view, at least to the extent that one cannot decide on whether a life has been good until it is over.[8]

The relationship between the badness of states of affairs and our reasons for action may not be straightforward. Perhaps small differences in age do not matter when assessing whom to aid, so that dying at 40 is not relevantly different from dying at 41? Francis Kamm (1993) has argued that some claims amount to *irrelevant utilities*. An example where the *principle of irrelevant utilities* seems plausible is in order. Imagine that you must decide between curing patient *A* and *B*, who both stand to die an equally premature death such that you have no reason to choose one over the other. Then you are told that by helping patient *B*, you would also be able to help patient *C* who has a temporary mild headache. It now seems that the option with the most utility, the best option, is helping *B* and *C*, rather than *A*. But Kamm argues that this is wrong. We cannot allow the choice between saving *A* or *B* to be decided by something as trivial as curing the mild headache of *C*. The utility associated with helping patient *C* is *irrelevant* in the context of whether we should help patient *A* or *B*. What is important to note here, and what is at the center of my argument in paper 1, is that irrelevant utilities are dynamic, in the sense that what is an irrelevant utility in one context is not irrelevant in all others. A mild headache might for instance be irrelevant in the context of life-threatening illness, but not when compared with a cold. Similar theories, collectively referred to as theories of *limited aggregation*, have been argued for by Alex Voorhoeve (2014) and Larry Temkin (2014). As we will see later in the dissertation, trying to conceptualize severity in line with a limited aggregation view may lead to a lack of a systematic priority ordering of illness severity.

The last option is to reject intra-personal aggregation all together as relevant for severity. We could claim that there is no difference, in terms of severity, between dying at the age of 20 or dying at the age of 80. This would amount to rejecting the relevance of deprivation for the severity of illness. Erik Nord (Nord, 2005; Nord & Johansen, 2014) seems to argue for such a conceptualization of severity, while still allowing aggregation to play a part in the larger cost-effectiveness system for priority setting. I will argue that rejecting aggregation for illness

---

[8] It is also possible to say this on utilitarianism (it is possible to suffer so much during the remainder of one's life that it ends up being negative even though things have been very good up until now). What Aristotle seems to have in mind is that the value of a life is not straightforwardly the sum om its component times. On could not for example, in trying to value my life, simply add up the value of each year seen in isolation. There might be holistic considerations as well (some have for example suggested that the shape of a life is relevant, where it is typically argued that a life with an upward trend is better than one with a downward trend, all else equal).

severity means that we cannot account for many of our intuitions on what makes an illness severe.

These issues are not only of theoretical interest. The absolute shortfall operationalization of severity, for example, relies on the notion of deprivation. An illness that kills 20-year-olds is more severe than an illness that kills 80-year-olds due to the greater loss of potential life: the greater loss compared to the counterfactual illness free life. Additionally, the absolute shortfall operationalization of severity presupposes specific answers to the three questions posed above: the comparator is the statistically remaining life span for a person of the same age, all life years matter equally (given the same quality of life), and small differences matter (that is the current operationalization presupposes *unlimited aggregation*).

Public Perspectives and Reflective equilibrium

Well-being and health are areas where public perspectives are prima facie important. But in what sense? To prepare for the discussion of publicly informed reflective equilibrium processes, I will first briefly discuss the rise of experimental ethics before describing the notion of a reflective equilibrium and considered judgments.

Experimental Philosophy

The recent decades have seen a rising interest in experimental and descriptive ethics. Ethical reasoning and ethical judgments are examined using experimental, quantitative, and qualitative methods. Joshua Greene (2014) has, for example, examined intuitions on classic trolley dilemmas and used the results to argue for the superiority of utilitarian modes of reasoning. Perhaps the largest empirical study to date is the Moral Machine experiment (Awad et al., 2018), where millions of views about how to best program the behavior of self-driving cars were collected by means of a gamified experiment. Empirical ethics is also growing within medical ethics. Studies have, for example, been conducted on health inequality aversion (Robson et al., 2017), perceived deservingness of care during Covid-19 (Reeskens et al., 2021), features of patients relevant to ventilator triage (Chan et al., 2022), and the meaning of illness severity (Stenmarck et al., 2023).

These studies, and others like them, prompt questions of the relevance of public perspectives for normative theory. Awad et al. argued that their findings have implications for what policies to implement, whereas John Harris (2020) disagrees, calling the experiment the "Immoral Machine" experiment. He argues that following public preferences without regard for moral argument would be monstrous. The tension between these perspectives points to a difficulty in the relationship between public perspectives and ethical policy making. One the one hand, it seems like public perspectives matter for ethical policymaking. On the other hand, it also seems obvious that questions of ethics cannot be decided by "majority vote" (Savulescu et al., 2021). People used to think slavery justified, but that does not mean that it ever was. In an attempt at bridging this chasm between empirical and normative ethics, Savulescu, Gyngell, and Kahane (2021), Alice Baderin (2017), and Gustavsson and Lindblom (2023) have recently explored the option of including popular views in reflective equilibrium processes. The role of popular views in reflective equilibrium processes is the focus of paper three and four of my dissertation and will be explored below.

Reflective Equilibrium

How do we "do ethics"? That is, how do we decide whether an ethical position is justified? Reflective equilibrium is a notion of justification in ethics, where ethical judgments are considered justified when in a resting state where our theory is in harmony with our considered judgments (Daniels, 1979; Tersman, 2018). Ideally, theory and considered judgments should not only cohere but be mutually supportive; that is, the theory should explain our considered judgments and the considered judgements should support the more abstract and theoretical judgements (Rechnitzer & Schmidt, 2022). Reflective equilibrium is also used to describe the process of reflection leading to such a state of coherence, i.e., a method for reaching justified ethical judgements. I primarily rely on the method of reflective equilibrium as it was developed by John Rawls (1951, 1999), and Norman Daniels (1979).

Reflective equilibrium, as worked out by Rawls and Daniels, is largely a coherentist method.[9] In principle, no theory or considered judgment is considered above suspicion, revision, or rejection during the reflective process. However, some considered judgments and

---

[9] One could argue that reflective equilibrium is a form of modest foundationalism. This is because some of the justifiability of the equilibrium comes from the independent credibility of the intuitions or considered judgments that are entered into the process. My focus on considered judgments in this dissertation arguably means that the form of reflective equilibrium I propose is modestly foundational in this sense. Alternatively, one could interpret my work on considered judgments as an effort towards getting at what truly matters to people, rather than what they report in unideal circumstances. Such an interpretation would be compatible with reflective equilibrium being a completely coherentist method of justification.

theoretical commitments have a firmer standing than others. A final equilibrium would presumably not be endorsed if it accepted slavery or random acts of violence, no matter how well it accounted for our other considered judgments. But still, we may accept some counterintuitive implications in order to arrive at an equilibrium. For example, many philosophers endorse utilitarianism even though it does not a priori rule out slavery. It is enough that the theory rejects slavery in all plausible scenarios. An equilibrium on priority setting in health may have similarly counterintuitive implications in strange circumstances.

In some shape or form, reflective equilibrium is the de facto method in ethics. Considered judgments or intuitions are taken as a starting point for philosophical reflection, and often given a prima facie credibility. While this can seem like a dubious process, it is closely connected to what kind of enterprise philosophical ethics is: it attempts both to give answers to important normative questions but also to make sense of our moral sensibilities (Rawls, 1999). We cannot plausibly escape the fact that ethics as practiced by us is a fundamentally human endeavor, where we are forced to attend to what we believe to be valid reasons and arguments.

Considered Judgments

> Considered judgments are simply those rendered under conditions favorable to the exercise of the sense of justice, and therefore in circumstances where the more common excuses and explanations for making a mistake do not obtain. The person making the judgment is presumed, then, to have the ability, the opportunity, and the desire to reach a correct decision […] the relevant judgments are those given under conditions favorable for deliberation and judgment in general.
>
> (Rawls 1999, p. 42)
>
> A considered moral judgment, even in a particular case, is in many ways far more like a "theoretical" than an "observation" statement. […] We readily give reasons for the moral judgments, and our appeal to theoretical considerations to support them is not mainly concerned with the conditions under which the judgments are made.
>
> (Daniels 1979, p. 270)

It is an important feature of reflective equilibrium, as practiced by Rawls and Daniels, that the judgments we attempt to make sense of in the reflective process are considered. “Bare” intuitions, gut feelings, and the first answers that come to mind do not have a strong normative standing. So, we need some criteria for considered judgments. As Rawls points out, the standing of a judgment is connected to whether those making the judgment approach an ideal judge: are they unbiased, are they considering all the evidence, do they have the relevant abilities, and so forth? Daniels points out that a considered judgment, due to being susceptible to reasons and arguments, is somewhat theoretical in nature. That is, unlike in matters of pure taste, we readily give reasons for our moral judgments. In general, a considered judgment should be impartial and be the result of, or have survived, deliberation and reflection.

Considered judgments have a few features that are essential for the preceding argument. First, they are ethical propositions or arguments that can be more or less normatively powerful. That is, a judgment can be weighty or not, and important or not. I may, for example, have a firm judgment on the importance of property rights that is outweighed by judgments about distributional fairness or desert. Secondly, considered judgments are susceptible to arguments. It is possible to criticize someone’s judgments, and we should be open to the possibility that our judgments are less well-founded than other contrasting judgements. Third, and this is very important for the ensuing discussion, it should not matter from whom the judgment originates. Provided that a person pronouncing a judgment is able to reflect, is unbiased, and has deliberated, it should not matter whether the originator is a philosopher or a member of the general public.

Publicly Informed Reflective Equilibrium Processes

On the method of reflective equilibrium justification is built on achieving a theoretical system in a coherent, and mutually reinforcing, equilibrium with our considered moral judgments (Daniels, 1979; Rawls, 1951, 1999; Tersman, 2018). Because of the importance of considered judgments, reflective equilibrium seems like an ideal method for including popular views in normative deliberations. I will now discuss when popular views[10] approximate considered judgments before examining the issue of how to work with the popular views we already have access to.

Before addressing questions about publicly informed reflective equilibrium processes, it is necessary to show that two related issues are not what is at stake. One is the issue of our normative reasons for being sensitive to democratic legitimacy. The other is how public perspectives are already accounted for in certain parts of the priority setting system by gathering public preferences for health states.

Democratic Legitimacy

What I will argue in this dissertation is that we sometimes have reason for including popular views in reflective equilibrium processes. Importantly, these are reasons for taking popular views seriously as moral arguments or judgments when we do moral theorizing. A separate question is whether we have independent reasons for taking democratic legitimacy into account when constructing priority setting criteria. The case for utilitarianism or prioritarianism might for example be overwhelming in terms of normative reasons, while the public rejects these theories. We would then (perhaps) have moral reasons, rooted in democratic norms, for rejecting the morally optimal option. I will not discuss these issues in the dissertation but rather the issue of when we should take popular views seriously while discussing normative theory.

A related point is that severity and other thick concepts get part of their normative plausibility from common notions about it. When people hear that we are to prioritize those who are severely ill, this seems plausible to them. If we use these concepts, we get a certain “normative credibility” for our arguments. Changing the contents of these concepts can thus, if we are not very careful, amount to a sort of fraud. If the population in a country has accepted a

---

[10] I use the term *popular views* to refer to all views gathered from the public using any type of empirical research.

priority setting system based on what seems to be a naturally important normative factor, defining this factor in a way that does not resonate with how the concept is used seems dishonest. This also gives an extra argument for a procedure such as a collective reflective equilibrium. Gathering population views ensures that the arguments promoted in normative literature remain true to the concepts as they are employed in ordinary language.

Aggregation of Preferences

To some extent, the public is already involved in setting healthcare priorities through their input on the value of various health states (Weinstein et al., 2009). This is arguably a way of allowing the public to shape priority setting policy. But importantly, it is a way of accounting for people's preferences, not their moral reasons and judgments. How we distribute medical resources remains an issue that is decided by experts, policymakers, and clinicians. Given the problems with measuring people's preferences over health states, such as the problems relating to adaptation and focusing illusions, it is even questionable whether measuring the preferences of the public gives rise to moral reasons. My hope is that including considered judgments elicited from the population will allow the public to influence priority setting in a manner that clearly gives rise to moral reasons.

Paper 3: When do Popular Views Approximate Considered Judgments?

While the reflective equilibrium process should be open to considered judgments, it is striking that these judgments are almost only the judgments of philosophers and other academics. Typically, in a normative paper, the author identifies their intuitive responses to specific morally charged cases and works with these intuitive responses. It is often explicitly assumed that the readers share these intuitive responses. Importantly, these intuitive responses are seldom taken at face value. Their reliability is checked by considering other cases, important consequences of accepting the intuition are examined, and debunking arguments are entertained. Sometimes, a philosophical paper will end with a skepticism regarding the intuition rather than towards relevant theory.

It is part of the rationale for reflective equilibrium that a position is strengthened by entertaining and responding to considered judgments and arguments (Rawls, 1999). As such, it makes sense to open the reflective equilibrium process as wide as possible. The more arguments and considered judgments that have been either defeated or incorporated into the ethical position, the more credence we can have in the equilibrium. One potential source of such judgments is the increasing number of studies on popular views on moral topics. It seems to be in the spirit of reflective equilibrium to incorporate views from such studies on popular views in the reflective process. One plausible way of doing this, argued for by Savulescu, Gyngell and Kahane (2021), Baderin (2017), and Gustavsson and Lindblom (2023), is to incorporate popular views in the reflective process as considered judgments or intuitions. Savulescu, Gyngell, and Kahane argue that while popular views are not considered judgments in the proper sense, we can screen popular views to ensure that we are working with empirically reliable intuitions. Baderin argues that we should use popular intuitions in moral cases, such as trolley dilemmas, because popular intuitions are less likely to be polluted by theory than the intuitions of philosophers. In contrast, I argue that the only way of properly respecting popular views is to incorporate them into the reflective process as considered judgments proper. This requires an investigation into when we are justified in assuming that a view approximates a considered judgment. This investigation also reveals why popular views should be taken into account only when they could reasonably be seen as considered judgments proper.

Considered judgments are characterized by being impartial and by being the result of, or having survived, deliberation and reflection. As such, one feature of empirical evidence to be sensitive to when assessing consideredness is whether the method used to extract views encourages reflection and deliberation. Another is whether the topic of the investigation is one that we can expect participants to be familiar with.

Deliberation and reflection are the obvious factors when deciding whether a judgment is a considered one. This mirrors the rationale for reflective equilibrium itself: Have we considered alternative views, the ramifications of the judgment (both theoretical and practical), or whether we are biased for some reason? These components of deliberation are important due to two important considerations. First, our intuitions strike us as *prima facie* plausible, but we do not know why we see them in this way. It could easily be due to bias or because we have not considered all relevant aspects of the issue at hand. Secondly, as Daniels (1979) points out, we should be able to give reasons for our considered judgments. By reasoning about a judgment, we understand more about what it means and what taking it seriously would require. Before we have gone through the motions of deliberation and reflection, we do not know for sure whether we would continue to hold the judgment after being exposed to more information or the ramifications of the judgment.

In my paper, I also argue that familiarity can be a useful proxy for deliberation and reflection. I argue that this seems plausible because familiar topics are more likely to have been the subject of deliberation and reflection in the lives of participants. Most people, for example, have personal experience with the ramifications of illness or what it takes to live a good life, and may have discussed this with friends or family. They are much less likely to have gone through the same process regarding, for example, questions about aggregation or the basis of moral standing. When assessing the consideredness of popular views on a topic we should have a higher initial credence in the views when they are about familiar topics.

I also argue that we should not have a higher initial credence in prevalent intuitions. One could argue that prevalent intuitions have been the subject of much deliberation and reflection *by someone*. There would then be someone who had considered and accepted the intuition. But on second thought, the opposite may equally be true. An intuition may be less considered because it is widespread. It is much easier to accept a view unreflectively if it is accepted by “everyone”.

These reflections on the consideredness of popular views can guide us in the construction of, and examination of, studies of popular views on moral topics. The importance of deliberation and the unimportance of prevalence means that some types of studies are more suited to extracting considered judgments. Studies that encourage deliberation, such as group discussions, citizen panels, and in-depth interviews are suited for the purpose of examining morally relevant judgments. Quantitative studies, such as surveys and the Moral Machine Experiment, are less well suited.

When assessing to what extent popular views approximate considered judgments, we should consider to what extent the method of elicitation allows for deliberation and reflection and whether the topic is familiar to the participants. I now turn to working with less ideal data; data where the degree of deliberation and reflection is harder to assess.

Paper 4: Working with Non-Ideal Empirical Data

I have argued that for popular views to approximate considered judgments, we should look at whether the research method allows for deliberation and reflection, and at whether the topic of inquiry is one with which the participants are likely to be familiar. The ideal process would be to guide participants through the relevant arguments, counterarguments, and evidence. But, for much of the already existing data on popular views, this standard is not met. This is also a problem for the views which were gathered in the study of which my work is a part (for details about this work, see Stenmarck et al., 2023).

In paper 4, I examine a set of popular views on illness severity published in Stenmarck et al. (2023) and argue that we can make progress towards extracting considered judgments from these non-ideal data. First, we must make sure that the popular views are relevant for the normative issue in question; that is, are the views about severity in a manner that is relevant to the theoretical debate on severity as a priority setting criterion. Second, we should explore whether the popular views express views like those found in the philosophical literature. If popular and philosophical views build on similar intuitions, we may be able to bolster the popular views with arguments from the literature, to approximate considered judgments. In paper 4, I sum up what the popular accounts have to say about relevant good- and bad-making factors in the following manner:

Natural Life Span

On the "natural life span" account, severity is about (negative) deviations from the natural course of life. Death in old age is natural and thus not severe. A 'good death', including dying on your own terms, is less severe. Pain, being a natural part of life, is not necessarily severe either. Mental illness is considered severe, as is loss of dignity.

Severity is Subjective

On this account severity is defined – completely subjectively – by the person with a health problem. Most potentially objective measures of severity, such as death, ability to work, pain, or prognosis are rejected as being independently relevant for severity. What matters is the experience and preferences of the person in question.

Functioning and Quality of Life

On this account severity is about the effects of illness on an individual's quality of life. Illness that impacts your ability to work, enjoy your hobbies, or otherwise participate in society is severe. Still, while severity is about how an illness affects a person, it is not up to the person in question to decide whether the illness is severe. Severity must account for the individual features and traits of an individual but is not ultimately decided by the person's preferences. This account thus works with a more objective concept of quality of life, one with a focus on functioning in daily life for an individual.

Objective Measures and Triage

On this account, severity appears to be defined by objectively measurable factors such as age, diagnosis, urgency, and prognosis. Healthcare professionals and experts should be involved in defining the severity of different health states. The severity of an illness cannot, and should not, be decided by the afflicted person. Importantly, this does not mean that subjectively experienced states such as mental illness and pain cannot be severe but that the severity of these conditions must be objectively measurable or decidable from a third-person perspective. This account of severity is thus in opposition to the 'severity is subjective' view, where the subjective experience and preferences of the individual is the sole arbiter of severity assessments.

In paper 4 I argue that three of the popular accounts can be understood in line with established theories of well-being from the philosophical literature, and that the fourth also contains a judgment that we should take very seriously when doing normative work on priority setting.

The "natural life span" account of severity can be understood in line with a broadly perfectionist theory of well-being. The badness of illness is connected to what kinds of life we ideally lead, and what kinds of creatures we are. Severe illness prevents us from leading the kind of life we should be able to expect.

The "severity is subjective" account can be understood in line with a hedonist or preferentialist theory of well-being. Subjective experience is an important factor for the badness of illness, and so is illness that prevents us from leading the kind of life we want.

The "functioning and quality of life" account can be understood in line with an objective list theory or a capability approach.

The "objective measures and triage" account underscores the importance of objective measurability and fairness in conceptualizing severity. It also contains a skepticism towards relying on the subjective experiences and reports of individuals when assessing severity.

These are the popular judgments that I believe merits careful consideration in a publicly informed reflective equilibrium on severity of illness: "From the "natural life span" account: that severity is connected to what kinds of beings we are, or with what a human life can and should look like. The measurement of these factors does not naturally lend themselves to aggregation but allow for some differentiation between worse- and better off individuals. From the "severity is subjective" account: that the subjective aspects of illness are important for the severity of an illness. From the “functioning and quality of life” account: that while some factors are important for everyone, these factors may be instantiated in different ways for different individuals, and that the badness of illness can also come from restrictions of our opportunities and capabilities. From the “objective measures and triage” account: that a conception of severity should be objective and fair, and that there are reasons to suspect that allowing individuals to define their severity is difficult to combine with objectivity and fairness.

A complication

In focusing on the importance of the consideredness of the views entered in the reflective equilibrium process, I have positioned my arguments in line with theorists that argue that we should somehow filter intuitions, such as Rawls (1951, 1999) and Daniels (1979). Kelly & McGrath (2010) argue that without some way of vetting the intuitions that are entered into the reflective equilibrium process we risk ending up with a position that reflects faulty start points. Against this Rechnitzer & Schmidt (2022) argue that we can secure reasonable commitments without strong filtering conditions, because the process of reflective equilibrium requires more than simple coherence. Thiel & Delden (2009) similarly argue that we should rely on a good reasoning-justified strategy rather than a credible input-justified strategy in the reflective equilibrium process. While I am sympathetic to the idea that we should filter intuitions in some way, in line with how normative work usually proceeds, I may of course be wrong. In that case, I believe that the work I have done on the consideredness of popular views, and especially the work on how to strengthen popular views, can be of value. If any intuition, perhaps barring those that do not survive even a minimal filter, can be entered into a reflective equilibrium, that still does not speak for the chance that the intuitions will survive the reflective process. Entering popular views into the reflective equilibrium process in their strongest form would be essential for ending up with a reflective equilibrium that builds on popular views in any meaningful way. We should thus still, when designing studies to gather popular views, make sure that we allow for deliberation and reflection. We should also attempt to strengthen popular views. These are essential elements when constructing normative theory that takes our responsibility to the views of others seriously.

The Reflective Process: What Stays and What Goes?

