Bracketing Cluelessness

Sylvester Kollin* Jesse Clifton†
Anthony DiGiovanni† Nicolas Macé†

September 8, 2025

Abstract

Consequentialists must take into account all possible consequences of their actions, including those in the far future. But due to the difficulty of getting a grasp on these consequences and producing non-arbitrary probabilities for them, it seems that consequentialists should often consider themselves *clueless* about which option is best. Contrary to orthodox consequentialism, however, there is a common-sense intuition that one should bracket those consequences which one is clueless about. Building on a model involving imprecise probability, we develop two novel alternatives to orthodoxy which capture this intuition. On *bottom-up bracketing*, we set aside those beneficiaries for whom we are clueless what would be best, and then base the overall verdict on the remainder. On *top-down bracketing*, we instead base the overall verdict on what would be best for the largest subsets of beneficiaries relative to which we are not clueless. The two are not equivalent: the former violates statewise dominance, whereas the latter does not. The main objection which applies to both kinds of bracketing is that they do not rank prospects acyclically. Our response includes showing how a natural way of generalising bracketing to the dynamic setting avoids value-pumps. Finally, we argue that bracketing has important implications for real-world altruistic decision-makers, favouring neartermism over longtermism.

*Stockholm University. ORCID: 0009-0008-6976-0127. Part of this work was also carried out at the Center on Long-Term Risk. Comments welcome at: sylvesterkollin@gmail.com.

†Center on Long-Term Risk

Contents

1 Introduction

Consequentialists face an epistemic challenge: when making choices, they must take into account not only the immediate consequences of the options available to them, but *all* of their possible consequences—even those in the far future. But in practice we are unaware of most of these consequences; and even when we are aware, we cannot usually assign anything close to the precise probabilities that expected value theory demands, at least not non-arbitrarily. As a result, it appears that consequentialists are, or should often consider themselves to be *clueless* about which course of action is best (cf. Kagan, 1998, pp. 64-5; Lenman, 2000; Greaves, 2016).

To illustrate the worry, let us consider an example from Mogensen (2021, pp. 141-2). Suppose you are choosing between donating to the Make-A-Wish Foundation (MAWF), which grants wishes to critically ill children, and the Against Malaria Foundation (AMF), which provides insecticidal nets to prevent malaria. On the face of it, AMF appears far more effective, saving a life for a few thousand dollars rather than funding wish experiences (Singer, 2015, pp. 5-6). However, it is arguably very ambiguous how to weigh up the potential long-term and indirect effects of these donations. Saving lives by donating to AMF could lead to increased population growth, which might accelerate economic development and technological progress, but could also contribute to environmental degradation, depletion of resources, or political instability. Even small demographic changes could have vast and unforeseen effects. And so if these far-reaching consequences make up most of the consequentialist value of one's actions (as they quite clearly do), yet remain inscrutable, it seems that consequentialists should be clueless about whether donating to MAWF or AMF is best.

Contrary to orthodox consequentialism, however, there is a common-sense intuition that one should base one's decisions on the consequences that one is not clueless about.[1] For instance, in Mogensen's case, there is an intuition

[1] Another response is to introduce a discount rate on future goods, reflecting a 'rate of pure time preference'. However, first of all, as Ramsey (1928, p. 543) put it, such "a practice [....] is ethically indefensible and arises merely from the weakness of the imagination". Second, in any case, pure discounting does not capture the intuition we are pointing to here. For one can

that one should make decisions based on the clear and immediate effects of donating, i.e., saving lives or granting wishes, and bracket out speculative, long-term consequences like population growth and resource depletion. In this paper, we develop two alternatives to orthodoxy which capture this intuition, building on a model involving imprecise probability. On the first theory, *bottom-up bracketing*, we set aside those beneficiaries for whom we are clueless what would be best (which will soon take on a precise meaning), and then base the overall verdict on the remainder. On the second theory, *top-down bracketing*, we instead base the overall verdict on what would be best for the largest subsets of beneficiaries relative to which we are not clueless. The two theories are not equivalent, and we will favour the top-down approach on the grounds that the bottom-up approach violates statewise dominance.

We address four objections which apply to both theories of bracketing, three of which are based on the fact that the theories do not generally rank prospects in an acyclic way. Chief among these is that the two theories are thus open to being turned into value-pumps, i.e., that they recommend sequences of choices which amount to giving up value for free, and should hence be rejected. Or so the argument goes. We respond by showing how a natural way of generalising bracketing to the dynamic or sequential setting avoids value-pumps.