Paper 4 ends by mentioning, somewhat in passing, that when we have identified which considered judgments should be included in the reflective process, we still have the task of reaching an equilibrium ahead of us. Given the coherentist characteristics of reflective equilibrium, it seems unlikely that all considered judgments can remain in the final theory. This seems like a particularly difficult challenge when working towards a publicly informed reflective equilibrium. When constructing a reflective equilibrium as a sole researcher, or perhaps together with like-minded others, we have the benefit of a limited supply of considered judgments. Plausibly, completely contradictory considered judgments will be less prevalent than when constructing a publicly informed reflective equilibrium. Given the widely differing moral outlooks in even fairly homogeneous societies such as the Norwegian one, it seems likely that a publicly informed reflective equilibrium will have to consider a wider range of considered judgments. But how do we choose what stays and what goes?

Unfortunately, it does not seem possible to give a more precise answer than that we must go through the same process as during a “regular” reflective equilibrium process. That is, we must consider the various arguments, theories, and judgments and weigh their relative merits. Importantly, given that I have argued that the prevalence of a given judgment is not relevant for its consideredness and subsequent weight during the reflective process, we should not be inclined to accept commonly held judgments in favor of less commonly held ones. The process thus requires sensitivity to reasons, an eye for coherence, and a proclivity for reasonableness. These are all traits that are nurtured by philosophical training, and in particular by experience in moral argument. The process cannot, at least for now, be an algorithmic one. What remains for us is the task of weighing opposing reasons and judging the merits and demerits of proposed equilibriums. I will approach the issue by examining two questions that must be answered in order to conceptualize severity. The first question is which goods and bads are relevant for illness severity; that is, what is the currency of severity. The second question is how to aggregate these goods and bads within a life and whether all bad are relevant in assessing severity; that is, to what extent is severity an intra-personally aggregative concept, and to what extent should we accept unlimited aggregation. I begin with the question of the currency of severity, before addressing the question of aggregation.

The Currency of Severity

I will now address the first question discussed in the introduction: what are the relevant goods and bads when assessing severity? I will attempt to make progress by drawing on popular views and theoretical arguments.

The Public Perspectives

The following table, from paper 4, summarizes the considered judgments from the popular account of severity.

Table 1: Mapping accounts of severity to moral theory

Account	**Considered judgements**
Natural life span	Severity is connected to what kinds of beings we are, or with what a human life can and should look like. Some relevant goods and bads do not naturally lend themselves to aggregation but allow for some differentiation between worse- and better off individuals.
Severity is subjective	The subjective aspects of illness are important for the severity of an illness.
Functioning and quality of life	While some factors are important for everyone, these factors may be instantiated in different ways for different individuals. The badness of illness can also come from restrictions of our opportunities and capabilities.
Objective measures and triage	A conception of severity should be objectively measurable and fair. Allowing individuals to define their severity is difficult to combine with objectivity and fairness.

Table 1. This table details views from the popular account that approximate considered judgments after they have been bolstered using relevant theory from the philosophical literature on well-being. The popular accounts are in the left column and corresponding considered judgments in the right column.

As I argue in paper 4, popular accounts can be interpreted as giving voice to perspectives echoing theories from the philosophical literature on well-being, namely from both the subjective and the objective tradition of well-being theory. Since we ought to avoid, on reflective equilibrium as a method, discarding considered judgments that have passed initial scrutiny without good reason, a pluralistic theory of well-being seems like the most plausible candidate for conceptualizing illness severity. This is because it is the only theory that can account for all the popular intuitions. That is, we should plausibly claim that illness severity is a function of the badness of an illness and that this badness involves both subjective components, such as suffering, and objective components, such as mobility and autonomy.

Such a conceptualization of severity brings with it significant theoretical problems. Most importantly, including both subjective and objective components of well-being in the theory opens it up to objections directed at both subjective and objective theories of well-being. That is, the theory will inherit both the good and the bad from all its “parent theories.” The theory will, for instance, be open to the objection that it allows for a discrepancy between what someone's experienced well-being and the well-being we assign to them. It will also be open to the objection that someone’s subjective experience can be affected by adverse circumstances in ways that makes them poor indicators for claims of distributive justice. I will argue that an objective theory of severity, with subjective components, does better than other theories when it comes to balancing these competing concerns.

Paper 2: Adaptation and How we Need a Pluralistic Theory of Well-Being to Account for our Intuitions

If we are to avoid discarding considered judgements that has passed initial scrutiny, we should arguably make use of an objective theory of well-being with a heavy emphasis on subjective well-being. I will now argue that we have reasons independent of these popular views for accepting such a theory.

The argument starts from the premise that illnesses vary in the extent to which they are adaptable; that is, people do not adapt similarly to all illnesses. People for example adapt significantly to many illnesses that impact bodily functions such as paraplegia but do not adapt significantly to many illnesses characterized primarily by their subjective effects, such as depression or chronic pain.[11] This difference has been noted often in the literature, primarily because of a fear that accounting for the effects of adaptation will lead to a lower priority to illnesses that people do adapt to. If, for example, we combine a subjective theory of well-being

---

[11] Talking about the effects of adaptation on chronic pain is complicated by the fact that pain can plausibly be interpreted in two ways: as referring to the subjective experience of pain, or as the neurological activity associated with pain. One could maybe argue that adaptation to the latter is possible, leading to a reduction in the former. Adaptation to depression seems more like a conceptual impossibility. If I adapt to my depression to the extent that my quality of life is better, has not my depression improved?

with a welfarist conception of distributional justice and note that people with paraplegia are more or less just as happy as non-paraplegic people, the result would be that we do *not* have (strong) reason for using resources on treating paraplegia. For many people this is an unacceptable result, like the unacceptability of allowing adaptive preferences to influence distributional decisions. This phenomenon is therefore considered an argument against accounting for adaptation when assessing the severity of an illness.

What has not received the same attention in the literature is that there is a flipside to the problem of adaptation: the problem of non-adaptation. What I mean by the problem of non-adaptation is the fact that there are illnesses, such as depression, that people do not significantly adapt to. In my paper I argue that this is morally relevant for two primary reasons. The first is that adaptation is, all else equal, clearly a good thing. Adapting to illness reduces suffering, and not adapting thus means more suffering compared to adapting. Suffering is (as I argue in depth in the paper) a relevant factor according to any plausible theory of well-being relevant for health. Secondly, not counting adaptation leads to a systematic failure in tracking the effects that illness has on well-being. This is because whether a person adapts to their illness impacts how much they suffer from the illness, and this is relevant according to any plausible theory of well-being.

In conclusion, adaptation is clearly relevant according to any theory of well-being relevant for health. Illness severity is, at least in part, a function of how bad an illness is for a person. It follows that we should be sensitive to adaptation when assessing the severity of an illness.

This leaves us with the fact that adaptation is, sometimes, intuitively not a good thing all things considered. Menzel et al. (2002), for instance, list the following forms of adaptation to illness: cognitive denial of dysfunction, suppressed recognition of full health, skill enhancement, activity adjustment, substantive goal adjustment, altered conception of health, lowered expectations, and heightened stoicism. Some of these forms of adaptation are intuitively bad. Cognitive denial seems, for example, not to be the kind of adaptation that we ought to encourage. And yet, there is a reduction in suffering associated with even these forms of problematic adaptation that is clearly relevant to someone's health related well-being, and thus to the severity of their condition. In my paper I argue that the best way to account for these opposing considerations is to accept an objective theory of well-being, albeit one with a strong focus on subjective well-being. Our task them becomes to weigh these opposing concerns. Accepting such a theory of well-being will also help with understanding our intuitive judgments on the severity of a number of illnesses where there is a large discrepancy between subjective well-being and the severity of the condition. It can for example account for why Downs's syndrome is bad for a person, even though people with Downs's syndrome have a high degree of subjective well-being (Skotko et al., 2011).

Accepting an objective theory of illness severity will arguably mean that we fail to show people secondary recognition respect, in the sense that people are not the final arbiters of their own good (Terlazzo, 2014). Against this I follow Serene Khader (2012) in arguing that it is not disrespectful to question some of a person's values. On the theory I propose, we should not doubt a person when they report their subjective well-being, but we should not necessarily accept their verdict on their objective well-being. A person is the final arbiter of some but not all aspects of their health-related well-being.

Congruence Between Theoretical and Public Perspectives

On the issue of the currency of severity the popular views and my theorical arguments are in alignment. To account for the plurality of considered judgments on what makes an illness severe we must accept a pluralistic theory of well-being that includes both subjective and objective factors. To account for our intuitions on the normative relevance of adaptation we must similarly accept a theory of well-being that includes both objective and subjective factors: in effect, an objective-list theory of severity.

Interestingly, this way of thinking about health- related worse off-ness, as composed of objective and subjective factors, is somehow in line with how health- related quality of life is measured for purposes of health economic analysis. In the EQ-5D descriptive systems, for instance, health states are defined through five dimensions: mobility, self-care, usual activities, pain and discomfort, and anxiety and depression (Dolan, 1997). Two of these, pain and discomfort, and anxiety and depression, are clearly subject dependent (subjective), whereas the remaining three are subject independent (objective).[12] On the other hand, the values used for health economic analysis are based on people's preferences over various of these health states. So, what we have with the EQ5Ds and associated quality adjusted life years is something of a mix of various perspectives on health-related prudential value.

---

[12] While some of the objective factors can be open to interpretation and highly individual, they are still subject independent. See footnote 5.

My proposed objective theory of illness severity is in many ways similar to Daniel Hausman's theory of the value of health. Hausmann (2015) argues that capturing the value of health requires us to account for both activity limitations and health related feelings. His theory of the value of health includes opportunities (an objective factor) and health-related feelings (a subjective factor). Importantly, Hausman would not use the terms objective and objective factors to describe his theory. He does not believe that the value of health should be measured by the impact of health on well-being. He is instead primarily interested in the *public value* of health. The public value of health is different from the private value of health because the state has interests stemming from its responsibilities to its citizens: facilitating opportunities for worthwhile lives and reducing suffering (Hausman, 2015). But yet, there are interesting parallels between his theory and the objective well-being-based theory that I have argued for. The opportunities that Hausman argues are important are to a large extent co-extensive with the objective factors that I argue are important regardless of their impact on subjective well-being. Objective factors such as mobility are for example important partly due to the opportunities granted by being mobile. While Hausman argues that the fact that many people with mobility issues have a high degree of subjective well-being makes well-being a bad way of accounting for the value of health, I think we can argue that people with mobility issues have a lower degree of (objective) well-being than if they were mobile, all else equal. At least a part of the difference between Hausman's theory and the theory that I argue for is thus due to a difference in how we use the term 'well-being'. There are also interesting parallels in the way that the value of health states should be assigned. Hausman rejects preference elicitation and instead favors a process that is primarily interested in the public due to the importance of their values, arguments, and experience. I believe that a publicly informed reflective equilibrium process is plausible for the same reasons as Hausman's proposal.

Aggregating Severity

As I argued in the introduction, we must answer two questions in order to conceptualize the badness of severe illness. The first is the question of what the relevant kinds of bads are. The second is the question of *if*, and if so, *how* to aggregate these bads. I have already argued that a satisfactory account of the relevant kinds of bads must include both objective and subjective factors. I now turn to the question of aggregation: how to add up these goods and bads. I will first describe how the popular accounts think of aggregation, before making a theoretical argument. I will argue that severity is, or ought to be, an aggregative concept in the unlimited sense.

The Public Perspectives

The public perspectives I draw on for this section are published in Stenmark et al. (2023), which can be consulted for a more detailed account. They are not included in paper 4, which focuses solely on relevant goods and bads. I here summarize what the public perspectives are on aggregation. In contrast to the public perspective on specific goods and bads, our participants did not provide perspectives explicitly about aggregation. I therefore use the participants views on the relationship between age and illness severity as a proxy for views on aggregation.

On the "Natural life span" account, illness in the very young and very old is severe and not severe respectively, but age of onset is not continuous with severity. The death of a 40-year-old would for example not be considered more severe than the death of a 45-year-old.

On the "Severity is subjective" account, age is considered important. In addition to agreeing that illness that affects the young is more severe, this viewpoint is associated with accepting a continuous relationship between age and severity. That is, participants agree with the statement "the younger, the more severe".

On the "Objective measures and triage" account, age is relevant for severity, but participants are not as concerned with the difference between the young and old as those in the "severity is subjective" account. Participants agree with the statement that "the younger, the more severe".

On the "Functioning and quality of life" account, age is fully rejected as a relevant determinant of severity.

I focus on the views about the importance of age for severity in the different viewpoints because they to some extent can be interpreted as reflecting differences on whether, and to what extent, severity is an *intra-personally aggregative* concept. They also seemingly reflect different perspectives on the limited aggregation account, in the sense that at least one of the accounts claim that small differences in age do not matter for severity. On unlimited theories of aggregation small goods and bads can add up to outweigh larger goods and bads (Carlson, 2000; Hirose, 2007). On an unlimited view, the relative severity of an illness should be a result of *all* the consequences of the illness that affect the individual. If an illness leads to pain of the same intensity for four years rather than for three years, it is a more severe illness.

Severity operationalized as absolute QALY-shortfall implies unlimited aggregation, in that any (measurable) time lost to illness or spent in reduced health counts toward the severity of a condition. If we accept unlimited aggregation, severity indirectly tracks age, at least in terminal or chronic illness, because younger people have more years to lose to illness. In this sense, the "Severity is subjective" and "objective measures and triage" views seem to presuppose an aggregative view where an illness is more severe the younger the patient. This is in line with unlimited aggregation, as favoured by utilitarian and prioritarian theories. On the "natural life span" view, illness is particularly severe in the very young and less severe in the very old. This is in line with a threshold view or a limited aggregation view. On a threshold view it matters greatly whether someone is above a specified threshold; perhaps, for example, all illness below 18 years is particularly severe, whereas all illness above 80 years is less severe (in line with a 'fair innings' type argument (Harris, 1985; Williams, 1997)). On a limited aggregation view (see for instance Kamm, 1993; Temkin, 2014; or Voorhoeve, 2014), what matters are relative differences in age. An age difference of 20 years may merit a differential judgment of severity whereas a difference of five years may not, so that a terminally ill 40-year-old is as severely ill as a terminally ill 30-year-old, but the same would not be said about a 20- and 40-year-old.

Alternatively, a complex theory of aggregation is necessary due to the importance of goods that defy simple aggregation. Being a parent, having a successful career, or living a meaningful life are goods that seemingly belong to a life as a whole. It is, for example, not obvious that it makes sense to argue that one year of meaning is better than two years of meaning, or that the value of parenthood can be decomposed into discrete "time slices" of being a parent. David Velleman (1991) has argued that we need to look at a life as a whole to determine its quality, somewhat in line with Aristotle's (2012) conception of eudaimonia. The "functioning and quality of life" view, rejecting the relationship between age and severity while endorsing complex goods, perhaps also requires a complex theory of aggregation in line with Velleman (1991) or Aristotle (2012). Alternatively, it could be interpreted to reject aggregation in line with Nord's (2005) proposal. One could also make an argument to the extent that the "natural life span" account is somewhat aligned with a complex theory of aggregation.

It may strike the reader that the theoretical reading of the popular accounts in terms of aggregation is much more speculative than the reading in terms of well-being. What I mean by this is that while the participants have provided straightforward statements on the various bads of illness, they have not provided straightforward statements on whether these bads can be aggregated. In fact, I may simply have read all the possible positions on aggregation into their statements. At this point it seems worthwhile to return to one of the points I made during the theoretical work on popular views as considered judgments. I argued that we have reason to suspect that a view is more considered if it regards a topic we can assume the participants to be familiar with. Issues of aggregation are notoriously theoretical and technical, and as such do not easily lend themselves to being carefully considered by laypeople. Aggregation may be an issue where gathering considered judgments is not particularly useful.

Still, in line with the argument for "bolstering" the popular views, it seems that popular views can be interpreted in line with various theories of aggregation. It is thus possible to argue for the inclusion of considered judgments supporting all the relevant theories of aggregation. This is a problem for a reflective equilibrium approach to conceptualizing severity because a reflective equilibrium cannot plausibly keep most of the considered judgments on aggregation. For example, limited and unlimited aggregation are mutually exclusive theories. This thus seems to be a case where we must dismiss at least some of the popular judgments. I now turn to arguments to the effect that we should conceptualize severity as a fully aggregative concept; that is, we should reject all judgments except those endorsing an unlimited aggregation approach.

Paper 1: Why Limited Aggregation Should be Rejected when it Comes to Age and Severity

While gathering perspectives on illness severity, a specific notion came up several times: it seemed natural to many that it is more severe for a young person to be ill than for an older person. Many would for instance endorse the conclusion that terminal illness is more severe for a 30-year-old than for a 60-year-old. Still, the same individuals were reluctant to claim that a terminal illness is more severe for a 45-year-old than for a 50-year-old. Many seem to accept something like Small Differences:

SD: small differences in age of life do not matter for illness severity.

SD cannot be accounted for within a utilitarian, prioritarian, or egalitarian theory of distribution. It is, all else equal, worse to die earlier rather than later, and to be ill for longer rather than shorter. On utilitarianism, our reasons for action straightforwardly track our opportunities to promote the good. On prioritarianism, our reasons for action track our opportunities to promote the good moderated by how badly off the individuals we aid are. What matters on prioritarianism is how badly off an individual is and how much we can aid them. The 45-year-old is clearly worse off than the 50-year-old, in the sense that it is worse to die five years earlier. On an egalitarian theory, our reasons for action similarly follow our opportunities to promote the good. The general problem is that utilitarianism, prioritarianism, and egalitarianism are consequentialist, in the sense that our reasons for action straightforwardly track worse-off-ness or goodness and badness. There is therefore no room for a distinction between the good and our reasons for intervention.

Sufficientarianism may seem to give the right answer when discussing the 30- and 60-year-old and the 45- and 50-year-old. Perhaps 30 years is below a relevant threshold, whereas 60 is not? And perhaps the 40- and 45-year-old are both either above or below the same threshold? But sufficientarianism has trouble accounting for SD when we introduce new cases. Let us say that we want to say that a 52-year-old is less severely ill than a 32-year-old (all else being equal besides age, e.g., expected years of survival etc.). We would then need a threshold between these two. We could continue to construct cases, requiring ever new thresholds. As we construct an endless number of new cases, the number of thresholds would approach a continuous line, approaching an unlimited aggregation approach. Sufficientarianism also struggles with problems of arbitrariness and sharpness. The thresholds seem arbitrary (why 30 years, if that is where you would like to put the limit, and not 29 or 31? etc.), and the difference in treatment between people just above and below a threshold seems difficult to accept.

I argue that the only way of making sense of SD is by drawing on the literature on limited aggregation, also called the relevance view, that was discussed in the section on aggregation. Perhaps the difference between the 45-year-old and the 50-year-old is too small to be relevant in deciding whom to aid? The relevance view has the advantage of allowing us to both claim that the 45-year-old is worse off, while not accepting that this has the implication that the 45-year-old should have priority over the 50-year-old on this ground. Severity understood in this way amounts to *morally relevant worse off-ness.*

SRW: an illness is more severe than another if it is worse, and if this worse-ness is relevant for deciding whom to aid.

The problem with understanding severity as morally relevant worse off-ness is that SRW is not a transitive relation. More precisely, the less severe than- and the equally severe as-relations are intransitive. Let us say that a difference in age must be at least nineteen years for it to be morally relevant in deciding whom to aid. We should save a 30-year-old over a 50-year-old but not a 45-year-old over a 50-year-old. But now image that we must decide on the severity of terminal conditions of a 30-, 40-, and 50-year-old. The 30- and 40-year-old are equally severely ill, as are the 40- and 50-year-old. So, 30=40=50, but also 30≠50. Clearly, SRW is an intransitive relation. In addition, as one sees from the example, SRW is sensitive to irrelevant alternatives: whether a 40-year-old is present changes our judgment of the relative severity of the 30- and 50-year-old. This intransitivity and sensitivity to irrelevant alternatives makes the task of providing a privileged priority ordering for the purposes of priority setting difficult.

I explore some options for constructing a priority order that have been suggested in the literature in paper 1, but I want to address an option that I did not consider in the paper. One possibility would be to first exclude from priority those that are dominated by someone else in the set, and then choose between the remaining options by some criteria. In our above example, with the 30-, 40-, and 50-year-old, the 50-year-old is the only person who is dominated by someone else in the set (the 30-year-old). We could then exclude the 50-year-old and subsequently decide between the 30- and 40-year-old by some procedure. On such an account the acceptability of differential treatment would be based on whether you are relevantly older than the youngest person in the set. While this may be a technically feasible procedure, I am not convinced that it can be morally justified while accepting the relevance view. It is true that the 50-year-old is dominated by the 30-year-old, but they are not dominated by the 40-year-old. If we only had to deal with the 40- and 50-year-old, we would treat them equally in terms of severity. But simply because there is a 30-year-old in the mix, the 50-year-old is treated differently from the 40-year-old, even though we also claim that there is no morally relevant difference between the 40- and 50-year-old. In this sense this possibility resembles an option I explore in paper 1: create a system where severity is based on a person's ratio between less severe than-relations and more severe than-relations in the population. The youngest are more severely ill than most people in the population, whereas the reverse is true for the oldest. Prioritizing on such a basis would allow us to prioritize the young over the old. As I show in paper 1, doing this "saves" the relevance view by making it practically irrelevant. Prioritizing based on dominance suffers from the same problem: it does not respect the point of the relevance view. We end up treating people differently whom we claim should *not* be treated differently.

Given that endorsing severity as morally relevant worse off-ness requires us to either adopt a sprawling system of thresholds or accept that severity is an intransitive relation unsuitable for systematic priority setting, perhaps we should reject SD? There are several reasons why we should be skeptical of SD. First, our intuition may be epistemically rather than normatively grounded. When asked to judge the relative severity of a 45- and 50-year-old, we may simply be unsure as to whether the difference is large enough to reliably track what matters. Perhaps the 50-year-old has a longer life ahead of them, have had a harder life so far, or have younger children. The last point, about children, leads the way towards the idea that we may also be influenced by a life stages kind of intuition. Perhaps we are willing to say that the 60-year-old is less severely ill than the 30-year-old because they so clearly are in different life stages. The same is difficult to say about the 45- and 50-year-old. Even though we have stipulated that the patients are similar in all relevant aspects in the cases we may not be able to control for these other sources of our intuitive response. Finally, we should also be skeptical of the intuition because there is something suspect about the idea that time spent alive can be irrelevant. Our years of life are fundamentally all we have. I argue that SD should be rejected and that we should accept a continuous relationship between age and illness severity. We should accept, all else equal, that the 45-year-old is more severely ill than the 50-year-old.

Can the Argument be Extended to Well-Being?

Assuming that my argument is sound, limited aggregation should be rejected when it comes to the relationship between age and severity. I will now examine the possibility of generalizing the argument to good life years. Generalizing the argument requires arguing that small differences in (health related) well-being should matter when assessing severity.

First off, the same problems of transitivity and sensitivity to irrelevant alternatives show up if we attempt to use a limited aggregation account to rank alternatives that differ in terms of well-being. Let us say that health-related well-being ranges from +1 to 0. Consider the following pairs of patients:

Case 1

- Patient *A* has a health-related well-being of 0.2.
- Patient *B* has a health-related well-being of 0.8.

Is one of these cases more severe, and if so, which one? Next, consider Case 2 below:

Case 2

- Patient *C* has a health-related well-being of 0.4.
- Patient *D* has a health-related well-being of 0.45.

Is one of these cases more severe, and if so, which one?

My intuitive judgments on these cases, as in the cases in paper 1, is that patient *A* is more severely ill than patient *B*, but that patient *C* and *D* are equally severely ill. This seems to follow the pattern of SD_{WB}:

SD_{WB}: small differences in health-related well-being does not matter for illness severity.

This principle is structurally the same as the principle for age (henceforth SD_A). It once again leads to a lack of systematic priority ordering (if accepting something like SRW) or requires an infinity of thresholds (if we go the way of thresholds). We thus seem to have the same exact theoretical reasons for being skeptical to SD_{WB} as to SD_A.

There seem to be a couple of important differences between well-being and age. Firstly, it is controversial whether we can make interpersonal comparisons of well-being.[13] Age is simpler in this regard in that it is easy to make comparisons as to how long someone has been alive. Being able to make comparisons of well-being is obviously required for even making sense of the cases in the manner that I have done. In the following I assume that we can make these kinds of comparisons.

Secondly, I argued that one reason for discarding SD_A is that irrelevance seems implausible when discussing life-years: time alive is, in a very real sense, all that we have. The same can maybe be said in some shape or form about well-being (particularly if we argue that

---

[13] See John Broome (2004) and Timothy Scanlon (1991) for discussion of the problem of interpersonal comparison and two (very) different responses.

well-being essentially means “what makes a life go better or worse”) but seems obviously wrong when applied to health-related well-being. While health is important, there is more to life. Small differences in health-related quality of life may not always have significant effects on overall quality of life, because other aspects of life, sometimes and to some extent, are more important. While we have the same theoretical reasons for being skeptical to SD_{WB} as to SD_{A}, we may not be as sure that the differences are truly relevant.

However, while the “all we have”-argument does not apply for health-related quality of life, the argument that our intuition may be epistemically grounded seems stronger when discussing health-related well-being. I previously suggested that the intuition underlying SD_{A} may be epistemic rather than normative. That is, we may simply be uncertain of whether small differences in age accurately tracks what matters in terms of life lived. Similarly, a parallel intuition underlying SD_{WB} may be epistemic rather than normative. Because health-related quality of life interacts in complex ways with quality of life, it seems even more likely that the intuition is fundamentally about uncertainty rather than the irrelevance of small differences. If our intuition is shaped by epistemic reasons rather than normative ones, it may be unstable as a source of normative reasons.

As already noted, SD_{WB} leads to a similar lack of systematic priority setting as SD_{A}, and hence, we have reason to distrust our intuitive responses. Nevertheless, other differences between age and quality of life may be relevant. Describing quality of life as I have done in the cases, as on a fine-grained linear scale, seems more plausible on some conceptualizations of well-being than others. Particularly, it seems possible on hedonistic or preferentialist theories of well-being. As I will argue that accounting for our intuitions on illness severity requires adopting objective well-being factors as well as subjective ones, we must discuss whether these factors are amenable to aggregation.

Accepting my argument against SD_{WB}, as well as my argument against $SD_{A,}$ amounts to rejecting limited aggregation when considering illness severity. However, what about factors that cannot straightforwardly be aggregated? David Velleman (1991) argues that life must be judged as a whole. If, for example, my misery during the writing of my PhD-dissertation ends up with it being accepted this misery is in some sense justified. The value of this period of my life will therefore be different than if my dissertation is rejected. Can the same be said for the objective components of health? Some factors seem straightforwardly amenable to aggregation: the badness of several years of lack of mobility is presumably some function of the number of years and the badness of separate times.[14] The “objective” factors included in the EQ-5D, mobility, usual activities, and self-care, are seemingly temporal goods. Thus, there is no reason to assume that they cannot be aggregated across time. Other factors may work differently. Being unable to work for a couple of years may, for example, be qualitatively different from not being able to work at all during a life: the latter may present a different type of deprivation, such as the lack of a fulfilling career, meaning, or prestige. These deprivations are not clearly temporal, in the sense that we can locate them at specific times in a person’s life. While I lean towards the idea that all relevant badness can be located in time, it is possible to accept that non-temporal goods contribute to the badness of illness while endorsing unlimited aggregation. Plausibly, a life spent without a career may represent an additional bad to be added to the badness located in time. Our conception would then be more complex, but not in opposition to the idea that all badness contributes to illness severity. While we should hold on to the notion that most of the relevant factors for illness severity can be fully aggregated, we may want to allow for the possibility that there are factors that can only be aggregated in more complex ways. In line with Velleman’s or Aristotle’s arguments, there may be components of health-related quality of life that must be seen more holistically. We must, if this is the case, figure out what these components are, why they matter, and how they should influence assessments of severity. If these factors are to influence priority setting, we have no choice other than to assign them some value commensurable with the factors that allow for aggregation.

Rejecting Non-Aggregation

I have argued that we should reject a limited aggregation approach to the relationship between badness and severity. To fully make the case that we should accept unlimited aggregation, I must argue that we should not accept a non-aggregative conception of severity.

Rejecting aggregation completely leads to a conception of severity that neither tracks our intuitions on severity or who the worse off are. First, many of our intuitions on the severity of illness track how much one is deprived of due to illness. A concept of severity that is not aggregative cannot track this deprivation. I have also argued that severity is important because we care about the worse off. On a deprivationist theory of the badness of death, the most likely

---

[14] The function does not necessarily have to be strictly linear. Presumably, for instance, lack of mobility is worse when it last longer than your number of paid sick days leading to worse economic prospects. Perhaps some forms of illness make you depressed after a certain time etc. What is important is that we will be able to say that 2X years of the disability is worse than X years of the disability, and also how much worse.

contender and the only one that can explain our intuitions on the badness of dying, you are clearly worse off when dying earlier rather than later. In addition, on any plausible account of what is good or bad for us it matters for how long something lasts. Being in pain for a year is worse for you than being in pain for a day. Conceptualizing severity in a fully non-aggregative manner means that it neither tracks who the worse off are in terms of health nor our intuitions on what constitutes severe illness.

Rejecting Thresholds

I have already discussed how thresholds cannot account for SD_A and SD_{WB}, but I believe that we should reject the idea of thresholds almost entirely for conceptualizing illness severity.

First, threshold theory seems ill-aligned with how we use the term severity. We commonly say such things as "*A* is more severely ill than *B*, but *C* is more severely ill than *A*". Operationalizing severity using a limited number of thresholds will require seemingly arbitrary restrictions on how we use the term, and using an unlimited number of thresholds will be indistinguishable from operating with a continuous function between severity and worse off-ness. This descriptive problem mirrors the normative problem of operationalizing severity in line with a threshold model. A threshold marks a boundary where we treat people over and below the threshold differently. For the threshold model to be different from a continuous function, the difference must be significant. This differential treatment presents two problems for the proponents of thresholds. The first problem is that it seems intuitively wrong to treat people who are extremely similar but just on different sides of the threshold differently. Should we, for instance, treat 29- and 30-year-olds very differently, while not making the same separation between, for example, a 31- and 35-year-old? The difference in treatment close to the threshold seems much too sharp. There seems to be an interesting lack of a parallel between thresholds in different settings. For example, the thresholds for criminal responsibility or for voting rights do not seem intuitively problematic in the same sense as a threshold for severity. While we may argue about the placing of the thresholds for voting rights or criminal responsibility (which are probably somewhat arbitrary), we seemingly accept a sharp distinction between rights and responsibilities based on small differences that cross the threshold. One reason for this difference may be pragmatic: if we conceptualize these rights and responsibilities in an either/or fashion we cannot escape the need for a threshold. The same is not true for illness severity. We are not forced to accept a threshold-theory because severity is easily conceptualized as ranging along a continuum. If we can avoid the problems associated with a threshold theory, I do not see why we should choose not to.

The second problem is how to specify the locations of the threshold in a non-arbitrary manner. Why should we, for example, place a threshold at 30 years old rather than 18 years old? This challenge may be overcome to some extent if there are “natural” boundaries within a life, perhaps corresponding to life stages such as childhood, young adulthood, childbearing age, productive age, old age and so forth. It may be that certain perfectionist goods follow such a pattern. If this is the case, we may have to accept some form of threshold theory to value these goods correctly. It may still be impossible to place thresholds in a fully non-arbitrary manner (when, for example, does old age begin *specifically*?). However, we may choose to accept a certain arbitrariness to be able to value these goods. As such, I will leave it as an open question whether *some* goods should be operationalized using a threshold model. This would mean that most goods, such a subjective well-being and mobility, should be aggregated straightforwardly, whereas some goods, such as the capability of having children, may have to be aggregated in line with a threshold theory. But if these goods are to be part of our overall assessment of severity, we must still be able to add them to the overall amount of goods and bads; that is, our overall theory should still be one accepting unlimited aggregation.

We should (mostly) only keep the popular judgments supporting unlimited aggregation and reject the other ones. This conclusion is based on sound normative arguments, theoretical considerations, and reasons of coherence. We should accept unlimited aggregation when assessing severity.

What Should Our Conception of Severity Look Like?

Our examination seems to allow us to make at least some claims about what a publicly informed reflective equilibrium on illness severity will look like. First, severity is a measure of health-related worse off-ness. Second, the associated distributive aspects of severity are in line with some form of (fully or partially) consequentialist distributional theory; severe illness is morally significant because it is bad *for patients.* Third, we should accept unlimited aggregation when assessing severity. Fourth, a publicly informed equilibrium on severity should include a pluralistic theory of the badness of illness; it should include both objective factors such as functioning and subjective well-being.

The Way Forward

Before concluding, I want to briefly discuss an alternative to discussing severity only in relation to well-being and make some preliminary remarks about the possibility of making further headway on an objective theory of illness severity. Two difficult questions must be answered for my theory to be operationalizable: (1) which factors should be on the list, and (2) how we should measure the factors and decide on their relative weights.

Well-Being or Societal Concerns?

I have explicitly couched by arguments in terms of well-being. That is, I have argued that illness is bad because it is bad for people and advanced an objective list theory to account for this badness. If we do not accept that objective features are part of well-being, it still seems possible to argue for a similar conclusion from the perspective of a mix of care for individual well-being and societal concerns. Hausman (2010) has for example argued that "the criterion of evaluation from a public perspective should be how severely health states limit the range of alternative lives and pursuits that are open to individuals.". The importance of both subjective and objective features of illness could perhaps be argued for on the basis that the subjective features of illness care for individual well-being, whereas the objective features cover the public's interest in reducing activity limitations associated with illness. Alternatively, we could argue, like Hausman (2015), that the state has both an interest in opportunities for its citizens *and* a duty of care for their well-being (in particular in reducing suffering).

Which Factors Should be on the List?

While I am sympathetic to Hausman's suggestion that we should value health based on the impact illness has on opportunities and subjective well-being, my theory is somewhat different in that I argue solely from the perspective of how illness affects the affected person. To make headway on operationalizing my theory I therefore cannot rely on societal concerns. The questions of which objective factors to include on the list must be decided by way of an acceptable conception of well-being. This may represent a problem, because creating a non-arbitrary list of factors is a challenge for any objective list-theory of well-being. But, we may be able to use a suggestion from Rice (2013): we can make headway on an objective list by relying on peoples considered judgments. In line with my arguments in paper 3, I would even suggest that questions of well-being are questions where public input is likely to be of value given adequate processes of deliberation. Hausman (2015) suggests, when discussing how to value health, that valuing opportunities involves some decision on which opportunities have more value. To what extent should we be concerned with ability to have children, participate in the workforce, or enjoy the arts? These are the kinds of questions that are relevant when assessing the importance of factors such as fertility, mobility, and cognitive- and perceptual abilities, and they are not questions where experts are well positioned to answer on their own. The public should be consulted, but since these are fundamentally normative questions, we should query the public to elicit their values, their experience and their (considered) judgments.

Measuring and Weighing the Factors

To prioritize based on severity, we must be able to measure the relevant factors and weigh them against each other. How should we measure the subjective factors? Which objective factors are more important? And how do we weigh them all to end up with a verdict on total severity of illness? When assessing subjective well-being, I suggest that we should rely on the reports of people suffering from relevant illnesses and empirical studies of subjective well-being. When it comes to the questions of how to value objective factors and weigh them against subjective factors, I suggest that we rely on a publicly informed reflective equilibrium. If such a process involves adequate information, we can, thorough deliberation and input from people with relevant experience, make headway towards an operationalizable objective theory of illness severity.

Subject- dependence and -independence

Measuring severity in line with an objective theory of well-being with a subjective component would ideally involve measuring the different factors in ways that respect their character. Concretely, this means that the subjective factors should be measured in a way that respects the subject-dependent properties of these factors and that objective factors should be measured in a way that respects the subject-independent properties of these factors.

Currently, the most common way to judge the value of health states is to elicit preferences from the population. Many have expressed skepticism towards relying on preferences or towards specific preference elicitation methods. Daniel Kahneman (2008) has for example pointed out problems with current health state valuation practices, there are arguments to be made against relying on preferences from the perspective off all the relevant theories of well-being that are not based on preferences (see the previous section on well-being), and many have raised problems related to specific elicitation methods (see for example Rand & Kesselheim (2021) for an overview of problems with QALYs). Hausman (2015) offers several arguments, of which I believe that three are particularly important: (1) Elicited preferences are a guide to well-being if they are rational, informed, and self-interested. This is not the case for the preferences typically elicited in healthcare valuation. As Hausman argues, "Those surveyed know too little and have thought to little" (Hausman, 2015, p89). Argument 1 gets some of its force from the next argument. (2) People make mistakes in assessing health states due to biases and faulty information. As we have seen, people are for example prone to focusing illusions, they underestimate the effects of adaptation, and they rely on faulty information about what it is like living with illnesses. (3) Even if the two previous arguments could be countered by improving evaluation practices, preferences for health states are not basic. When reporting their preference for a health state, people provide preferences based on what they think about the health state. When reporting my preference between a longer life in a wheelchair and a shorter life fully mobile, I, for example, think about what it would be like to need a wheelchair. Would I be able to play sports? Would I be able to continue my work? How would it affect my dating life? My preference for the health state is informed by my assessment of how my life, and the lives of people I care about, would be impacted by illness. These are the things that matter. That is, we should be concerned with "the factors to which preferences depend" (Hausman, 2015, p103). Scanlon (1975) argues similarly, stating that what we are after is not how people feel about different benefits, but rather "the reasons for which these benefits are considered desirable." (T. M. Scanlon, 1975).

These arguments against the relevance of preferences are pertinent to the question of how to assess severity. What we are after, on both the subjective and objective theories of well-being, is how much someone's well-being is impacted by the illness in question. We are not interested in *how people believe* that someone's well-being is impacted by the illness. As Hausman argues, there is no reason to believe that people have more insight into these matters than experts. So why defer to them?

If we cannot defer to people's preferences, how should we measure the subjective value of health? If we are to measure the subjective components of illness, and in particular the suffering associated with illness, we need a measure of the impact illnesses have on subjective well-being. That is, to what extent does a given illness lead to suffering? This is not plausibly a question that the public is in a good position to answer. As I argued in paper 2, given the discrepancy between hypothetical and experienced health and the fact that (at least some of) this difference is due to adaptation, we ought to privilege experienced health measures when assessing the subjective effects of illness. If we are interested in suffering for its own sake, as I argue that we should, we could also rely on measures of quality of life, knowledge about the pain caused by various illnesses and other sources of knowledge. Psychologists have for example made strides in measuring happiness, emotional well-being, life satisfaction, and moment to moment happiness (Diener et al., 1999; Kahneman & Krueger, 2006) What matters is not, in and of itself, who provides the information, but rather whether the information is accurate.

Objective factors present their own challenges. On the surface, some of them are easier to measure than the subjective factors. It is, for example, rather obvious whether a person is mobile or not. What makes these factors challenging is thus not measurement per se but rather the valuation of the factors. How do we evaluate these factors when they are supposed to be subject-independent. That is, when preferences and subjective experience are *not* what matters?

First, let me make an important qualification. These factors, of course, *also* matter because of subjective experience. A part of what makes these factors important is their effects on subjective well-being and this is presumably in large part the parts that are affected by adaptation. These effects will presumably, be adequately measured *as subjective well-being.* What is difficult is assessing the subject-independent badness of these factors. If neither people's experiences nor preferences are important, what do we do? Presumably, if preferences or experiences do not matter individually, they do not matter when aggregated, either.

In line with my examination of how to integrate public perspectives in normative work, I suggest that the ideal process would be a publicly informed reflective equilibrium process. The public should be involved, but not in the sense of polling for preferences. What we are after is not how important people think mobility is. We care about how important mobility is for people. Why should we believe that the public is in a better position to answer this question than researchers? Just as it does not make sense to poll the public in how bad delusional psychosis is, it does not make sense to poll them on how a lack of mobility affects well-being. We should attempt to elicit considered judgments from the public. This process could involve group discussions, making sure that participants deliberate, and that they are exposed to relevant information. We would be interested in the reasons for people's preferences rather than preferences themselves. When we have elicited considered judgments, the work of constructing a theory remains. Fortunately, reflective equilibrium is precisely the kind of methodology that lends itself to assessing these kinds of questions. We need reasons and arguments. We need theoretical reflection informed by considered judgments on what matters in these cases. And we need the ensuing theory to display coherence. This reflective equilibrium would be about the kind of topic where the public is likely to have views approximating considered judgments. People make decisions about what matters in their lives, discuss it with others, and are exposed to views through literature, film, and through other means. I propose that the best way, and perhaps the only satisfactory way, of deciding on the relative weights of the objective factors would be a publicly informed reflective equilibrium process, informed by all relevant considered judgments.

I have not made anything resembling a complete argument to the effect that the weighting of the objective factors of severity must be decided on through a reflective equilibrium process, but it seems difficult to argue for any other method. Preferences are ruled out because of the importance of subject-independence. Measures of subjective well-being are ruled out because what we are after is precisely factors that are not subjective. Aggregation of preferences or subjective well-being is similarly ruled out. One option that seems plausible at first sight is to vote; that is, people decide collectively. But on second consideration, voting also seems problematic. There is no reason to suppose that a poll will reveal the true badness of a health state, and no reason to suspect that a judgment is sound due to being accepted by a majority.

We are left with the question of how to weigh the subjective and objective factors to assess illness severity. Measuring the subjective factors will, as I have argued, involve either measuring the preferences of people who have experienced the illness in question or measuring subjective well-being directly. The weights of the objective factors will have to be decided on in a (publicly informed) reflective equilibrium process. How to combine the two is an inherently normative question, that should be answered by a process involving our considered judgments, arguments, and theory: a reflective equilibrium process.

Conclusion

I believe that I have made some progress on the two questions I posed in the introduction: that is (1) what is the relevance of subjective and objective factors of well-being to illness severity, and (2) to what extent is severity an aggregative concept. I also believe that I have made progress on how to include popular views in reflective equilibrium processes and that I have exemplified how to prepare views for this inclusion.

I have argued that popular views should play a role in reflective equilibrium processes if the views approximate considered judgments. The ideal process, as argued for in paper 3, involves eliciting views in a setting that promotes deliberation and reflection. Properly moderated focus groups, citizen panels, and in-depth interviews where people are exposed to relevant arguments and perspectives would be ideal. Input from the public would be particularly valuable when the topic is familiar to people, such as when constructing a theory of illness severity.

When we do not have data on popular views that are gathered under these ideal circumstances, we can still make progress. We can search the theoretical literature for arguments and views that reflect the same underlying intuitions that are found among the public to attempt to bolster popular views before including them in the reflective equilibrium process. Given that the reflective process roots out incoherent and ill-founded judgments, this would be helpful when constructing a publicly informed reflective equilibrium even if we end up accepting a model of reflective equilibrium without any substantial filtering of intuitions. I have examined and bolstered a set of popular views and argued that they approximate considered judgments that merit inclusion in a reflective equilibrium on illness severity. I have argued that these judgments can be accounted for most effectively by an objective list-theory of illness severity with a subjective component.

Severe illness is bad for people suffering from illness. I have argued that to account for this badness in a way that respects popular views on illness severity, we should accept an objective list-theory of severity. An objective theory is also in line with my argument stemming from the relevance of adaptation for illness severity. Adaptation reduces suffering, and suffering is relevant on any plausible conception of health-related well-being. We should therefore count adaptation when assessing the badness of illness. Still, accepting a purely subjective theory of the badness of illness could have problematic distributional consequences. Accepting an objective list-theory of the severity of illness would allow us to acknowledge both concerns and weigh them accordingly. An objective theory of illness severity with a subjective component is thus supported both by popular considered judgments and theoretical argument.

I have argued that we should accept a fully aggregative theory of illness severity in that all goods and bads should be counted. While accepting unlimited aggregation involves rejecting many popular considered judgments, it is in the nature of the aggregation question that the options are exclusionary. We cannot accept both limited and unlimited aggregation. We should reject limited aggregation accounts due to problems with transitivity and sensitivity to irrelevant alternatives. We should also reject non-aggregation accounts because they cannot account for our judgments on what severe illness is. If there are complex goods, we must still add the value of these goods to the goods that can be straightforwardly aggregated. I remain open to the possibility that some goods should be conceptualized using thresholds, perhaps in line with a life stages-view, but these goods must still somehow be included in the overall assessment. The best option for a theory of illness severity is to accept unlimited aggregation.

In conclusion, I have made the case for a fully aggregative objective theory of illness severity with a substantial focus on subjective well-being. Such a theory is supported by popular considered judgments and theoretical arguments. Much more work must be done on the specifics of such a theory, and I have made some preliminary points about how to proceed. I suggest that the questions of which factors to include on the list of relevant factors and how to weight these factors to assess severity should be tackled by going through a publicly informed reflective equilibrium process.