The remainder of the paper is structured as follows. In §2, we put forward a model of cluelessness in terms of imprecise probabilities and incomplete betterness. We make this model more precise in § 3, introducing a formal framework and some important definitions. In §4, we set the stage by considering some relevant literature on the so-called Super-Strong Pareto principle, and arguments against said principle which also apply to theories of bracketing. We formally define and discuss our two theories of bracketing in §5. In §6, we take up the challenge of cyclicity in the dynamic setting, developing the aforementioned generalisation. We discuss the challenge that bracketing poses to longtermism, roughly the view that what matters most is how our actions affect the far future, in §7, and finally conclude in §8. Both §6 and §7 may be skipped without much loss of central content.

be more clueless about some effects than others, even when those effects are temporally prior. A similar reply would apply if we instead tried to account for this intuition by discounting outcomes with very small probabilities, à la Monton (2019).

2 Cluelessness and credal imprecision

We now want to be clear about how we understand and model consequentialist cluelessness. Our first claim is that the predicament faced by consequentialists should be thought of as an instance of *deep*, *severe* or *great uncertainty* (Hansson, 1996; R. Bradley and Drechsler, 2013; R. Bradley, 2017, pp. 227-32; Helgeson, 2020). There is no consensus with regards to how this kind of uncertainty should be defined, and we will not attempt a definition here. (The three aforementioned terms are in any case sometimes treated as technical ones.) Nonetheless, situations of deep uncertainty—which is the term we will use going forward—are generally understood to involve a severe lack of (or unclear balance of) evidence, no knowledge of frequencies or chances, and unawareness of decision-relevant variables.

Second, we claim that the credences or degrees of belief of an agent in cases of deep uncertainty should be modelled by a non-singleton *set* of probability measures $\mathbf{P}$, known as the *credal set* or *representor* (Levi, 1974; Jeffrey, 1983a; van Fraassen, 1990; Mahtani, 2019). Such a model is a significant improvement on the orthodox Bayesian model on which degrees of belief are represented by a single probability measure, which appears wholly unable to accurately represent our credal states in a wide range of cases.[2] Alternatively, if one takes a binary relation ranking propositions in terms of their subjective expectedness as basic (Koopman, 1940; Stefánsson, 2017; Konek, 2019), known as the *comparativist* approach to credence, the claim would be that it is not appropriate to assume nor demand its completeness. See e.g. Insua (1992) and R. Bradley (2017, p. 236) for how such a relation can be represented by a credal set given that it satisfies certain conditions.

There are generally two main motivations for credal imprecision or incompleteness, which apply especially strongly in cases of deep uncertainty (compared to, say, standard cases of ambiguity with a well-defined space of possibilities and associated chances, such as those in Ellsberg, 1961). First of

[2] We set aside other potential improvements on the precise model, such as those involving additional second-order structure (see e.g. Skyrms, 1980; Gärdenfors and Sahlin, 1982; and Hill, 2019). Although, in agreement with e.g. Mogensen and Thorstad (2022, pp. 15-6), we do note that demanding precise second-order credences while recognising the arbitrariness of precise first-order credences borders on incoherence.

all, as already alluded to, it is often clear that precise probability assignments are arbitrary, and do not accurately represent the uncertainty of many agents.[3] This has long been recognised. Perhaps one of the earliest examples is Boole in *Laws of Thought*:

> Though our expectation of an event grows stronger with the increase of the ratio of the number of the known cases favourable to its occurrence to the whole number of equally possible cases, favourable or unfavourable, it would be unphilosophical to affirm that the strength of that expectation, viewed as an emotion of the mind, is capable of being referred to any numerical standard. (Boole, 1854/1940, p. 258)

This kind of arbitrariness is especially salient when we are conscious of our own unawareness of relevant possibilities, and have to form judgements about the likelihood of 'unknown unknowns' (Steele and Stefánsson, 2021, pp. 69-81). (More on this in §7.) Relaxing the requirement of precision then seems like a natural enrichment of our formal model of credence. Second, imprecision is arguably a more appropriate way for one's credences to reflect the available evidence in many situations. Following Joyce (2005), one might for instance argue that P should be in the credal set $\mathbf{P}$ for an agent S at a time t just in case P is compatible with the evidence S has at t. While it is non-trivial to explicate the 'compatibility with evidence' notion (see e.g. Joyce, 2005, 2010, White, 2009), it should be clear that in cases like Mogensen's, where unawareness and evidential poverty looms large, there is a broad range of probability measures compatible with the available evidence. Putting together the first and second claim, then, we agree with Greaves (2016, pp. 327-34) and Mogensen's (2021) use of imprecise probability to model the credences of the consequentialist in the context of cluelessness.[4]

[3] Drawing on Mahtani (2018; 2020), even if credences are defined or grounded dispositions to choose, the imprecise model is still an improvement on the precise one. For the situations which provide the motivation for imprecision in the first place will plausibly be situations in which choice is unstable across worlds.