Literature

Anderson, E. S. (1999). What Is the Point of Equality? *Ethics*, *109*(2), 287–337. https://doi.org/10.1086/233897

Aristotle. (2012). *Aristotle's Nicomachean Ethics* (R. C. Bartlett & S. D. Collins, Trans.; Reprint edition). University of Chicago Press.

Arneson, R. J. (1989). Equality and Equal Opportunity for Welfare. *Philosophical Studies: An International Journal for Philosophy in the Analytic Tradition*, *56*(1), 77–93.

Arneson, R. J. (1999). Human Flourishing Versus Desire Satisfaction. *Social Philosophy and Policy*, *16*(1), 113–142. https://doi.org/10.1017/S0265052500002272

Awad, E., Dsouza, S., Kim, R., Schulz, J., Henrich, J., Shariff, A., Bonnefon, J.-F., & Rahwan, I. (2018). The Moral Machine experiment. *Nature*, *563*(7729), Article 7729. https://doi.org/10.1038/s41586-018-0637-6

Baderin, A. (2017). Reflective Equilibrium: Individual or Public? *Social Theory and Practice*, *43*(1), 1–28.

Barra, M., Broqvist, M., Gustavsson, E., Henriksson, M., Juth, N., Sandman, L., & Solberg, C. T. (2020). Severity as a Priority Setting Criterion: Setting a Challenging Research Agenda. *Health Care Analysis*, *28*(1), 25–44. https://doi.org/10.1007/s10728-019-00371-z

Begon, J. (2020). Disability, Rationality, and Justice: Disambiguating Adaptive Preferences. In A. Cureton & D. T. Wasserman (Eds.), *The Oxford Handbook of Philosophy and Disability* (pp. 342–359). Oxford University Press. https://doi.org/10.1093/oxfordhb/9780190622879.013.27

Bradford, G. (2015). Perfectionism. In *The Routledge Handbook of Philosophy of Well-Being*. Routledge.

Bradford, G. (2017). Problems for Perfectionism. *Utilitas*, *29*(3), 344–364. https://doi.org/10.1017/S0953820816000418

Brandt, R. B. (1979). *A Theory of the Good and the Right*. Amherst, N.Y.: Prometheus Books.

Brock, D. (2002). The Separability of Health and Well-Being. In C. J. L. Murray, J. A. Salomon, C. D. Mathers, & A. D. Lopez (Eds.), *Summary measures of population health: Concepts, ethics, measurement, and applications*. World Health Organization.

Broome, J. (2002). Measuring Health by Aggregating Well-Being. In C. J. L. Murray, J. A. Salomon, C. D. Mathers, & A. D. Lopez (Eds.), *Summary measures of population health: Concepts, ethics, measurement, and applications*. World Health Organization.

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

Carlson, E. (2000). Aggregating Harms—Should We Kill to Avoid Headaches? *Theoria*, *66*(3), 246–255. https://doi.org/10.1111/j.1755-2567.2000.tb01167.x

Chan, L., Schaich Borg, J., Conitzer, V., Wilkinson, D., Savulescu, J., Zohny, H., & Sinnott-Armstrong, W. (2022). Which features of patients are morally relevant in ventilator triage? A survey of the UK public. *BMC Medical Ethics*, *23*(1), 33. https://doi.org/10.1186/s12910-022-00773-0

Cohen, G. A. (1989). On the Currency of Egalitarian Justice. *Ethics*, *99*(4), 906–944.

Colburn, B. (2011). Autonomy and Adaptive Preferences. *Utilitas*, *23*(1), 52–71. https://doi.org/10.1017/S0953820810000440

Crisp, R. (2003). Equality, Priority, and Compassion. *Ethics*, *113*(4), 745–763. https://doi.org/10.1086/373954

Crisp, R. (2017). Well-Being. In E. N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy* (Fall 2017). Metaphysics Research Lab, Stanford University. https://plato.stanford.edu/archives/fall2017/entries/well-being/

Damschroder, L. J., Zikmund-Fisher, B. J., & Ubel, P. A. (2005). The impact of considering adaptation in health state valuation. *Social Science & Medicine*, *61*(2), 267–277. https://doi.org/10.1016/j.socscimed.2004.11.060

Daniels, N. (1979). Wide Reflective Equilibrium and Theory Acceptance in Ethics. *The Journal of Philosophy*, *76*(5), 256–282. https://doi.org/10.2307/2025881

Diener, E., Suh, E. M., Lucas, R. E., & Smith, H. L. (1999). Subjective well-being: Three decades of progress. *Psychological Bulletin*, *125*(2), 276–302. https://doi.org/10.1037/0033-2909.125.2.276

Dolan, P. (1997). Modeling Valuations for EuroQol Health States. *Medical Care*, *35*(11), 1095–1108.

Dorsey, D. (2017). Adaptive Preferences Are a Red Herring. *Journal of the American Philosophical Association*, *3*(4), 465–484. https://doi.org/10.1017/apa.2017.36

Dreyfus, H., & Kelly, S. D. (2011). *All Things Shining: Reading the Western Classics to Find Meaning in a Secular Age*. Simon and Schuster.

Eftekhari, S. (2021). The Irrationality of Adaptive Preferences: A Psychological and Semantic Account. *Utilitas*, *33*(1), 68–84. https://doi.org/10.1017/S0953820820000321

Elster, J. (1982). Sour Grapes–Utilitarianism and the Genesis of Wants. Sen A, Williams B. In A. K. Sen & B. A. O. Williams (Eds.), *Utilitarianism and Beyond* (pp. 219–238). Cambridge University Press.

Fletcher, G. (2015). Objective List Theories. In *The Routledge Handbook of Philosophy of Well-Being*. Routledge.

Frithjof Norheim, O. (2019). The Badness of Death: Implications for Summary Measures and Fair Priority Setting in Health Care. In J. McMahan, E. Gamlund, & C. T. Solberg (Eds.), *Saving People from the Harm of Death* (p. 0). Oxford University Press. https://doi.org/10.1093/oso/9780190921415.003.0003

Greene, J. D. (2014). Beyond Point-and-Shoot Morality: Why Cognitive (Neuro)Science Matters for Ethics. *Ethics*, *124*(4), 695–726. https://doi.org/10.1086/675875

Gregory, A. (2015). Hedonism. In *The Routledge Handbook of Philosophy of Well-Being*. Routledge.

Griffin, J. (1986). *Well-Being: Its Meaning, Measurement and Moral Importance*. https://philarchive.org

Gustavsson, E., & Lindblom, L. (2023). Justification of principles for healthcare priority setting: The relevance and roles of empirical studies exploring public values. *Journal of Medical Ethics*, jme-2022-108702. https://doi.org/10.1136/jme-2022-108702

Hall, A., & Tiberius, V. (2016). Well-being and subject dependence. In G. Fletcher (Ed.), *The Routledge Handbook of Philosophy of Well-Being*. Routledge.

Harris, J. (1985). *The Value of Life: An Introduction to Medical Ethics*. Routledge. https://doi.org/10.4324/9780203005606

Harris, J. (2020). The Immoral Machine. *Cambridge Quarterly of Healthcare Ethics*, *29*(1), 71–79. https://doi.org/10.1017/S096318011900080X

Hausman, D. M. (2010). Valuing health: A new proposal. *Health Economics*, *19*(3), 280–296. https://doi.org/10.1002/hec.1474

Hausman, D. M. (2015). *Valuing Health: Well-Being, Freedom, and Suffering*. Oxford University Press.

Haybron, D. M. (2016). Mental State Approaches to Well-Being. In M. D. Adler & M. Fleurbaey (Eds.), *The Oxford Handbook of Well-Being and Public Policy* (p. 0). Oxford University Press. https://doi.org/10.1093/oxfordhb/9780199325818.013.11

Heathwood, C. (2005). The problem of defective desires. *Australasian Journal of Philosophy*, *83*(4), 487–504. https://doi.org/10.1080/00048400500338690

Heathwood, C. (2015). Desire-fulfillment theory. In *The Routledge Handbook of Philosophy of Well-Being*. Routledge.

Hirose, I. (2007). Review Article: Aggregation and Non-Utilitarian Moral Theories. *Journal of Moral Philosophy*, *4*(2), 273–284. https://doi.org/10.1177/1740468107079270

Hirose, I. (2014). *Egalitarianism*. Routledge.

Hurka, T. (1993). *Perfectionism*. Oxford University Press.

Jølstad, B., & Juth, N. (2022). Age and Illness Severity: A Case of Irrelevant Utilities? *Utilitas*, 1–16. https://doi.org/10.1017/S0953820822000024

Kahneman, D. (2008). *QALYs versus Experience: A Perspective from Experimental Economics* [Monographs]. Office of Health Economics. https://econpapers.repec.org/bookchap/ohemonogr/000243.htm

Kahneman, D., & Krueger, A. B. (2006). Developments in the Measurement of Subjective Well-Being. *Journal of Economic Perspectives*, *20*(1), 3–24. https://doi.org/10.1257/089533006776526030

Kamm, F. M. (1993). *Morality, Mortality. Volume I: Death and Whom to Save From It*. Oxford University Press.

Kelly, T., & McGrath, S. (2010). Is Reflective Equilibrium Enough? *Philosophical Perspectives*, *24*, 325–359.

Khader, S. J. (2012). Must Theorising about Adaptive Preferences Deny Women's Agency? *Journal of Applied Philosophy*, *29*(4), 302–317. https://doi.org/10.1111/j.1468-5930.2012.00575.x

Kitcher, P. (1999). Essence and Perfection. *Ethics*, *110*(1), 59–83. https://doi.org/10.1086/233204

Lemaire, S. (2021). Subjectivism without Idealization and Adaptive Preferences. *Utilitas*, *33*(1), 85–100. https://doi.org/10.1017/S095382082000031X

Martha, N. (1988). Nature, Function, and Capability: Aristotle on Political Distribution. *Oxford Studies in Ancient Philosophy*. https://philpapers.org/rec/NUSNFA-2

Mayerfeld, J. (1996). The Moral Asymmetry of Happiness and Suffering. *The Southern Journal of Philosophy*, *34*(3), 317–338. https://doi.org/10.1111/j.2041-6962.1996.tb00795.x

McMahan, J., & McMahan, A. P. of P. and R. S. J. (2002). *The Ethics of Killing: Problems at the Margins of Life*. Oxford University Press.

Menzel, P., Dolan, P., Richardson, J., & Olsen, J. A. (2002). The role of adaptation to disability and disease in health state valuation: A preliminary normative analysis. *Social Science & Medicine*, *55*(12), 2149–2158. https://doi.org/10.1016/S0277-9536(01)00358-6

Michel, Y. A., Engel, L., Rand-Hendriksen, K., Augestad, L. A., & Whitehurst, D. G. (2016). "When I saw walking I just kind of took it as wheeling": Interpretations of mobility-related items in generic, preference-based health state instruments in the context of spinal cord injury. *Health and Quality of Life Outcomes*, *14*(1), 164. https://doi.org/10.1186/s12955-016-0565-9

Mitchell, P. (2018). Adaptive Preferences, Adapted Preferences. *Mind*, *127*(508), 1003–1025. https://doi.org/10.1093/mind/fzy020

Nagel, T. (1970). Death. *Noûs*, *4*(1), 73–80. JSTOR. https://doi.org/10.2307/2214297

National Institute for Health and Care Excellence. (2022, January 31). *NICE publishes new combined methods and processes manual and topic selection manual for its health technology evaluation programmes* [NewsArticle]. NICE; NICE. https://www.nice.org.uk/news/article/nice-publishes-new-combined-methods-and-processes-manual-and-topic-selection-manual-for-its-health-technology-evaluation-programmes

Nord, E. (2005). Concerns for the worse off: Fair innings versus severity. *Social Science & Medicine*, *60*(2), 257–263. https://doi.org/10.1016/j.socscimed.2004.05.003

Nord, E., & Johansen, R. (2014). Concerns for severity in priority setting in health care: A review of trade-off data in preference studies and implications for societal willingness to pay for a QALY. *Health Policy*, *116*(2), 281–288. https://doi.org/10.1016/j.healthpol.2014.02.009

Nozick, R. (1974). *Anarchy, state, and utopia*. Basic Books.

Nussbaum, M. C. (2011). *Creating capabilities: The human development approach*. Belknap press of Harvard University press.

Omsorgsdepartementet, H. (2014, November 12). *NOU 2014: 12* [NOU]. Regjeringen.no; regjeringen.no. https://www.regjeringen.no/no/dokumenter/NOU-2014-12/id2076730/

Parfit, D. (1984). *Reasons and Persons*. OUP Oxford.

Parfit, D. (1991). *Equality or Priority*. University of Kansas, Department of Philosophy. https://kuscholarworks.ku.edu/handle/1808/12405

Pyne, J. M., Fortney, J. C., Tripathi, S., Feeny, D., Ubel, P., & Brazier, J. (2009). How Bad Is Depression? Preference Score Estimates from Depressed Patients and the General Population. *Health Services Research*, *44*(4), 1406–1423. https://doi.org/10.1111/j.1475-6773.2009.00974.x

Railton, P. (1986). Facts and Values. *Philosophical Topics*, *14*(2), 5–31.

Rand, L. Z., & Kesselheim, A. S. (2021). Controversy Over Using Quality-Adjusted Life-Years In Cost-Effectiveness Analyses: A Systematic Literature Review. *Health Affairs*, *40*(9), 1402–1410. https://doi.org/10.1377/hlthaff.2021.00343

Rawls, J. (1951). Outline of a Decision Procedure for Ethics. *The Philosophical Review*, *60*(2), 177–197. https://doi.org/10.2307/2181696

Rawls, J. (1999). *A Theory of Justice: Revised Edition*. The Belknap Press of Harvard University.

Rechnitzer, T., & Schmidt, M. W. (2022). Reflective Equilibrium is enough. Against the need for pre-selecting "considered judgments." *Ethics, Politics & Society*, *5*(2), 59–79.

Reeskens, T., Roosma, F., & Wanders, E. (2021). The perceived deservingness of COVID-19 healthcare in the Netherlands: A conjoint experiment on priority access to intensive care and vaccination. *BMC Public Health*, *21*(1), 447. https://doi.org/10.1186/s12889-021-10488-3

Rice, C. M. (2013). Defending the Objective List Theory of Well-Being. *Ratio*, *26*(2), 196–211. https://doi.org/10.1111/rati.12007

Robeyns, I., & Byskov, M. F. (2023). The Capability Approach. In E. N. Zalta & U. Nodelman (Eds.), *The Stanford Encyclopedia of Philosophy* (Summer 2023). Metaphysics Research Lab, Stanford University. https://plato.stanford.edu/archives/sum2023/entries/capability-approach/

Robson, M., Asaria, M., Cookson, R., Tsuchiya, A., & Ali, S. (2017). Eliciting the Level of Health Inequality Aversion in England. *Health Economics*, *26*(10), 1328–1334. https://doi.org/10.1002/hec.3430

Savulescu, J., Gyngell, C., & Kahane, G. (2021). Collective Reflective Equilibrium in Practice (CREP) and controversial novel technologies. *Bioethics*, *35*(7), 652–663. https://doi.org/10.1111/bioe.12869

Scanlon, T. (2000). *What We Owe to Each Other*. Harvard University Press.

Scanlon, T. M. (1975). Preference and Urgency. *The Journal of Philosophy*, *72*(19), 655–669. https://doi.org/10.2307/2024630

Scanlon, T. M. (1991). The moral basis of interpersonal comparisons. In J. Elster & J. E. Roemer (Eds.), *Interpersonal Comparisons of Well-Being*. Cambridge University Press.

Segall, S. (2016). Incas and Aliens: The Truth in Telic Egalitarianism. *Economics & Philosophy*, *32*(1), 1–19. https://doi.org/10.1017/S0266267115000309

Sen, A. (1980). Equality of What? *The Tanner Lecture on Human Values*, *I*, 197–220.

Sen, A. K. (2009). *The Idea of Justice*. Harvard University Press.

Sidgwick, H., & Rawls, J. (1981). *The Methods of Ethics, 7th Edition* (7th edition). Hackett Publishing Company.

Skotko, B. G., Levine, S. P., & Goldstein, R. (2011). Self-perceptions from people with Down syndrome. *American Journal of Medical Genetics Part A, 155*(10), 2360–2369. https://doi.org/10.1002/ajmg.a.34235

Solberg, C. T. (2019). Epicurean challenges to the disvalue of death. In C. T. Solberg & E. Gamlund (Eds.), *Saving People from the Harm of Death*. Oxford University Press.

Stenmarck, M. S., Jølstad, B., Baker, R., Whitehurst, D. GT., & Barra, M. (2023). A severely fragmented concept: Uncovering citizens' subjective accounts of severity of illness. *Social Science & Medicine*, *330*, 116046. https://doi.org/10.1016/j.socscimed.2023.116046

Sumner, L. W. (1999). *Welfare, Happiness, and Ethics*. Oxford University Press. https://doi.org/10.1093/acprof:oso/9780198238782.001.0001

Tännsjö, T. (1998). *Hedonistic Utilitarianism*. Edinburgh University Press.

Tännsjö, T. (2007). Narrow Hedonism. *Journal of Happiness Studies, 8*(1), 79–98. https://doi.org/10.1007/s10902-006-9005-6

Tännsjö, T. (2019). *Setting Health-Care Priorities: What Ethical Theories Tell Us*. Oxford University Press.

Temkin, L. S. (2014). *Rethinking the Good: Moral Ideals and the Nature of Practical Reasoning*. Oxford University Press.

Terlazzo, R. (2014). The perfectionism of Nussbaum's adaptive preferences. *Journal of Global Ethics*, *10*(2), 183–198. https://doi.org/10.1080/17449626.2014.931874

Terlazzo, R. (2015). Adaptive preferences: Merging political accounts and well-being accounts. *Canadian Journal of Philosophy*, *45*(2), 179–196. https://doi.org/10.1080/00455091.2015.1059252

Tersman, F. (2018). Recent work on reflective equilibrium and method in ethics. *Philosophy Compass*, *13*(6), e12493. https://doi.org/10.1111/phc3.12493

Thiel, G. J. M. W. van, & Delden, J. J. M. van. (2009). The justificatory power of moral experience. *Journal of Medical Ethics*, *35*(4), 234–237. https://doi.org/10.1136/jme.2008.026559

Ubel, P. A., Loewenstein, G., Schwarz, N., & Smith, D. (2005). Misimagining the unimaginable: The disability paradox and health care decision making. *Health Psychology*, *24*(4, Suppl), S57–S62. https://doi.org/10.1037/0278-6133.24.4.S57

Väyrynen, P. (2021). Thick Ethical Concepts. In E. N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy* (Spring 2021). Metaphysics Research Lab, Stanford University. https://plato.stanford.edu/archives/spr2021/entrieshick-ethical-concepts/

Velleman, J. D. (1991). Well-Being and Time. *Pacific Philosophical Quarterly*, *72*(1), 48–77. https://doi.org/10.1111/j.1468-0114.1991.tb00410.x

Versteegh, M. M., & Brouwer, W. B. F. (2016). Patient and general public preferences for health states: A call to reconsider current guidelines. *Social Science & Medicine*, *165*, 66–74. https://doi.org/10.1016/j.socscimed.2016.07.043

Voorhoeve, A. (2014). How Should We Aggregate Competing Claims? *Ethics*, *125*(1), 64–87. https://doi.org/10.1086/677022

Weinstein, M. C., Torrance, G., & McGuire, A. (2009). QALYs: The Basics. *Value in Health*, *12*, S5–S9. https://doi.org/10.1111/j.1524-4733.2009.00515.x

Williams, A. (1997). Intergenerational Equity: An Exploration of the 'Fair Innings' Argument. *Health Economics*, *6*(2), 117–132. https://doi.org/10.1002/(SICI)1099-1050(199703)6:2<117::AID-HEC256>3.0.CO;2-B

1

Utilitas (2022), **34**, 209–224
doi:10.1017/S0953820822000024

CAMBRIDGE
UNIVERSITY PRESS

ARTICLE

Age and Illness Severity: A Case of Irrelevant Utilities?

Borgar Jølstad[1]* and Niklas Juth[2]

[1]Akershus University Hospital, The Health Services Research Unit – HØKH, Lørenskog, Norway and Centre for Medical Ethics (CME), Institute of Health and Society, Faculty of Medicine, University of Oslo, Oslo, Norway and [2]Karolinska Institutet, Center for Healthcare Ethics, dep. of LIME, Stockholm, Sweden
*Corresponding author: E-mail: borgarjolstad@gmail.com

Abstract
Illness severity is a priority setting criterion in several countries. Age seems to matter when considering severity, but perhaps not small age differences. In the following article we consider Small Differences (SD): small differences in age are not relevant when considering differential illness severity. We show that SD cannot be accommodated within utilitarian, prioritarian or egalitarian theories. Attempting to accommodate SD by postulating a threshold model becomes exceedingly complex and self-defeating. The only way to accommodate SD seems to be to accept some form of relevance view, where some age differences are irrelevant. This view can accommodate SD, but at the expense of consistent priority orderings. Severity thus becomes unsuitable for systematic decision-making. We argue that SD should be dismissed and that we should accept a continuous relationship between severity of illness and age.

What is illness severity and why should we care?

Given the ubiquities of health problems and scarcity of health resources, priority setting is a necessity. One natural desire is for our resources to be used as effectively as possible; given a restricted health care budget, we want to "produce" as "much health" as possible. This strategy amounts to a maximizing ethics with a focus on the cost-benefit ratio of interventions. In several countries, this maximizing ethic is adjusted by a focus on the notions of need or *illness severity*[1] (Barra and others 2020; Franken and others 2015). These notions reflect the view that it does not only matter that we get as much health as possible but also who the recipients of these benefits are. For instance, in the Norwegian priority setting system, illness severity is a criterion alongside cost and benefit (Barra and others 2020). That is to say, Norway is willing to spend more on the same health benefits when they accrue to those who are severely ill. This use of an illness severity criterion leads to the question of what it means to be severely ill. How do we judge that an illness, or state of illness, is more severe than another? Plausible candidates include the level of pain, quality of life, life years lost, risk of

---

[1] Severity of illness and illness severity are used interchangeably throughout the article.

death, and the age of the patient, but there is as yet no consensus on what severity is or should be (Barra and others 2020). There seems to be a close relationship between the desire to prioritize those who are worse off and prioritizing those who are severely ill. Clearly those who have severe illnesses are worse off in some sense, but in precisely what sense is not obvious. Current operationalizations differ as to the role that age plays in illness severity. In the Norwegian system illness severity is operationalized as absolute shortfall: the severity of an illness is a function of how many good life years are lost. Age is thus indirectly a factor when assessing severity in the sense that younger patients typically have more life years to lose than older patients. More generally, it is probable that any severity-measure that operationalizes severity as life years lost will be sensitive to age. This can potentially be considered a form of age discrimination. In the Swedish priority setting system the Human Dignity Principle precludes considering chronological age in priority setting, and thus prevents age from being a factor in severity of illness (Barra and others 2020). The relationship between severity and age is thus not an uncontroversial one. The aim of this article is to explore some intuitions regarding the relationship between illness severity and age that seem to imply that age differences have to be relatively large to merit a differential judgement of illness severity. We will argue that this implication can only be accommodated at the expense of a consistent priority ordering.

But first, why care about illness severity? Why not simply get as much health as possible? There are several reasons for caring about illness severity. First, unlimited aggregation leads to unintuitive consequences. An unlimited aggregation view entails prioritizing minor benefits for large numbers of well-off individuals over large benefits for small numbers of worse off individuals, which can seem morally counterintuitive (Gustavsson and Juth 2019). A second reason is that the worse off are morally more important in some way. Such a concern for the worse off is asserted by both prioritarian, egalitarian, and, arguably also, sufficientarian theories of distribution (Hirose 2014). A third reason, somewhat more speculative, is that a public health care system is viewed by many as a form of safety net. Utilizing resources on relatively well-off individuals that could have been spent on someone who has a desperate need may feel wrong in this context (Gustavsson and Juth 2019). All of these reasons highlight the affinity between a focus on severity and a desire to prioritize the worse off.

In this article, two assumptions are made about severity of illness: that severity of illness is (at least partly) a function of how *bad* the illness is for the person,[2] and that severe illness constitutes a reason to prioritize a patient. It is thus assumed that severity can be viewed in the following two ways: (Severity 1) as a measure of health-related worse off-ness, and (Severity 2) as a marker for which patients should be given priority. Severity seems to be a thick concept with a descriptive dimension (Severity 1) and a normative dimension (Severity 2). We also assume a deprivation account of the badness of death, where how bad it is to die is a function of how much you lose because of death (Nagel 1970; Solberg and Gamlund 2016). It seems that if such a deprivation account on the badness of death is not assumed the notion of worse off-ness is difficult to make sense of in the context of health priorities. This is also how health loss is assessed in current health priority settings. Leaning on the deprivation account ensures that what we are aiming at is not maximizing the number

---

[2] In focusing on the badness for the patient, we are intentionally ignoring other factors that some might consider relevant to illness severity, such as whether the patient has children, is socially important, saves lives for a living, etc.

of good life years in the population (this could entail replacing sick individuals with new individuals) but rather maximizing the good life years of already existing members of the population. A further point is that we are here attributing severity to patients, or cases of illness, rather than to illnesses per se.[3] A given illness can thus, for many different reasons, be more or less severe *for a person*. Lastly, the subject of this article is severity as a priority setting criterion with a moderating role in a largely maximizing priority setting system. Severity is thus a part of a larger decision-making scheme that also includes the cost and benefit of the intervention.

Severity and age: four cases

A potential and probably commonly held intuition about the relationship between age and severity of illness is that terminal illness is, *ceteris paribus*, less severe when faced by the very old and more severe when faced by the very young. In this regard, consider the following two cases:

Case 1
- Patient A is a 20-year-old who has an expected survival of one year due to an illness.
- Patient B is an 80-year-old who has an expected survival of one year due to an illness.

Is one of these cases more severe, and if so, which one? Next, consider Case 2 below:

Case 2
- Patient C is a 40-year-old who has an expected survival of one year due to an illness.
- Patient D is a 45-year-old who has an expected survival of one year due to an illness.

Is one of these cases more severe, and if so, which one?

First, it might be helpful to standardize the cases further by positing that the patients enjoy the same level of welfare, would prefer to live on, are childless and are equally socially valuable. All of this is meant to ensure that what we are trying to tease out is how bad the illness is from the perspective of the individual patient and not for (for example) a family, or society at large. By assuming that all the patients have the same level of welfare we also mean to focus solely on the length of life, and not on the quality of life. Further, note that all the patients have the same amount of life ahead of them (one year). Hence, the difference between the patients is to what extent they are *worse off* in two ways: (1) they *stand to lose* more or less time compared to a hypothetical “normal” lifespan, and (2) they *have had* more or less time to live.[4] We thus assume that the 40-year-old would be similar to the 45-year-old in all relevant aspects if she were 10 years older. Our intuitive judgement is that patient A is more severely ill than patient B, but that patients C and D are equally severely ill. Age thus

---

[3] We would like to thank an anonymous reviewer for pointing this out to us.

[4] Note that this does not follow from a deprivation account but seems to be an intuitive reason for the badness of early death. Nothing in the article hinges on the reader accepting this as a reason for the badness of dying early.

seems to make a difference for illness severity in case 1 but not in case 2.[5] Can this difference in intuitive judgement be grounded in moral theory? In a series of interviews on severity that the corresponding author has recently taken part in, several participants expressed views on severity and age of this kind. More specifically, they expressed confidence in assigning differing illness severity when differences in age were large but considered cases equally severe, or were unwilling to make a judgement, when age differences were smaller.[6] We believe that these intuitions are similar to the intuitions that Francis Kamm claims to have about differences in years saved and whom to aid. She argues for the relative merits of equal chances, a weighted lottery and choosing on the basis of differences in outcome: for a 19-year difference in outcome she leans towards choosing on the basis on this difference, whereas a 2-year difference merits equal chances (Kamm 1993, pp. 287–88). Some differences in years saved are thus enough to merit differential treatment, whereas others are not. An interesting real world example of something similar to this is found in the triage recommendations at the critical care unit at the University of Pittsburgh (University of Pittsburgh 2020).[7] These recommendations were made while expecting cases of Covid-19 to exceed the resources at the hospital. Patients are here categorized into age groups based on life stages. Whereas being in different life stages can result in differential treatment, a lottery is recommended for other cases. Clearly, some age differences are considered relevant while others are not.[8] We will therefore assume that this differential judgement is worthy of consideration and investigate what kind of ethical theories can account for the divergence. First, we will consider common ethical frameworks for priority setting in health care (utilitarianism, egalitarianism, prioritarianism, and sufficientarianism). We will then argue that the differential judgement, and its implications, cannot be accommodated within any strictly consequentialist framework but must rather be seen from deontological perspectives. Lastly, we will then discuss problems with these perspectives.

Before we can consider the merits of the different relevant theories, it will be helpful to scrutinize these intuitive judgements. The intuitive judgement that patient A is more severely ill than patient B seems to us to be the strongest. This intuition can be explained in many ways. Illness in old age seems more *natural* (whatever this means), an older person has already had a long life (we are assuming a reasonably happy one) and the person typically does not have many years of life left regardless of illness (at least not compared to what 20-year-olds typically have). Illness, and especially terminal illness, in the young strikes us as tragic. Patient A has not had a chance to live a life and we consider death at this young age as something *unnatural* (whatever this means). Compare this with case 2. Patients C and D are both more like patient A in the mentioned respects. They are both in the prime of their lives, have a lot to look forward to and death at this age strikes us as tragic (but maybe not as tragic as death at 20).

---

[5] Note that in current priority setting contexts, a 40-year-old would be prioritized over a 45-year-old if the 40-year-old could be helped to a larger extent, for example in an emergency situation, where we have to choose between saving one or the other. In this article, we assume that whatever health gain can be provided is the same for all patients, thus assuming that cost and benefit are equal.

[6] This research, part of the Severity and Priority Setting in Health Care (SEVPRI)-project, is still a work in progress.

[7] We would like to thank an anonymous reviewer for directing our attention to this report.

[8] This resembles a sufficientarian system with multiple thresholds, something that will be discussed later in the text.

All this leads us to ask whether there is something about the old that leads us to consider an illness less severe. Maybe there is a cutoff – say, at 80 – where death is “natural” or at least not as bad? This would be in line with the “fair innings” argument where “everyone is entitled to some ‘normal’ span of health … The implication is that anyone failing to achieve this has in some sense been cheated, whilst anyone getting more than this is ‘living on borrowed time’” (Williams 1997: 119). The “fair innings” is a threshold after which additional years of life have less moral importance. The fair innings argument can account for our intuition by positing that patient B has had a full life (or at least a chance at one), whereas none of the other patients have. Patients A, C and D are worse off than patient B (severity 1) in a sense that makes them more worthy of moral concern by being prioritized (severity 2). On a strict “fair innings”-threshold interpretation patient B thus stands out as less severely ill than the other three patients who are equally severely ill because neither of them has reached the threshold.[9]

But now compare cases 1 and 2 with the following:

Case 3

- Patient E is a 20-year-old who has an expected survival of one year due to an illness.
- Patient F is a 50-year-old who has an expected survival of one year due to an illness.

Is one of these cases more severe, and if so which one?

Case 4

- Patient G is a 77-year-old who has an expected survival of one year due to an illness.
- Patient H is an 80-year-old who has an expected survival of one year due to an illness.

Is one of these cases more severe, and if so which one?

Our intuitive judgement here is that patient E is more severely ill than patient F, but that patients G and H are equally severely ill. This cannot be accounted for by the fair innings idea strictly conceived. If one assumes that we are all entitled to 80 years of life, then patient H is the only one whose illness is less severe. Our intuitive judgement is that the difference between patient G and H is not *large enough* to warrant a differential judgement of illness severity and that the difference between patient E and F is clearly significant. The intuitive judgements seem to lead to the following principle:

> Small differences (SD): small differences in age are not relevant when considering differential illness severity.

SD implies that there is no continuous relationship between age and severity. The severity of an instance of illness, rather that reflecting an underlying ranking of states of affairs (or a consistent ranking of the severity of illness based strictly on how the illness in question affects age of dying), seems to be relative to other instances of illness.

---

[9] If this seems strange, it is perhaps because if the 45-year-old has been cheated in comparison to the fair innings, then it seems that the 20-year-old has been cheated to a larger extent. John Harris, the originator of the fair innings argument, discusses this in *The Value of Life* (Harris 1985).

Terminal illness does not seem more severe for a 20-year-old than for a 25-year-old, the illness of a 25-year-old does not seem to be more severe than the illness of a 30-year-old etc., but the illness of a 20-year-old seems to be more severe than the illness of, for instance, a 50-year-old. We thus have a case where $20 \ngtr 25$, $25 \ngtr 30$, $30 \ngtr 35$, $35 \ngtr 40$, $40 \ngtr 45$ and $45 \ngtr 50$, but $20 > 50$. If these intuitive judgements are taken at face value, it seems that the "not more severe than"-relation is not a transitive one. The same will follow for the equally severe as-relation.

One problem when trying to make sense of SD is that it flies in the face of the worse off-aspect of severity. Assuming you are not a "hard epicurean" it is difficult to argue that death at 40 does not leave you worse off than death at 45: you have lost an additional 5 years (assuming that you would otherwise live for the same amount of time), and you have had 5 fewer years. If we are to keep both SD and admit that death at 40 is worse than death at 45, then severity cannot be a function of worse offness alone. This difficulty will be explored later on, but for now the problem seems to be that severity has a priority aspect and unless one is a consequentialist, worse off-ness (severity 1) and moral significance (severity 2) do not completely overlap. This claim will be explored further in what follows.

Utilitarianism and prioritarianism

From the perspective of utilitarianism, the morally right thing to do is to maximize the net sum of goodness over badness. Utilitarianism is both *consequentialist* in that what matters is the resultant states of affairs and *aggregative* in the sense that all utility is counted, no matter the size of the utility. We will here ignore the differences between different forms of utilitarianism and simply assume that there is some form of health-related good that can be maximized, most naturally good life years. Whether the good accrues to someone who is worse off or not does not matter in itself according to utilitarianism. Axiologically, utilitarianism is committed to a neutral weighting between individuals, and typically assumes a continuous ranking of states of affairs. All of this ensures that from a utilitarian point of view the state of affairs of dying as a 20-year-old (rather than as a 21-year-old) is, *ceteris paribus*, neither better nor worse than dying as an 80-year-old (rather than as an 81-year-old), but the same can be said for the difference between 40 and 45 (or even 40 and 41). If we were able to avoid one of these states of affairs we should, all else being equal, be indifferent between the patients. Where all the patients have one year to live, there is no reason to prioritize the worst-off, since they cannot be helped to a higher degree. A year of life is, all else equal, equally worthy of consideration regardless of whom it accrues to. On the most natural reading of utilitarianism, given the neutrality to worse off-ness, the notion of severity thus seems superfluous. Or, at least, it has no independent moral weight – at most it is a way to talk about levels of value. What matters is the utility that can be produced. Any amount of badness and good would simply be factored into the calculation, and whether we call it severe or not is irrelevant.

If one is attracted to the idea of maximizing *and* prioritizing the worse off, then *prioritarianism* seems to be a natural next step. From a prioritarian perspective we should give more, but not absolute, priority to helping those who are worse off (Hirose 2014; Parfit 1991). Prioritarianism builds on the idea that the claims of the worse off are morally weightier to some degree. Well-being is claimed to have marginally decreasing moral importance. This is typically assumed to take the form of a concave value function of moral goodness, such as this (see Figure 1):

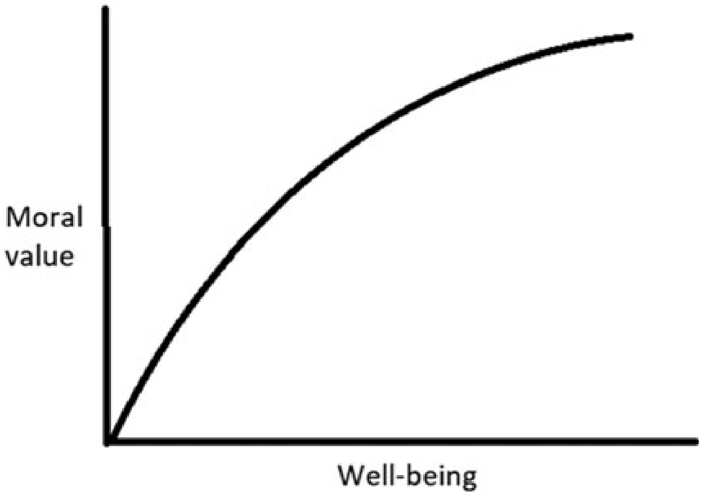

Figure 1. Example of a prioritarian value function.

As is evident from the figure there is more value associated with increasing the well-being of someone lower on the well-being scale. It is also clear that the relationship between well-being and moral value is continuous. Due to the continuity of worse off-ness and of the function of worse off-ness and moral importance we end up with a system where those who are worse off are given priority proportional to how badly off they are. But it is obvious that this system cannot accommodate SD. If patient C is worse off than patient D (and we are assuming this), then it is obvious that we should prioritize patient C. Patient C is thus more severely ill than patient D on both the worse off- and priority-aspects of severity (severity 1 and 2). Prioritarianism, like utilitarianism, assumes a continuous relationship between worse off-ness and moral importance, and therefore cannot accommodate SD.

Egalitarianism

The four most influential egalitarian theories of distribution are *telic egalitarianism*, *luck egalitarianism*, *sufficientarianism*, and *Rawls maximin or leximin principles* (Hirose 2014). For the purposes of this article, we will disregard luck egalitarianism, because illness would in most cases be considered a form of brute luck.[10] We will first consider telic egalitarianism and maximin/leximin.

Telic egalitarianism is a family of theories that claim that inequality contributes to the badness of an outcome over and beyond the effects that this inequality has on specific individuals. According to Parfit (1991) telic egalitarians believe that equality has intrinsic value or, rather, that the more inequality, the worse the outcome (all else being equal). In our examples, this means that in addition to the badness contributed by the illness of the patients there is a separate badness contributed by the fact that the illness is unequally distributed in the population. The illness of patient A is thus a worse outcome than the illness of patient B both because patient A is worse off and because the illness of patient A contributes more to the separate badness of inequality. The same can be said for the illness of patient C and D. Patient C is both worse off than patient D and contributes more to the badness of inequality. The difference in moral importance caused by differential worse off-ness is thus even greater according to telic

---

[10] Some instances of illness, such as illnesses caused by smoking or obesity, can arguably be considered instances of option luck. We are disregarding this complicating factor for the purposes of this article.

egalitarianism, and there is no reason to disregard the difference in worse off-ness between patient C and *D*. Telic egalitarianism is thus also unable to accommodate SD.

Rawls' theory on maximin or leximin dictates giving either absolute (maximin) or lexical (leximin) priority to the worse off (Rawls 1999).[11] On a maximin reading of case 1, patient A is not just more morally important than patient B, she is the only morally important person in this case. On a leximin reading, patient A is morally more important as long as something can be done for her, and only when this is not possible does patient B acquire moral importance. But the same can be said for the remaining three cases. There is, following the logic of Rawls theory, no reason not to give weight to the differences in worse off-ness, and every reason to do so. Rawls' theory thus seems, to an even larger extent than previously discussed theories, to be unable to accommodate SD.

The reason that these theories cannot accommodate SD seems to be that these theories are consequentialist;[12] moral importance (priority) is a function of consequences (either worse off-ness or worse off-ness and the badness of inequality). The two views on severity (worse off-ness and being marked for priority) are either collapsed into one, or on the utilitarian reading rendered superfluous. The ones who should be prioritized (severity 2) are the ones who are worse off to the largest degree (severity 1) (on the utilitarian reading worse off-ness means a greater opportunity to do good) and there is no room for SD.

Sufficientarianism

Sufficientarianism is a family of theories that claim that what is morally important is that everyone has enough, or more accurately that we should prioritize "those whose well-being is below a certain threshold" (Hirose 2014: 112). The already discussed fair innings argument can thus be viewed as an essentially sufficientarian argument, albeit one with a high threshold. This theory might be able to make sense of the intuitive judgements of the cases if we interpret them as solely pointing to the lower severity of the very old and higher severity of the young. To be able to account for all of the cases we would have to assume a sufficientarian theory with multiple thresholds. The result might be something akin to a curve that looks like this (see Figure 2):

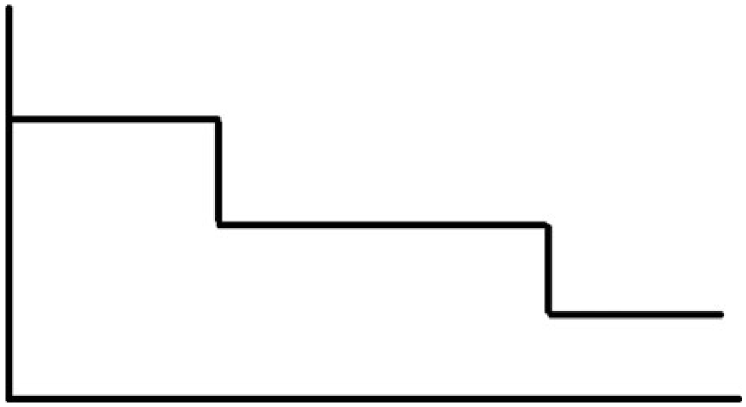

Figure 2. Example of a sufficientarian two-threshold value function.

---

[11] We are here disregarding the complications around the focus on the worst-off groups and worst-off individuals.

[12] We are here using the term consequentialist in the Rawlsian (Rawls 1999) sense to include all theories that view the right as following from the good, rather than in Kagan´s (Kagan 1992) sense, where the difference between consequentialist and non-consequentialist theories consists in whether they limit moral

where age is on the *X*-axis and degree of severity is on the *Y*-axis. This would imply considering the illnesses of the young as very severe,[13] the illnesses of the middle-aged as medium severe and the illnesses of the old as less severe.[14] This could account for some of our intuitions but would lead to a new set of problems. Say that we define young as under-30-year-olds. Then what about 31-year-olds? Is the difference between a 30-year-old and a 31-year-old large enough to merit a difference in kind? And, similarly, is the difference between a 79-year-old and an 80-year-old relevant when considering the severity of an illness? The thresholds seem arbitrary and the value difference between people on different sides of the threshold seems too sharp. This problem would also appear if we attempted to avoid the discontinuity-difficulties by claiming that the relationship between severity and age takes the form of categories (for instance 10–20, 21–30, 31–40 years old, etc.). We could perhaps avoid the problem of sharpness by making the thresholds sloped. The curve would then look something like this (see Figure 3):

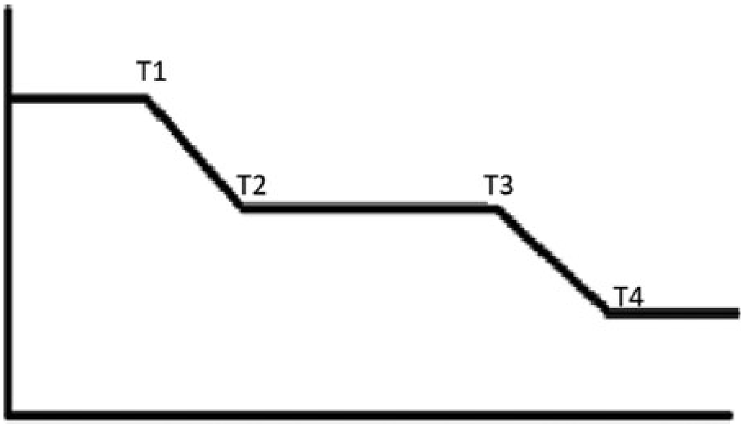

Figure 3. Two-threshold value function with sloped thresholds.

Say the line from T1 to T2 represents ages 30–40 and the line from T3 to T4 represents ages 70–80. This would remove the problem of sharpness but would reintroduce a continuous relationship between age and severity on parts of the line (the sloped lines). One could then claim that there is, all else being equal, no difference in severity between a 40- and 60-year-old. But one would be forced to accept that there is a difference in severity between a 30-year-old and a 30-year-and-a-day-old. What could be the reason for accepting the relevance of such small differences on certain parts of the line while ignoring large differences on other parts of the line? Reintroducing continuity on parts of the line seems more arbitrary than accepting it on all parts of the line.

In addition to the problems associated with arbitrariness and sharpness it also seems like a multiple threshold system will become very complex if we are to accept SD. If we

---

aggregation. Note that within Rawls' theory it is the maximin/leximin decision procedure that is clearly consequentialist in this sense. Rawls' system as a whole, especially considering his principle on basic liberties, is plausibly a hybrid view.

[13] We are ignoring the difficult questions regarding the disvalue of death for the very young, such as in Jeff McMahan's (McMahan 2019) Time-relative Interest Account. We are simply stipulating that the graphs start at 20 years.

[14] This would essentially be a simplified version of a life stages view. We believe that the arguments discussed here would apply to a more complicated life stages view as well.

make a series of cases such as 30 vs. 60, 40 vs. 70, etc. we would have to include thresholds between 30 and 60, 40 and 70, etc. It seems like no number of thresholds would be enough to accommodate our intuition. Introducing enough thresholds would approximate a continuous line, and thus fail to accommodate our intuition.

These two contrasting difficulties, between on the one hand having to find more or less arbitrary cut-offs and accepting that a difference of one year (or even a day) is relevant for considerations of illness-severity on the other, seem to be impossible to avoid within a consequentialist system with an underlying complete ranking of states of affairs. This leaves us with the options of either dismissing SD and our intuitive judgements or exploring (at least partly) deontological views. We first turn to deontological views.[15]

Severity as relevant worse off-ness

In the current Norwegian priority setting system funds are allocated on the premise of optimizing the numbers of Quality Adjusted Life Years (QALYs) per monetary unit, with extra weight on QALYs accrued to those with the most severe conditions.[16] Even though they are weighed more heavily, the QALYs accrued by curing or treating, for instance, lethal cancer are weighed against the QALYs gained by curing or treating less severe conditions. If one were to choose between an intervention that saved the life of one patient from lethal cancer and an intervention that cured the once-a-month moderate headache of a thousand individuals the answer is not straightforward and would involve calculating the QALYs gained (weighed by severity) for each intervention. We believe that many would object to this calculation on the grounds that once-a-month moderate headaches are irrelevant when compared to the cancer patient's prospect of shortening her life, no matter how many people suffer from these headaches.[17] In *Morality, Mortality* (1993) Frances Kamm introduces the notion of *irrelevant utilities*. An irrelevant utility is a utility that should not be considered in the process of deciding between competing claims for a resource. The notion of irrelevant utilities is dynamic. Whether a utility is irrelevant is dependent on what it is measured up against: in the context of who should live, a broken arm might be irrelevant. A broken arm, however, is not an irrelevant utility when measured against a broken leg.

The notion of irrelevant utilities seems to be able to buttress the intuition expressed in SD. Consider the option (or impossible choice) of either saving the life of one person, or the life of another plus, say, curing the intermittent headache of a third person. Is it right, in the context of choosing who lives, to prioritize based on the added utility of curing the intermittent headache? This issue goes to the heart of the current priority-setting framework and might also aid us in our quest to make sense of SD. If we assume that severity of illness has normative implications, then the question of who is more severely ill may not reduce to a question about who is worse off (severity 1), but also a question of who should be given higher priority when distributing scarce resources (severity 2). When claiming that the terminal illness of a 20-year-old is more severe than the terminal illness of an 80-year-old we are not *only* saying that the 20-year-old is worse off, but also that this worse off-ness is significant when deciding whom to aid. When we are claiming that the terminal illness of a 40-year-old is equally severe as for a

---

[15] Deontological in the sense of limits or constraints on moral aggregation. See Kagan (1992).

[16] A QALY is a life year adjusted for the health-related quality of life in that year. A QALY number is thus a function of the number of years and the health-related quality of life in those years.

[17] Varieties of this view can be found in Kamm (1993), Scanlon (2000) and Temkin (2014).

45-year-old, we might be claiming that even though the 40-year-old is worse off than the 45-year-old (severity 1) this difference in worse off-ness is not significant when deciding whom to aid (severity 2). The extra years are in this context something akin to an irrelevant utility, or rather a form of *irrelevant worse off-ness*. We can thus conceptualize *severity as relevant worse off-ness (SRW)* as illness severe enough to justify different priority setting. An advantage to conceptualizing the relationship between severity and age in this way is that we can assume a full ranking worse off-ness while keeping our intuitions that seem to display intransitivity. What we mean by a full ranking of worse off-ness is that it should be possible to rank all people from worst to best off in such a way as to make the worse off-relation a transitive relation (probably on a linear scale). On this scale the soon-to-be-dying 40-year-old is worse off than the soon-to-be-dying 45-year-old and would also be worse off than a soon-to-be-dying 41-year-old. If we consider some additional years as irrelevant utilities the question of whether the 40- or 45-year-old is more severely ill is not a question reduced to who is worse off, but rather whether one is worse off enough to make unequal treatment legitimate. The question of who is more severely ill will then be a function of whether there is a *morally relevant* difference in worse off-ness *in a priority setting decision*.

SRW thus has appeals but is not without problems of its own. Perhaps most notably, SRW works in the context of pairwise comparison, but not when confronted by a context where multiple decisions are be made over time. To see why this is so imagine that you are making decisions on which treatments to fund in a publicly funded health care system. In addition to questions of cost and effectiveness, severity of illness is a priority setting criterion. Treatment 1 can give 30-year-old patients with a terminal illness an additional year of life. The cost is high but given the severity of the illness (surely terminal illness in 30-year-old patients is severe) you decide to fund the treatment. Now imagine that treatment 2 is like treatment 1 but aimed at 35-year-old patients. Illness at 35 is equally severe as illness at 30 ($30 = 35$), so (all else being equal) you should fund the treatment. Then treatment 3 shows up, aimed at 40-year-old patients. Our intuition is that terminal illness at 40 is equally severe as illness at 35 ($35 = 40$). Then treatment 4, and all the way up to treatment *n* aimed at 90-year-old patients, shows up. We thus end up assigning the same severity to 30- and 90-year-olds, even though we wish to claim that the illness of the 30-year-old is more severe than that of the 90-year-old; we want to say $30 > 90$ but end up saying not $30 \ngtr 90$ or alternatively $30 = 90$. Note that this problem is not dependent on a specific degree of overlap. If one thinks that 5 years is more than enough to consider one case of illness more severe than another, the same problem can be shown to result from any difference (one year for instance) given enough treatment options (treatment for 31-year-old patient, 32, 33 … 80, etc.). SRW is thus an intransitive relation.[18]

Interestingly, the “collapse” of the severity relation would not come up if we only considered, for example, a 20-, 50- and 80-year-old patient.[19] Let us stipulate here that the age difference has to be at least 20 years for a differential judgement of illness severity to be merited. We would then have no problem claiming that the 20-year-old is more severely ill than the 50-year-old, who is more severely ill than the 80-year-old. But once we include more alternatives, for instance a 35-year-old and a 65-year-old, the

---

[18] Strictly speaking the “more severe than”-relation and the “equally severe as”-relation within SRW will be intransitive. “More severe than” will be a transitive relation.

[19] How large the differences have to be depends on the intuitive judgements of the reader.

severity ordering collapses[20] and we are no longer able to prioritize the 20-year-old over the 80-year-old. This is clearly a violation of the independence of irrelevant utilities (my preference of apples over oranges should not be dependent on whether there are pears present). A few clarifications are in order. First, the problems just mentioned are dependent on the severity-relation being a total ordering. This means that, given any two patients A and B with any possible level of illness related worse off-ness, the illness of patient A is either more severe than the illness of patient B (A > B), less severe than the illness of patient B (A < B) or (either precisely or imprecisely) equally severe as the illness of patient B (A = B). This essentially means that all patients can be compared in terms of severity of illness.[21] This might not be a completely uncontroversial assumption, but it is difficult to see how a severity-relation that fails to compare all relevant alternatives can function as a criterion for systematic priority setting. Second, intransitivity and sensitivity to irrelevant alternatives are only issues for the normative aspect of severity (severity 2). We are thus not arguing that worse off-ness is an intransitive relation.

The need for a complete ranking for priority setting purposes also rules out the concept of parity. According to Ruth Chang (Chang 2016) two items can be comparable in the sense of being *on a par*; they are comparable without being either more, less or equally X. Two instances of illness could thus be related such that neither is more severe than the other, without this implying that they are equally severe, yet still be comparable. The important aspect of parity for our purposes is that parity is an intransitive relation. A being on a par with B, which is on a par with C, does not imply that A is on a par with C. This intransitivity will give rise to the same problems regarding multiple decisions as SRW.[22]

Some might be tempted to save SD by assuming that severity is what Temkin has called an essentially comparative concept. The essentially comparative view "implies that an outcome may have a certain value relative to one alternative-set, and another value relative to another alternative-set" (Carlson 2013). In our context this would entail that when we are comparing a 20- and 80-year-old this is one alternative set, with its own relevant factors. When comparing a 20-, 30- … 80-year-old, this is another alternative set, with its own factors. However, in the context of priority setting we need a consistent frame of reference, a "privileged alterative set" (Carlson 2013). Systematic use of the more severe-than relation seems difficult, and probably unjust, if there is no common reference frame for priority setting.

What might a priority setting system, sensitive to severity and age, look like if one were to accept some form of relevance view while still maintaining transitivity and insensitivity to irrelevant alternatives for the purposes of multiple decisions? Here is a suggestion. Let A, B, C, D and E be five patients who are 20, 35, 50, 65 and 80

---

[20] In the sense that all the illnesses are equally severe, and neither is more severe than the others.

[21] Depending on your definition of commensurability, full commensurability might also require precise cardinal comparability as well. This is probably also necessary if severity is to function as a systematic priority setting criterion on a large scale. Say that we claim that the illness of a 20-year-old is, *ceteris paribus*, more severe than the illness of a 40-year-old. We can do this without specifying how much more severe, but if severity is to function as a systematic criterion the difference has to be cardinally measurable and commensurable with other criteria of our health care setting. These difficulties will not be addressed in the current article.

[22] Chang (2016) suggests that there are pragmatic pressures to accept commensurability where there might be no real commensurability. Our desire to make health care priority setting systematic might function as such a pragmatic pressure.

years old respectively, all with an expected survival rate of one year due to illness. Let us once again stipulate that 20 years is enough to warrant a differential judgement of illness severity. If we had to choose between giving an extra year of life to either A or B on the grounds of severity, we would not be permitted to choose A because she is younger, but we would be permitted, or perhaps even required, to save A over C, D or E. So as far as illness severity goes A = B & A > C, D and E. Similarly for B: B = A and C & B > D and E. For C: C < A, C = B and D & C > E. So, C should be given priority over E, be treated equally to B and D, and be prioritized less than A. This would let us make single decisions, but how would this play out as a priority setting criteria with multiple decisions? Let cohort stand for each group of patients who are the same age. Each cohort X would stand in one of three relations to all other cohorts: X = Y, X > Y or X < Y. Let us call the X = Y relations equality-relations, the X > Y relations more severe than-relations and the X < Y less severe than-relations. For any cohort X, the number of these relations will be a function of the cohort's age. A young cohort will have few (if any) less severe than-relations, comparatively few equality-relations, and many more severe than-relations. A middle-aged cohort will have more less severe than-relations, the highest number of equality relations and fewer more severe than-relations. An old cohort will have many less severe than-relations, comparatively few equality-relations, and few (if any) more severe than-relations.

Figure 4. Ratio of more severe than- to less severe than-relations as a function of age.

Figure 4 shows how the ratio of more severe than-relations to less severe than-relations changes as a function of age. In other words, it maps how severe terminal illness is in a cohort relative to all other cohorts. The graph would look different if we included equality relations because the middle-aged have comparatively more equality-relations than either the young or the old. Prioritizing based on these relations would mean giving higher priority to 40-year-olds than to 50-year-olds because the 40-year-olds should be given priority relative to a larger part of the population than the 50-year-olds. This is so even though if this were a singular decision, we would not choose between them based on age. This would paradoxically reintroduce the continuous function of age and severity on a priority setting level. We would in effect be mindful of the fact that priority setting is a continuous process and always bear in mind not just singular decisions, but the larger framework. It seems that this does not preserve the spirit of the relevance view. Or, to put it somewhat differently, it "saves" the relevance view by making it practically irrelevant.

It thus seems that the relevance view and SRW can accommodate SD, but that SRW is not suitable for priority setting on a systemic level. The possibility of devising a cohort-based system remains but seems to make both our intuitive judgements and SD irrelevant. The last possibility to be discussed is whether we should dismiss our intuitive judgements and SD.

Should we dismiss SD?

So far, we have explored two ways of making sense of SD, neither of which seems satisfactory. A multiple threshold model will either be plagued by arbitrariness or approximate a continuous line. SWR is troubled by intransitivity and sensitivity to irrelevant alternatives and is as such not suited for systematic priority setting. A third alternative is to dismiss SD and accept a linear relationship between age and severity.

What reasons do we have for dismissing SD? First, there is the possibility that the intuitive judgements regarding our four cases are not about relative age differences per se, but rather stem from other attributes that we have failed to adequately control for. We might simply not be comfortable making relative judgements of worse off-ness unless the differences are obviously large. If faced with a 40- and 45-year-old we might not be sure whether the 45-year-old has really either had more or stands to lose more in terms of good life years (for instance, the 45-year-old may have had more life-years with really poor quality of life compared to the 40-year-old). When faced with a 20- and 80-year-old, the difference is, or at least seems, obvious. Thus, even though we have assumed all else equal to a maximum extent in the cases, and thus made it clear who is worse off, our intuitions might not be sensitive to this kind of "fine tuning". If this is true, then our intuitions might be a result of uncertainty regarding worse off-ness rather than a result of beliefs regarding relevance. A similar argument can be made regarding life stages. Even though we have stipulated that the patients are similar in all relevant aspects other than age we might not be able to exclude intuitions based on ideas of the life stages of the different age groups. We might for example be unable to make a differential judgement of severity between the 40- and 45-year-old because they are, at least in our minds, in the same stage of life. Excluding this by stipulating that the patients are alike in all relevant aspects makes theoretical sense but might be difficult in practice. These are essentially debunking arguments; our intuitive judgements of X might be better explained without postulating Y (Tersman 2008). Debunking is a contested field, but the specific intuitions discussed in this article seem especially suitable for a debunking explanation. This is so, again, because there are obvious confounders, such as uncertainty regarding value difference between "close enough" alternatives.

A second reason for dismissing SD is that there seems to be something wrong with the very notion of irrelevance when applied to life years. Life years (and preferably good life years) are, in a very basic sense, all we have. A year can mean watching your child learn to walk, mastering a skill, saying a proper goodbye to your loved ones, reconciliation and so much more. Even a day can be significant in the scheme of an entire life. From this perspective, claiming that any significant span of time is insignificant in the scheme of justice may seem perverse.

A third reason, perhaps the weightiest, is that SD forces us to accept either some form of exceedingly complex threshold model or a conceptualization of severity that is intransitive and sensitive to irrelevant alternatives. In light of these consequences, accepting a continuous relationship between severity and age might not seem so unappealing.

Conclusion

This exploration of various ways of accounting for our intuitive judgements seems to suggest that we are faced with three options, none of them completely satisfactory. The first option is to accept that the judgements of severity are intransitive and sensitive to irrelevant alternatives. This would involve either a great deal of difficulty when applying the concept in priority setting work or using the concept in a way that makes the relevance view irrelevant for priority setting (as in the cohort system described earlier). Given that transitivity and insensitivity to irrelevant alternatives are typically considered prerequisites for rational decision-making, it also seems like a theoretically problematic option. The second option is to develop a complex threshold model. If the core of the intuition is SD, then the number of thresholds would be high, and theoretically limitless. The third option is to disregard SD and assume a continuous relationship between severity and age. This would force us to accept that the illness of a 40-year-old is, all else equal, more severe than the illness of a 41-year-old. Considering that time spent alive matters to us, this seems like a better option than the alternatives.[23]

Declaration of competing interests. We have no competing interests to declare.

References

Barra, Mathias, Mari Broqvist, Erik Gustavsson, Martin Henriksson, Niklas Juth, and others. 2020. Severity as a Priority Setting Criterion: Setting a Challenging Research Agenda. *Health Care Analysis*, **28**.1: 25–44 <https://doi.org/10.1007/s10728-019-00371-z>.

Carlson, Erik. 2013. Intransitivity. In *International Encyclopedia of Ethics* (American Cancer Society) <https://doi.org/10.1002/9781444367072.wbiee410>.

Chang, Ruth. 2016. Parity, Imprecise Comparability and the Repugnant Conclusion. *Theoria*, **82**.2: 182–214 <https://doi.org/10.1111/theo.12096>.

Franken, Margreet, Elly Stolk, Tessa Scharringhausen, Anthonius de Boer, and Marc Koopmanschap. 2015. A Comparative Study of the Role of Disease Severity in Drug Reimbursement Decision Making in Four European Countries. *Health Policy*, **119**.2: 195–202 <https://doi.org/10.1016/j.healthpol.2014.10.007>.

Gustavsson, Erik, and Niklas Juth. 2019. Principles of Need and the Aggregation Thesis. *Health Care Analysis*, **27**.2: 77–92 <https://doi.org/10.1007/s10728-017-0346-6>.

Harris, John. 1985. *The Value of Life: An Introduction to Medical Ethics* (London: Routledge) <https://doi.org/10.4324/9780203005606>

Hirose, Iwao. 2014. *Egalitarianism* (London: Routledge) <https://doi.org/10.4324/9781315772004>.

Kagan, Shelly. 1992. The Structure of Normative Ethics. *Philosophical Perspectives*, **6**: 223–42 <https://doi.org/10.2307/2214246>.

Kamm, Frances M. 1993. *Morality, Mortality. Volume I: Death and Whom to Save From It*, BRL:KIE/42748 <https://repository.library.georgetown.edu/handle/10822/743570> [accessed 7 June 2021].

---

[23] We would like to thank two anonymous reviewers for helpful comments. The thesis of this article was presented at the Philosophy of Priority Setting seminar on 1 June 2021. We thank the participants for valuable discussion and comments. We would also like to thank Mathias Barra and Carl Tollef Solberg for helpful comments on a previous version of this article. The work of the authors is funded by two research grants: Norwegian Research Council 2020-303724: Severity and priority setting in health care, and Swedish Research Council 2021-01266: Distributive justice in health-care – why should severity of ill health matter? (JUST SEVERITY).

McMahan, Jeff. 2019. In Defense of the Time-Relative Interest Account: A Response to Campbell. In *Saving People from the Harm of Death* (Oxford University Press), pp. 267–78 <https://doi.org/10.1093/oso/9780190921415.003.0020>.

Nagel, Thomas. 1970. Death. *Noûs*, **4**.1: 73–80 <https://doi.org/10.2307/2214297>.

Parfit, Derek. 1991. *Equality or Priority* (University of Kansas, Department of Philosophy) <https://kuscholarworks.ku.edu/handle/1808/12405> [accessed 7 June 2021].

Rawls, John. 1999. *A Theory of Justice: Revised Edition* (Cambridge, MA: The Belknap Press of Harvard University).

Scanlon, Thomas. 2000. *What We Owe to Each Other* (Cambridge, MA: Harvard University Press).

Solberg, Carl Tollef, and Espen Gamlund. 2016. The Badness of Death and Priorities in Health. *BMC Medical Ethics*, **17**.1: 1–9 <https://doi.org/10.1186/s12910–016–0104–6>.

Temkin, Larry S. 2014. *Rethinking the Good: Moral Ideals and the Nature of Practical Reasoning* (Oxford University Press).

Tersman, Folke. 2008. The Reliability of Moral Intuitions: A Challenge from Neuroscience. *Australasian Journal of Philosophy*, **86**.3: 389–405 <https://doi.org/10.1080/00048400802002010>.

University of Pittsburgh, Department of Critical Care Medicine, School of Medicine. 2020. *Allocation of Scarce Critical Care Resources During a Public Health Emergency* <Allocation of Scarce Critical Care Resources During a Public Health Emergency> [accessed 23 December 2021].

Williams, Alan. 1997. Intergenerational Equity: An Exploration of the "Fair Innings" Argument. *Health Economics*, **6**.2: 117–32 <https://doi.org/10.1002/(SICI)1099–1050(199703)6:2<117::AID-HEC256>3.0.CO;2-B>.

Cite this article: Jølstad B, Juth N (2022). Age and Illness Severity: A Case of Irrelevant Utilities? *Utilitas* **34**, 209–224. https://doi.org/10.1017/S0953820822000024

2

Medicine, Health Care and Philosophy (2023) 26:413–423
https://doi.org/10.1007/s11019-023-10155-x

SCIENTIFIC CONTRIBUTION

Adaptation and illness severity: the significance of suffering

Borgar Jølstad[1,2]

Accepted: 16 April 2023

Abstract

Adaptation to illness, and its relevance for distribution in health care, has been the subject of vigorous debate. In this paper I examine an aspect of this discussion that seems so far to have been overlooked: that some illnesses are difficult, or even impossible, to adapt to. This matters because adaptation reduces suffering. Illness severity is a priority setting criterion in several countries. When considering severity, we are interested in the extent to which an illness makes a person worse-off. I argue that no plausible theory of well-being can disregard suffering when determining to what extent someone is worse-off in terms of health. We should accept, all else equal, that adapting to an illness makes the illness less severe by reducing suffering. Accepting a pluralist theory of well-being allows us to accept my argument, while still making room for the possibility that adaptation is sometimes, all things considered, bad. Finally, I argue that we should conceptualize adaptability as a feature of illness, and thereby account for adaptation on a group level for the purposes of priority setting.

Keywords Health care priority setting · Well-being · Consequentialism · Adaptation

> "… if adaptation to certain conditions takes place, this raises the difficult but unavoidable question if resource allocation decisions should take this into account as decision-makers may choose to prioritize conditions for which adaptation is less likely achieved." (de Hond et al. 2019).

> "I believe … that the hedonistic conception of happiness and especially of suffering, has enormous moral significance in its own right. It is what constitutes much if not most of the moral horror in experiences such as torture, starvation, military combat, disease, humiliation, clinical depression, and psychosis." (Mayerfeld 1996).

The phenomenon of adaptation to illness has received a fair amount of interest in the literature on health care priority setting. This is partially due to the debate on whether health state valuations should be based on patient valuations or hypothetical valuations by the public, and the skepticism towards priority setting on the basis on adapted valuations which to an extent fuels this debate. In this paper I intend to investigate an aspect of adaptation that seems so far to have been mostly overlooked: that adaptation reduces suffering, and that some health conditions are difficult, or even impossible, to adapt to. When allocating resources in health care, several countries rely on considerations of the severity of illness,[1] or other measures of who is worse off health-wise. Illness severity is influenced by many factors, one of which plausibly is the amount of suffering caused by the condition in question. Adaptation has the effect of reducing suffering, and suffering is a relevant factor on any reasonable theory of well-being relevant for health. If we disregard the effects of adaptation the result is a relative underestimation of the severity of conditions where adaptation is unlikely, resulting in a ranking of the severity of illness that does not accurately track the effects health conditions have on well-being.

The main objective of this paper is to make the argument that adaptation should count when considering to what extent someone is severely ill. Severity functions as a concept with both descriptive and normative features; it is both a description of an illness and it functions as a marker for priority. As such I will attempt to answer two questions: (1) does adaptation lead to a reduction in suffering that leaves a patient better off? And (2) is this reduction in suffering

✉ Borgar Jølstad
odd.borgar.jolstad@ahus.no

1 The Health Services Research Unit—HØKH, Akershus University Hospital (Ahus), Sykehusveien 25, Akershus Universitetssykehus HF, Postboks 1000, 1478 Lørenskog, Norway

2 Institute of Health and Society, Faculty of Medicine, University of Oslo, Oslo, Norway

---

[1] I use 'severity of illness', 'illness severity' and 'severity' interchangeably through the article.

relevant when allocating healthcare resources? It is possible that severity, rather than being sensitive to worse off-ness per se, is sensitive to worse off-ness relevant for healthcare priority setting. Jølstad and Juth (2022) have recently considered, and rejected, such a relevance claim regarding the relationship between severity and worse-off-ness as a function of small differences in age, but perhaps the suffering associated with non-adaptation should be considered irrelevant. As regards to (1) I will argue that the reduction in suffering associated with adaptation is good, but that we should accept that adaptation can still be bad for a person, all things considered, due to factors other than suffering. I argue that the best way to preserve this possibility is to accept a theory of well-being that includes both objective and subjective factors. As regards to (2) I will argue that the suffering and lack off suffering associated with adaptation and non-adaptation is relevant for illness severity, and that adaptation can, and should, be assessed on a group level. Accordingly, adaptability is a feature of an illness that is relevant for severity and priority setting in health. The claim is not that adaptation should be a separate priority setting criterion, but that in assessing illness severity we must be sensitive to the effects of adaptation.[2]

I begin by introducing the notions of illness severity, health states, suffering, and adaptation to illness. The concept of adaptive preference is then discussed briefly. I then argue that adaptation makes an illness less severe in at least one sense: it reduces suffering, and this is relevant on any plausible theory of well-being. I then address the question of whether, and how, this reduction in suffering is relevant for distributional justice.

Illness severity

When allocating health care resources, we typically wish to get as much health as possible, but it also matters to whom benefits accrue. Severity, or similar concepts such as need, is therefore used as a priority setting criterion in several countries, including the Netherlands, Norway and Sweden (Barra et al. 2020). The National Institute for Health and Care Excellence in the UK has also recently decided to adopt a severity modifier (National Institute for Health and Care Excellence 2022). In Norway, for instance, severity is a priority setting criterion alongside cost and benefit of interventions, with illness severity being operationalized as absolute (QALY) shortfall: the severity of a health condition is a function of how many good life years are lost compared to a reference life (Barra et al. 2020). The Norwegians and others who prioritize based on severity are willing to spend more resources *for the same benefit* to a severely ill person than to a less severely ill person. The importance of severity as a priority setting criterion is supported by the intuition that being badly- or worse off is morally important (Barra et al. 2020). Egalitarian, prioritarian and sufficientarian theories of distributive justice are all sensitive to the extent that someone is either badly- or worse off (Hirose 2014). A number of factors, including pain, disability, anxiety, loss of life years, reduced social functioning, age of the patient, risk of death, and others, could plausibly be claimed to contribute to illness severity (Barra et al. 2020; Jølstad and Juth 2022).

For the purposes of this article, I will assume two things about illness severity. The first is that illness severity is a descriptor of how bad the health condition is for a particular patient.[3] A patient suffering from a severe health condition is, ceteris paribus, worse off than if their health condition was less severe, and vice versa. Secondly, I assume that severity constitutes a pro tanto reason to prioritize a patient for treatment. Patients suffering from a severe health condition should, ceteris paribus, be prioritized relative to patients suffering from a less severe health condition.[4] Given these assumptions about severity my argument is relevant for the broader questions of *who is worse off when it comes to health and whom we should aid*. I develop the argument specifically in the context of illness severity for two reasons: first, as already mentioned, severity is used as a priority setting criterion in several countries. Answering the question of who is most severely ill is therefore an important practical question. Secondly, severity is a concept with intuitive moral significance in health care priority setting. Those with severe health conditions seem to merit our concern.

Illness and suffering

Neither human ailment nor suffering are straightforward to define. One way of conceptualizing human ailment is to operate with the concepts of disease, illness, and sickness. In this classification, disease refers to an organic phenomenon, illness to the subjective aspects of a health condition, and sickness to the social aspects of health problems (Hofmann 2002). I intend to talk about something that includes all the above, using the terms illness and health condition as synonyms. While using "illness" in this sense is not aligned with some conceptions in the literature (e.g., Hoffman's use), I

---

[2] I would like to thank an anonymous reviewer for making me aware that my argument could be read as arguing that adaptation should be a separate priority setting criterion.

[3] It may be a function of other parameters as well, but this is not anything I take a stand on here.

[4] Although there may be other criteria for priority setting, for example cost and benefit of intervention.

find other generic terms (e.g., malady) somewhat awkward and hope that the reader will not be confused. I will not take a stance on how we should define illness as an ontological category, but my argument presupposes that the illness itself and the experience and impact of the illness can to some extent come apart. This seems to be a prerequisite for speaking about adaptation in the first place. There is thus the illness itself, with its associated or defining properties, and the experience and impact of the illness for a patient, with adaptation to at least some extent mediating the relationship between the two. One feature of an illness is the degree to which it is possible to adapt to it.

Suffering is something we know when we experience it, but defining it is not trivial. As is clear from the at the start of the article, Jamie Mayerfeld (1996) relies on a hedonistic account of suffering in his discussion, where suffering is understood in terms of negative experiential states. According to Bjørn Hofmann (2015), definitions of human suffering can be grouped according to whether they focus on threats to human agency, profound losses that impair life, or experienced negative sensations. While these forms of suffering might be interesting, I will follow Mayerfeld and focus on hedonic suffering in this paper.

Adaptation to illness

The idea of measuring the value of different health states has been the subject of much scrutiny and criticism, especially by advocates of a mere-difference views on disabilities (Barnes 2009). A crucial point in this debate, and the essential point for the purposes of the argument presented here, is that the majority fails to understand the perspective of people living with disability. When considering life with a disability non-disabled people tend to imagine that life would be miserable. Research on the well-being of people with disabilities largely negates this view. People with disabilities report levels of happiness close to those of non-disabled people (Albrecht and Devlieger 1999). More generally, there tends to be a discrepancy between how the public and patients value health states, with patients typically valuing their health states more highly than healthy people value the same health states hypothetically (Damschroder et al. 2005). One reason for this seems to be that healthy people underestimate the level of adaptation (Ubel et al. 2005). Adaptation involves changing oneself in response to new circumstances (Menzel et al. 2002). This highlights that adaptation is adaptation to something, some *new circumstance.* The new circumstance remains unchanged during adaptation, whereas the persons attitude, experience, or handling of the circumstance changes. Menzel et al. (2002) list the following forms of adaptation: cognitive denial of dysfunction, suppressed recognition of full health, skill enhancement, activity adjustment, substantive goal adjustment, altered conception of health, lowered expectations, and heightened stoicism. Hedonic adaptation is also plausibly a factor in adaptation to illness (Mitchell 2018). Assuming that illness severity is a function of how bad a health condition is for a person, the public tends to view certain health conditions that patients adapt to as more severe than do patients suffering from these conditions.

Issues regarding this overestimation of the severity of conditions that patients adapt to has primarily been explored in the context of disability discrimination, especially through the notions of "double jeopardy" and "the QALY trap",[5][6] In this paper I focus on the converse aspect of this issue: namely the problem of non-adaptation. Some conditions are difficult, or maybe even impossible, to adapt to. Chronic pain is for example known to have a substantial and long-term effect on well-being (Daniel Kahneman and Krueger 2006). Patients suffering from depression value their health states lower than the general public (Pyne et al. 2009). It seems plausible to argue that these conditions are more severe than what the public thinks, and that this is at least partially because they have the feature of being difficult to adapt to.[7]

Adaptive preferences

Before delving into the relationship between adaptation and suffering, I wish to briefly consider the concept of adaptive preferences. There is an obvious affinity between adaptation to illness and the notion of adaptive preference formation, and this matters because adaptive preferences are generally

---

[5] «Double jeopardy" refers to the fact that people suffering from illness are, due to their lives containing fewer QALYs, less cost-effective to save when lives are at stake (Singer et al. 1995).

[6] The QALY trap is described in the following quote: "If we want to value interventions that raise people's quality of life, we are forced to give less value to saving the lives of those with lower quality of life. On the other hand, if we want to give equal value to saving the lives of those with lower quality of life, we are forced to give no value to raising people's quality of life" (John, Millum, and Wasserman 2017).

[7] Adaptation is separate from coping, which can be defined as "… the thoughts and behaviors mobilized to manage internal and external stressful situations. It is a term used distinctively for conscious and voluntary mobilization of acts…" (Algorani and Gupta 2022). While adaptation is a mostly relevant for long term illness, coping mechanisms are also important when dealing with illness in the short term. By focusing on adaptation, I do not intend to make any claims about the relevance of duration for the severity of illness. I focus on adaptation because it is a point of contention in the literature, where there are conflicting intuitions and theoretical commitment. While I do not discuss coping (denoting to short term illness), I believe that my arguments could be relevant for coping, in the sense that illnesses that are hard to cope with are extra burdensome. I thank an anonymous reviewer for bringing the distinction between coping and adaptation to my attention.

considered problematic for the purposes of distributive justice. Adaptive preferences have been the subject of extensive debate in the literature on ethics, just distribution, and rationality during recent years. Although intuitions on the matter are far from homogenous, there seems to be some agreement that adaptive preferences are "…shaped … by facts of, or perceptions of, availability or possibility." (Dorsey 2017). Put simply, adaptive preference formation is a process where what you want is influenced by what you get. Adaptive preferences seem most problematic when they are shaped by unfortunate circumstances, especially if these circumstances can be avoided by human action, such as oppression. Disregarding adaptive preferences has been argued for on the basis of their irrationality (Eftekhari 2021), their undermining autonomy (Colburn 2011; Elster 1982) and the effects that accounting for them would have as regards the just distribution of resources (Sen 2009). It is common in the literature for the term not only to refer to preferences proper, but also to '… desires, values, commitments, and some beliefs and features of a person's character.' (Terlazzo 2017). Polly Mitchell (2018) has argued that adaptation to illness is not, for the most part, an instance of adaptive preference. She argues for taking patient valuations influenced by adaptation into consideration based on their not being adaptive preferences and because we should be careful not to engage in "denial of testimony". My argument, though it engages with questions like the questions raised in the literature on adaptive preference formation, is neutral as to whether adaptation to illness involves adaptive preference formation. If they are, this does not necessarily undermine my arguments. To what extent we should disregard adaptive preferences is itself a matter of debate. Jessica Begon (2020) has recently defended a distinction between "well-being adapted preferences" that align with classic conceptions of adaptive preferences, and "justice adapted preferences", which are '… preferences that are a poor guide to individuals' entitlements.'. Some preferences might be adaptive in the first sense while being non-problematic for the purposes of the second, and vice versa. It is thus not clear that the questions of how a preference is formed and of whether a preference is relevant for resource claims come to the same things. So, even if problematic instances of adaptation are cases of adaptive preference, this does not necessarily justify dismissing the suffering entailed by non-adaptation.

Does adaptation make you better off?

The first argument I wish to develop is that adaptation is relevant for severity by making people better off. This involves a series of claims. The first is that adaptation reduces suffering, leaving patients, ceteris paribus, better off. This claim has previously been defended by Torbjörn Tännsjö (Tännsjö 2019). The second claim is that disregarding adaptation amounts to disregarding non-adaptation, and that this means disregarding suffering. The third is that no plausible theory of well-being relevant for health can disregard the importance of suffering. Together with the assumption that severity is (at least partly) a function of worse off-ness, these claims support the conclusion that adaptation, ceteris paribus, makes a health condition less severe.

Adaptation and suffering

One of the strongest reasons to be skeptical about accepting adaptation is that people adapt to circumstances they *should not* have to adapt to. Amartya Sen (2009) has forcefully argued this point regarding adaptive preference formation in response to poverty and deprivation. If all we consider relevant when determining how we should distribute is the happiness or preference satisfaction of individuals, then how are we to make a distinction between a poor person who is miserable due to lack of basic goods and a rich person who is, subjectively, equally miserable due to a lack of champagne and caviar? This is a compelling general argument against combining a pure subjective theory of well-being with a welfarist theory of justice. There is surely something to be said for not accounting for adaptation, for the sake of distributive justice. Consider the large number of destitute persons today. It is a tragedy that so many people are deprived of a proper education, adequate nutrition, and basic freedoms. It seems that this is cause for moral concern, regardless of the level of subjective well-being experienced by these people. But regardless of whether adaptation, or adaptive preferences, should be accounted for when distributing goods, it clearly seems to matter when considering someone's level of well-being. Consider two people: person *A* is poor but happy (i.e., has a high level of subjective well-being). Person *B* is poor and miserable. Who is worse off? I think that most people will agree that it is preferable to be poor and happy rather than poor and miserable.[8] In fact, this seems to be a clear case of one alternative dominating the other.

Sen's argument against adaptive preferences implies that not discounting for the effects of adaptive preference formation, i.e., including adaptation, will result in unfair distribution. A similar point has been made regarding "bad" forms of adaptation. Menzel et al. (2002) argue that some forms of adaptation, most notably suppressed recognition of full health, cognitive denial, and lowered expectations, should

---

[8] But for a contrary argument see Rickard (1995), who contends that adapting to degrading circumstances itself adds to the degradation, and therefore makes a person worse off.

not be considered when distributing resources.[9] It seems intuitive to exclude the effects of these forms of adaptation, because they are instances of poor reasoning or irrational acceptance of poor (or irrational denial of good) circumstances.[10] There seems to be two issues here: the familiar issue that including adaptation leads to an unfair distribution, and the intuition that some of these forms of adaptation are not good for people, all things considered. Some cases of adaptation, such as skill enhancement and goal adjustments, seem better than others, such as denial, for the people concerned. But while the argument that accounting for adaptation will lead to unfair distribution seems plausible, the argument that adapting does not increase someone's well-being is more suspect: there is of course something bad about people having to resort to suppressed recognition of full health, cognitive denial, or lowered expectations, to endure their situation, but the situation would be even worse if these forms of adaptation did not reduce suffering. Even though we would prefer a world where people did not have to adapt to illness, there is undoubtedly one good thing to be said for adaptation, namely that it reduces suffering. Note that this is purposefully a weak claim. Rosa Terlazzo (2017, 2022) has made the stronger claim that objects of adaptive preferences can start out being bad for us, but become good due to personal development over time. I might, for example, change my valuation of what a good life consists in in response to mobility issues. If we accept this claim, the case for counting adaptation is even stronger.

Discounting adaptation, "bad" or not, amounts to a disregard of a difference in suffering when considering who is worst off. This, I believe, is inherently implausible. But surely some cases of adaptation are bad for people, all things considered, even if suffering is reduced. I will return to this question towards the end of the section and argue that accepting that well-being has both subjective and objective components allow us to accept the subjective benefits of adaptation, while leaving room for the possibility that some forms of adaptation are, all things considered, bad for people.

Non-adaptation and suffering

Research suggests that adaptation to illness is highly domain-specific; some illnesses are easier to adapt to than others (Greene et al. 2016). Likely candidates for *non-adaptable health conditions* are certain forms of mental illness and chronic pain conditions. While not the only conditions that seem to resist adaptation (cancer and stroke are also difficult to adapt to), mental health conditions greatly impact subjective well-being and have empirically been found difficult to adapt to (Binder and Coad 2013). Chronic pain also seems to be difficult to adapt to (Greene et al. 2016), and has a substantial and lasting effect on subjective well-being (Daniel Kahneman and Krueger 2006). There are several theories as to why people do not adapt to these conditions: Binder and Coad (2013) suspect that the uncertainty involved in mental illness and pain is important, whereas Daniel Kahneman (2008) argues that adapting to depression and pain is difficult because "… the normal process of withdrawing attention from a steady situation is prevented." Pain and various forms of mental illness force us to attend to them, making it hard to move on. It even seems plausible to argue, given the inherently subjective nature of the illness, that adapting to depression is a contradiction in terms. If disregarding non-adaptation implies that we underestimate the severity of illnesses such as depression, the result is a distributive scheme that fails to prioritize people with these conditions to the extent that their life is worse.

So, does disregarding adaptation imply disregard of non-adaptation? I will present two arguments for this, one analytic and one practical.

The most straightforward argument seems to be this: if one views adaptation as a good thing (and I have argued that even "bad" forms of adaptation include some goodness in the form of reduced suffering) then it is reasonable to view non-adaptation as a bad thing, all else equal. If we deem adaptation valuable because of its effect of reducing suffering, then not adapting is a disvalue. Returning to a person who requires a wheelchair after an accident, many of us would consider this a tragic, or at least a significantly bad, state of affairs. If this was a friend, we would wish that this person was able to lead a life of meaning and enjoyment notwithstanding their poor luck. We would clearly not be indifferent as to whether they managed this or not. We would find joy in the relief of suffering found by acquiring new skills or a changed perspective, and pity them if they were unable to accomplish this, resulting in more suffering. If one is indifferent to whether people adapt to their conditions, then it seems that one is forced to be indifferent as to the suffering caused by non-adaptation. In the context of considering to what extent someone is severely ill, or more broadly worse off because of their health, this seems unreasonable.

---

[9] They are unsure as to whether heightened stoicism makes a case for or against adaptation.

[10] The question of the rationality or irrationality of adaptive preferences is complicated, especially because the process of character planning does not seem problematic or irrational to the same extent as other adaptive preferences (Elster 1982). If I change my preferences in response to my circumstances, there seems to be something rational about it to the extent that it improves my well-being. If I, on the other hand, kept my adaptive preference in the face of a change in circumstances, this might seem more irrational.

The second argument regards the relative severity of different health conditions, and the implications that these rankings of severity have for distributive matters. Priority setting is essentially a matter of relative distribution: some get more, some get less, and some get nothing at all. In health care systems such as the Dutch, Norwegian, or Swedish, who is prioritized is partly determined by who is most severely ill (Barra et al. 2020). When an illness is considered less severe, it will, all else equal, be outranked by illnesses considered more severe. Overestimating the severity of illnesses that patients typically adapt to means, relatively speaking, underestimating the severity of illnesses that patients do not typically adapt to. There is no neutral point when it comes to illness adaptation. Disregarding adaptation leads to disregarding non-adaptation in our rankings of illness severity.

Theories of well-being

It might be objected that I have implicitly been assuming a hedonist theory of well-being in the previous arguments. On a hedonist account peoples well-being is defined by their subjective experience. My argument requires that subjective well-being, or more precisely the degree to which someone is suffering, is a necessary component of any theory of well-being relevant for health. I will now consider whether the argument holds on a desire satisfaction theory or objective theory of well-being. Regarding the desire theories the crucial question is what it means for a desire to be “informed”. Regarding objective theories the central question is whether a plausible theory of well-being can do without a subjective well-being component. I first turn to desire satisfaction theories**.**

Desire satisfaction theory

Desire satisfaction, or preferentialist, theories claim that what a person’s well-being consists of is the satisfaction of their desires or preferences (Heathwood 2015). Much of the plausibility of these theories comes from the intuition that for *X* to be good or bad for a person, that person must care about *X*. As James Griffin (1986) has argued: Caviar might be considered great eating, but if I do not like caviar, then arguing that feeding it to me makes me better off is difficult. What is bad, on most forms of desire satisfaction theory, is frustrated desires. How to handle suffering on this account is not straightforward, because suffering and frustration do not seem synonymous. For the purposes of this paper, I will assume that a person is badly off from suffering because they have a frustrated desire not to suffer.[11] For the purposes of this paper the main issue is to what extent a desire must be informed for it to contribute to a person’s well-being.

Griffin (1986) argues that any plausible desire satisfaction basis for welfare will demand that desires are to some extent informed. If my argument succeeds there are important goods associated even with “uninformed” preferences. Note that preferences can be uninformed in two ways: they can be faulty, and they can be formed by a dubious process. There seem to be two relevant points to be made when considering to what extent adapted preferences are informed:

(1) Patients have a phenomenological closeness to their health states.
(2) Adapted desires might be the result of cognitive distortions, irrationality, denial, or similar factors.

Considering the first point, Versteegh and Brouwer (2016) make an interesting argument: the public (and patients before becoming ill) do not know how it is to have a certain disease, but patients do not have direct phenomenological access to how it is to be healthy either.[12] This argument is most plausible when considering patients who have never been in close to full health. However, many patients, in contrast to the public, have experienced both good health and the health states they are asked to evaluate. It thus seems likely that when considering tradeoffs between their health states and full health, these latter patients usually have a phenomenological edge when considering to what extent they are badly off. This leaves us with the question of cognitive distortions. Perhaps patients are in denial of their health or have shifted their expectations. At least, we cannot assume that this is never the case. Then, perhaps, the satisfaction of the desires that are based on these distortions should be considered less valuable. However, this is not the end of it. At this point we are entering familiar terrain. Whether or not patients have adapted in “bad” ways, surely their reduced suffering matters for their well-being. Perhaps patients are cognitively distorted when it comes to their health, but they are not plausibly completely misguided as to their subjective level of suffering, and we can assume that most people have a desire not to suffer.

---

[11] This seems somewhat forced. Bengt Brülde (2007) argues that what is bad on a desire theory is *aversion satisfaction*. This would make it easier to handle factors such as pain and suffering on a desire theory.

[12] Strictly speaking they frame their point in the wording of ex-ante and ex-post valuations, given that valuation exercises involve non-experienced health states.

Objective theories of well-being

Objective theories of well-being claim that at least some factors influence well-being regardless of whether people subjectively value them or not (Parfit 1984). Prominent theories include various forms of perfectionism and objective list-theories.

On perfectionist theories of well-being, what matters for well-being is to what extent someone either has, develops, or makes use of essentially human capacities (Bradford 2015). Bradford (2017) has argued that perfectionisms failure to deal with the importance of pleasure makes perfectionism less plausible as a theory of well-being. Importantly, this only holds for perfectionist theories that do not include a subjective component. Perhaps the capacity to be happy is a human capacity that can be had or developed to various extents? In that case perfectionism could make sense of the relevance of suffering. Unless such a factor is included in a perfectionist theory, I agree with Bradfords point and add to this that a failure to deal adequately with suffering would also make perfectionism an unreasonable theory of well-being.

Objective list accounts of well-being generally have two features in common: they are attitude independent, in the sense that whether *X* is good for a person can be independent on that person's attitude to *X*, and they are pluralistic. The second feature is not embraced by all theorists (some view hedonism as an objective list with one item) but the first is universally endorsed (Fletcher 2015). Attributes of certain illnesses might, on an objective list account, be bad for a person regardless of their attitude or experience; they may be attitude independent. Perhaps being blind, or having a mobility issue, is *intrinsically worse* than not, even if you do not value these things. If this is the case, there are factors making persons worse off health-wise that cannot in themselves be improved by adaptation. Importantly though, there seems to be a subjective component associated with these objective factors that is clearly improvable by adaptation. In fact, the concept of adaptation seems to require such a component. The previously mentioned candidates for non-adaptable health conditions, depression and chronic pain, are defined by their subjective experience, leaving less room for adaptation. The question then becomes whether it is plausible to assume that any account of well-being relevant for health can be indifferent as to subjective experience, and in particular suffering? An objective list without experienced well-being on the list seems implausible. If one where to come up with a list of factors that make your life go badly, suffering seems like an obvious contender. This is reflected in the works of philosophers working with an objective account. In Derek Parfit's (2011) development of a value based objective account of reasons, pain is used as a paradigmatic case of an objective reason. T. M. Scanlon (2000), in his arguments against consequentialism, still acknowledges that we have a prima facie reason to prevent or avoid causing pain. On Daniel Hausman's (2015) theory of the value of health, which can arguably be considered an objective theory, suffering/distress is one of two factors alongside functioning/limitations. This is not to claim that these authors accept a subjective theory of well-being, but simply to illustrate that it generally seems implausible to claim that you are not in some important sense worse off suffering than not suffering.

Returning to "bad" adaptation

I believe that I have made the case that there is a positive side even to bad forms of adaptation, and that this goodness should be taken into account when considering the severity of illness. But this still leaves the question of whether all forms of adaptation are good, all things considered. The primary argument against counting adaptation, or adaptive preferences, seems to be the sour grapes argument. This is, as I argued previously, essentially an argument against combining a subjective account of well-being with a welfarist theory of distributive justice. Surely people's objective circumstances matter. The argument seems to apply both to desire satisfaction and hedonist theories of well-being. The problem is, essentially, that there are important objective factors that the subjective theories fail to consider. I argued that it is worse to poor and miserable than poor and happy. This seems to be true, to the extent of being undeniable. But on a pure subjective account, the poverty involved can potentially loose its importance. In the same way, the fact that some forms of adaptation, such as denial or reduced expectation, are problematic, might be overlooked. On a pure subjective theory, we seem unable to differentiate between "good" and "bad" forms of adaptation in a way that aligns with our intuitions. In this regard an objective theory of well-being seems better poised to help us deal with the problem, provided it includes a measure of subjective well-being. Let us consider a new case of adaptation to illness: *A* has a massive amount of scarring to visible areas of their body because of a botched surgery. *A* previously had an active social and professional life, doing things that were very much worth doing. In response to the surgical scarring *A* becomes a recluse, spending their days indoors playing video games. After some time, *A* is happy spending their days gaming, and no longer has the wish to have an active social and professional life. Taking my argument seriously, there is obviously something good about *A* being happy doing what they are doing. It would clearly be worse if they were miserable. But at the same time is seems relevant that *A* no longer has an active life, doing things of substance and importance. On an objective theory we would be able to say that the adaptation is good in the sense that *A* is happy, but bad in the sense that they have given up things that are independently valuable. We could then weigh these factors when considering whether the adaptation makes the illness less severe. If we believe that the badness outweighs the goodness, the adaptation in question is bad, all things considered. This solution runs into the problem that we perhaps fail to show *A* what Rosa Terlazzo (2014) calls secondary recognition respect: we fail to recognize *A* as an authority on their own good. Any objective theory will plausibly run into this problem. While I am sympathetic to the idea that we owe people something like secondary recognition respect, I follow Serene Khader (2012) in rejecting the idea that '…it is disrespectful to treat another as though she has *some* bad values.'. As Khader reminds us, being subject to criticism is part of what it is to be a moral agent. On my account there are aspects of a persons' well-being that we cannot plausibly criticize, namely their subjective experience. Any increase or decrease in subjective well-being should plausibly lead us to reassess the severity of a patient's condition. Other aspects, the objective features, must be subject to debate and deliberation.

In assessing the overall severity of an illness or health condition, we must then weigh subjective and objective factors. While more needs to be said of such a theory, it would allow us to make sense of our intuitions in situations where subjective and objective features of well-being come apart or collide. Consider the following case reported by Mendez and Parand (2020), where a man presented with various cognitive impairments and a markedly joyful mood after a gunshot to the head. This was particularly striking because he had previously been disposed to suicidal ideation and depression. They write: "His mood appeared excessively good for the context, and his affect was congruent with mood but not overtly euphoric. He did not seem to be aware of the severity of his memory impairment, but he was aware that he had significant cognitive problems." (Mendez and Parand 2020). Both a purely subjective and a purely objective theory of the value of health would miss out on important aspects relevant for well-being in this case. The pattern of a seeming disconnect between objective and subjective factors can be found in more common health conditions. While many people with Down's syndrome struggle with various physical and cognitive problems, nearly 99% report being happy with their lives (Skotko et al. 2011). Cass Sunstein (2008) uses the case of Down's Syndrome to illustrate how legal damages can be justified based on capability losses even when subjective well-being is unaffected. Similarly, a mixed theory of the value of health allows us to recognize that people with downs syndrome have high levels of subjective well-being, while not disregarding their medical and social needs in the process of allocating health care resources. So, as well as enabling us to make sense of the good and bad of adaptation, a theory including both objective and subjective factors give us the right result when making sense of our intuitions in other cases where subjective well-being and other aspects of health diverge.[13]

Should adaptation influence healthcare priority setting?

So, I have established that adaptation can lead to a reduction in suffering that is relevant for assessing a person's well-being, making it clear that adaptation is relevant for the worse off aspect of severity. I will now consider the relevance of this reduction in suffering for healthcare priority setting, and thus for the priority aspect of severity. I will argue that we should focus on adaptability as a feature of an illness. If some illnesses are easier to adapt to, leading to less suffering, this is relevant for distributional concerns. As mentioned in the introduction, the argument is not that adaptability should be a separate criterion for priority setting, but rather that it is a feature that we should account for when assessing the severity of illness.

Recall that any theory of well-being relevant to matters of health that does not give at least some weight to suffering is deeply implausible. In fact, not being able to deal with subjective suffering seems to make a theory something else than a theory of well-being as commonly understood. On any plausible consequentialist theory based on well-being relevant for health, suffering should be part of our ethical calculus.[14] Allowing suffering to be overridden by other factors seem fine (if we accept an objective theory or are pluralists about well-being, it might often be the most reasonable option) but claiming that health related worse offness is not affected at all by whether a patient suffers seem unreasonable. We should therefore take illness adaptation into account when considering illness severity and when distributing resources, or so the argument goes. This can be problematic due to distributional considerations like those raised by Sen (2009), particularly if we restrict ourselves to

---

[13] I believe that Daniel Hausman's (2015) previously mentioned theory of the value of health could allow us to make sense of our intuitions on adaptation and illnesses without losses in subjective well-being, precisely because it includes both subjective (suffering/distress) and objective (functioning/limitations) factors.

[14] If we would like to distribute in a manner that does not take adaptation into account, while retaining theoretical consistency, perhaps the way to go is to focus on means for living well rather than on well-being per se. This focus on means rather than ends is shared by Rawls theory of justice (Rawls 1999) and Nusbaum and Sens capability approaches (Nussbaum 2011; Sen 2009). I will not consider this here, other than to note that any consequentialist theory focusing on a balance between the maximization of and fair distribution of well-being, will be susceptible to my arguments.

a purely subjective theory of well-being. Adopting a pluralist theory of well-being might allay some of these worries, by letting factors other than suffering, such as basic flourishing, play a major part in our judgments of severity. We would then be able to consider the suffering involved in illness, while also being sensitive to other concerns.

So why should we focus on adaptability as a feature of illness? Prioritizing based on individual adaptation would clearly be a contested issue, largely because of practical concerns, but perhaps also due to considerations of desert or responsibility. The distinction between adaptability of illness and individual adaptation can isolate an argument that is likely to be less contested. Prioritizing based on features of illnesses, such as pain, reduction in life span, or loss of functioning, is an essential and non-controversial aspect of health care priority setting. On the practical side, assessing the extent to which illnesses are adaptable seems like a tractable problem, whereas assessing individual adaptation looks more problematic.[15] On the question of desert or responsibility, the lack of adaptation to severe pain or depression does not seem to be as much a feature of particular people as of the illnesses in question. We are not plausibly responsible for features of our illness, making objections based on responsibility or desert mute. This argument should thus be plausible even to those who do not consider all forms of subjective well-being relevant for distributional concerns. Letting adaptability as a feature of illness influence our judgments of severity will result in sensitivity to the possibility of adaptation, while avoiding the problems associated with individual assessments. Accepting this argument might lead to prioritizing some conditions, such as depression and chronic pain conditions, to a larger extent that what we are currently doing. I mentioned in the introduction that the question of adaptation is important partly due to the controversy on whether we should use experienced or hypothetical measures when valuing health states. Another consequence of accepting the relevance of adaptability for severity is that adaptation is an argument *in favor* of including measures of experienced health when valuing health states for the purposes of priority setting. This is because the reduction in suffering achieved by adaptation is most effectively accounted for by asking patients to value their own health states. If we accept a pluralist or objective list conception of well-being, the most reasonable way forward is perhaps to include both measures of experienced health *and* measures of objective factors.

Conclusion

I have argued that suffering makes a person worse off on any plausible theory of well-being relevant for health. Given that severity is at least partly determined by to what extent an illness makes a person worse off, judgements of the degree of severity must be sensitive to suffering. Non-adaptation to illness means that a person suffers more than if she had adapted to her illness. Unless our judgements of illness severity are to be insensitive to suffering, we are thus compelled, all else equal, to consider illness that one adapts to less severe and illness that one does not adapt to more severe. If we wish to claim that some forms of adaptation are bad for people, the most promising option seems to be to accept an objective theory with a subjective component. On the distributional question, I have argued that there is a strong case for being sensitive to the adaptability of illness when making judgments of severity, at least on a group level.

Acknowledgment I would like to thank Niklas Juth, Mathias Barra, and Carl Tollef Solberg for helpful discussion, comments, and feedback on the article. Previous versions of the article have been presented at the early career workshop at Priorities 2022, the Philosophy of Priority Setting Tuesday seminar, the Just Severity and Severity in priority setting summer seminar, and at the Priorities 2022 conference. I am grateful for feedback and constructive discussion during these events.

Funding Open access funding provided by University of Oslo (incl Oslo University Hospital). The work is supported by NRF Grant No. 303724.

Declarations

Conflict of interest The authors have not disclosed any competing interests.

References

Albrecht, Gary L., and Patrick J. Devlieger. 1999. The Disability Paradox: High Quality of Life against All Odds. *Social Science*

---

[15] Assessing adaptation on a group level would plausibly require studies on how the quality of life of patients, including health related quality of life and subjective well-being, change over time. These studies could be quantitative, such as Binder and Coad's (2013) study on differing illnesses, but could presumably also include qualitative and mixed-methods studies. One way of accounting for the effects of adaptation could also be to include measures of experienced health in health valuations, since patient valuations avoid the underestimation of adaptation associated with hypothetical health valuations.

& Medicine 48 (8): 977–988. https://doi.org/10.1016/S0277-9536(98)00411-0.

Algorani, Emad B., and Vikas Gupta. 2022. Coping Mechanisms. In *StatPearls*. Treasure Island: StatPearls Publishing.

Barnes, Elizabeth. 2009. Disability, Minority, and Difference. *Journal of Applied Philosophy* 26 (4): 337–355. https://doi.org/10.1111/j.1468-5930.2009.00443.x.

Barra, Mathias, Mari Broqvist, Erik Gustavsson, Martin Henriksson, Niklas Juth, Lars Sandman, and Carl Tollef Solberg. 2020. Severity as a Priority Setting Criterion: Setting a Challenging Research Agenda. *Health Care Analysis* 28 (1): 25–44. https://doi.org/10.1007/s10728-019-00371-z.

Begon, Jessica. 2020. Disability, Rationality, and Justice: Disambiguating Adaptive Preferences. In *The Oxford Handbook of Philosophy and Disability*, ed. Adam Cureton and David T. Wasserman, 342–359. Oxford: Oxford University Press.

Binder, Martin, and Alex Coad. 2013. 'I'm Afraid I Have Bad News for You…' Estimating the Impact of Different Health Impairments on Subjective Well-Being. *Social Science & Medicine* 87: 155–167. https://doi.org/10.1016/j.socscimed.2013.03.025.

Bradford, Gwen. 2015. Perfectionism. In *The Routledge Handbook of Philosophy of Well-Being*. New York: Routledge.

Bradford, Gwen. 2017. Problems for Perfectionism. *Utilitas* 29 (3): 344–364. https://doi.org/10.1017/S0953820816000418.

Brülde, Bengt. 2007. Happiness and the Good Life. Introduction and Conceptual Framework. *Journal of Happiness Studies* 8 (1): 1–14. https://doi.org/10.1007/s10902-006-9002-9.

Colburn, Ben. 2011. Autonomy and Adaptive Preferences. *Utilitas* 23 (1): 52–71. https://doi.org/10.1017/S0953820810000440.

Damschroder, Laura J., Brian J. Zikmund-Fisher, and Peter A. Ubel. 2005. The Impact of Considering Adaptation in Health State Valuation. *Social Science & Medicine* 61 (2): 267–277. https://doi.org/10.1016/j.socscimed.2004.11.060.

de Hond, Anne, Pieter Bakx, and Matthijs Versteegh. 2019. Can Time Heal All Wounds? An Empirical Assessment of Adaptation to Functional Limitations in an Older Population. *Social Science & Medicine* 222: 180–187. https://doi.org/10.1016/j.socscimed.2018.12.028.

Dorsey, Dale. 2017. Adaptive Preferences Are a Red Herring. *Journal of the American Philosophical Association* 3 (4): 465–484. https://doi.org/10.1017/apa.2017.36.

Eftekhari, Seena. 2021. The Irrationality of Adaptive Preferences: A Psychological and Semantic Account. *Utilitas* 33 (1): 68–84. https://doi.org/10.1017/S0953820820000321.

Elster, J. 1982. Sour Grapes-Utilitarianism and the Genesis of Wants. In *Utilitarianism and Beyond*, ed. Amartya Kumar Sen and Bernard Arthur Owen. Williams, 219–238. Cambridge: Cambridge University Press.

Fletcher, Guy. 2015. Objective List Theories. In *The Routledge Handbook of Philosophy of Well-Being*. New York: Routledge.

Greene, Edie, Kristin A. Sturm, and Andrew J. Evelo. 2016. Affective Forecasting about Hedonic Loss and Adaptation: Implications for Damage Awards. *Law and Human Behavior* 40: 244–256. https://doi.org/10.1037/lhb0000181.

Griffin, James. 1986. *Well-Being: Its Meaning, Measurement and Moral Importance*. Oxford: Clarendon press.

Hausman, Daniel M. 2015. *Valuing Health: Well-Being, Freedom, and Suffering*. Oxford: Oxford University Press.

Heathwood, Chris. 2015. Desire-Fulfillment Theory. In *The Routledge Handbook of Philosophy of Well-Being*. New York: Routledge.

Hirose, Iwao. 2014. *Egalitarianism*. New York: Routledge.

Hofmann, Bjørn. 2002. On the Triad Disease, Illness and Sickness. *The Journal of Medicine and Philosophy: A Forum for Bioethics and Philosophy of Medicine* 27 (6): 651–673. https://doi.org/10.1076/jmep.27.6.651.13793.

Hofmann, Bjørn. 2015. Suffering: Harm to Bodies, Minds, and Persons. In *Handbook of the Philosophy of Medicine*, ed. Thomas Schramme and Steven Edwards, 1–17. Dordrecht: Springer.

John, Tyler M., Joseph Millum, and David Wasserman. 2017. How to Allocate Scarce Health Resources Without Discriminating Against People With Disabilities. *Economics & Philosophy* 33 (2): 161–186. https://doi.org/10.1017/S0266267116000237.

Jølstad, Borgar, and Niklas Juth. 2022. Age and Illness Severity: A Case of Irrelevant Utilities? *Utilitas*. https://doi.org/10.1017/S0953820822000024.

Kahneman, D. 2008. *QALYs versus Experience: A Perspective from Experimental Economics. Monographs*. London: Office of Health Economics.

Kahneman, Daniel, and Alan B. Krueger. 2006. Developments in the Measurement of Subjective Well-Being. *Journal of Economic Perspectives* 20 (1): 3–24. https://doi.org/10.1257/089533006776526030.

Khader, Serene J. 2012. Must Theorising about Adaptive Preferences Deny Women's Agency? *Journal of Applied Philosophy* 29 (4): 302–317. https://doi.org/10.1111/j.1468-5930.2012.00575.x.

Mayerfeld, Jamie. 1996. The Moral Asymmetry of Happiness and Suffering. *The Southern Journal of Philosophy* 34 (3): 317–338. https://doi.org/10.1111/j.2041-6962.1996.tb00795.x.

Mendez, Mario F., and Leila Parand. 2020. Positive Emotions from Brain Injury: The Emergence of Mirth and Happiness. *Case Reports in Psychiatry* 2020: e5702578. https://doi.org/10.1155/2020/5702578.

Menzel, Paul, Paul Dolan, Jeff Richardson, and Jan Abel Olsen. 2002. The Role of Adaptation to Disability and Disease in Health State Valuation: A Preliminary Normative Analysis. *Social Science & Medicine* 55 (12): 2149–2158. https://doi.org/10.1016/S0277-9536(01)00358-6.

Mitchell, Polly. 2018. Adaptive Preferences, Adapted Preferences. *Mind* 127 (508): 1003–1025. https://doi.org/10.1093/mind/fzy020.

National Institute for Health and Care Excellence. 2022. "NICE Publishes New Combined Methods and Processes Manual and Topic Selection Manual for Its Health Technology Evaluation Programmes." NewsArticle. NICE. NICE. January 31, 2022. https://www.nice.org.uk/news/article/nice-publishes-new-combined-methods-and-processes-manual-and-topic-selection-manual-for-its-health-technology-evaluation-programmes.

Nussbaum, Martha Craven. 2011. *Creating Capabilities: The Human Development Approach*. Cambridge: Belknap press of Harvard University Press.

Parfit, Derek. 1984. *Reasons and Persons*. Oxford: Oxford University Press.

Parfit, Derek. 2011. *On What Matters*. Oxford: Oxford University Press.

Pyne, Jeffrey M., John C. Fortney, Shanti Tripathi, David Feeny, Peter Ubel, and John Brazier. 2009. How Bad Is Depression? Preference Score Estimates from Depressed Patients and the General Population. *Health Services Research* 44 (4): 1406–1423. https://doi.org/10.1111/j.1475-6773.2009.00974.x.

Rawls, John. 1999. *A Theory of Justice*. Revised. Cambridge: The Belknap Press of Harvard University.

Rickard, M. 1995. Sour Grapes, Rational Desires and Objective Consequentialism. *Philosophical Studies: An International Journal for Philosophy in the Analytic Tradition* 80 (3): 279–303.

Scanlon, Thomas. 2000. *What We Owe to Each Other*. Cambridge: Harvard University Press.

Sen, Amartya Kumar. 2009. *The Idea of Justice*. Cambridge: Harvard University Press.

Singer, P., J. McKie, H. Kuhse, and J. Richardson. 1995. Double Jeopardy and the Use of QALYs in Health Care Allocation. *Journal* *of Medical Ethics* 21 (3): 144–150. https://doi.org/10.1136/jme.21.3.144.

Skotko, Brian G., Susan P. Levine, and Richard Goldstein. 2011. Self-Perceptions from People with Down Syndrome. *American Journal of Medical Genetics Part A* 155 (10): 2360–2369. https://doi.org/10.1002/ajmg.a.34235.

Sunstein, Cass R. 2008. Illusory Losses. *The Journal of Legal Studies* 37 (S2): S157–S194. https://doi.org/10.1086/595675.

Tannsjo, Torbjorn. 2019. *Setting Health-Care Priorities: What Ethical Theories Tell Us*. Oxford: Oxford University Press.

Terlazzo, Rosa. 2014. The Perfectionism of Nussbaum's Adaptive Preferences. *Journal of Global Ethics* 10 (2): 183–198. https://doi.org/10.1080/17449626.2014.931874.

Terlazzo, Rosa. 2017. Must Adaptive Preferences Be Prudentially Bad for Us? *Journal of the American Philosophical Association* 3 (4): 412–429. https://doi.org/10.1017/apa.2018.1.

Terlazzo, Rosa. 2022. Adaptive Preferences in Political Philosophy. *Philosophy Compass* 17 (1): e12806. https://doi.org/10.1111/phc3.12806.

Ubel, Peter A., George Loewenstein, Norbert Schwarz, and Dylan Smith. 2005. Misimagining the Unimaginable: The Disability Paradox and Health Care Decision Making. *Health Psychology* 24 (4 Suppl): S57-62. https://doi.org/10.1037/0278-6133.24.4.S57.

Versteegh, M.M., and W.B.F. Brouwer. 2016. Patient and General Public Preferences for Health States: A Call to Reconsider Current Guidelines. *Social Science & Medicine* 165: 66–74. https://doi.org/10.1016/j.socscimed.2016.07.043.

3

4