[4] See Thorstad and Mogensen (2020, pp. 12-5) for a critique of credal imprecision as a model of cluelessness. Their main contention is that choice rules based on imprecise probabilities and imprecise expectations fall short in various ways (e.g. being too permissive). However, this objection rests on an overly pragmatic view of credence and choice, whereby the adequacy of

Our third claim is that subjective consequentialist ex ante betterness should be generalised to the imprecise setting in a 'supervaluationist' way (drawing an analogy to the semantics of vagueness): prospect a is at least as good as prospect b just in case the expected value of b is not less than that of a on each measure in $\mathbf{P}$. This means that the more imprecise the credences of the agent, the more incomplete the consequentialist 'at least as good as' relation will be.[5] This is the straightforward and natural generalisation of orthodox expected value theory. Indeed, suppose one takes the analogy to vagueness more seriously, and interprets each of the measures in the credal set as a 'sharpening' of the benefactor's credences. (This follows especially naturally on the comparativist approach to credence, as it is merely a matter of formal representation—again, see R. Bradley, 2017, p. 236, for more). And suppose one also accepts that ex ante betterness is defined in the orthodox way relative to a single probability measure in the credal set $\mathbf{P}$. Then, on the supervaluationist approach, it seems to straightforwardly follow that one option is determinately at least as good as another if and only if that is the case relative to all measures in $\mathbf{P}$.

3 Formal set-up

While the previous section introduced many of the definitions and concepts we need going forward, we have to be slightly more formal. We will use a framework of decision-making broadly in line with Bolker (1966; 1967) and Jeffrey's (1965/1983b) theory, as well as R. Bradley's (2017, pp. 232-8) imprecise generalisations thereof. See Broome (1990) for an application of Bolker and Jeffrey's theory in social choice, involving a formalism not too dissimilar to our own (which culminates in a version of Harsanyi's, 1955, utilitarian theorem).

First, let $W = \{w_1, w_2 ..., w_m\}$ be a finite set of possibilities which the agent is aware of and takes as basic, and let $\wp(W)$ be the power set of propositions on W. We write $\wp^*(W)$ for $\wp(W) - \varnothing$. As a way to model the benefactor's

a formal model of credence is taken to be sensitive to intuitions about choice rules based on that model.

[5] It is beyond the scope of the paper to discuss how such incompleteness relates to incommensurability (Griffin, 1986, Raz, 1986; Chang, 1997; Andersson and Herlitz, 2022), parity (Chang, 2002), vagueness (Broome, 1997), evaluative imprecision (Parfit, 2016), et cetera.

conscious unawareness, the algebra may also include a 'catchall proposition', roughly standing for all the possibilities the agent is unaware of. See §7 for more. Second, let $I = \{1, 2, ..., n\}$ be a set of beneficiaries, playing the role of bearers of local value. We assume it to be finite for simplicity. Since one might want to be able to bracket out certain episodes *within* the life of a beneficiary, elements in I should perhaps really be understood as sufficiently short 'person-episodes' so as to allow for such bracketing.[6] We can alternatively think of I as being a partition of some other dimension of the good, e.g. as a sufficiently fine partition of space-time. This latter construal may be more suitable for cases of deep uncertainty about the long-term future, where it often-times seems extremely difficult or impossible to identify individual beneficiaries. A related issue is that our actions plausibly affect who will be be born in the future (see e.g. Parfit, 1984, pp. 351-5). This was indeed a big part of Lenman's (2000) framing of cluelessness. So I should plausibly be indexed to a given w, if understood as a set of persons, or person-episodes. However, this may cause trouble for our theories of bracketing, since they rely on ex ante comparisons relative to individual beneficiaries, or subsets thereof. About this kind of worry, Blackorby et al. (2007, p. 569) go as far as to write that "individual ex-ante assessments of prospects are meaningless if the person is not alive in all possible states".[7] On the other hand, Meacham (2012) defines a counterpart relation across worlds which can let us relate beneficiaries based on their role in moral reasoning rather than only based on qualitative similarity; and although this may temper or even defuse the worry expressed by Blackorby et al., again, construing I as a set of non-overlapping space-time regions might be more appropriate for our purposes.[8] But nonetheless, we will use the term "beneficiaries" going forward.

We assume that ex post goodness relative to each beneficiary admits a real-valued, ratio-scaled representation; i.e., each beneficiary $i \in I$ is associated with

[6] Thanks here to Lukas Finnveden.

[7] If one is able to successfully argue that the value relative to a non-existing individual is zero, then this might instead damage the degree to which theories of bracketing are action-guiding, compared to consequentialism. To see why this might be, consider the case in Table 2 which we will get to eventually.

[8] Space-time regions as bearers or containers of local value are commonly employed, as well as independently motivated, in infinitary ethics. See e.g. Vallentyne and Kagan (1997), and in particular Wilkinson (2020; 2021).

a ratio-scaled value function $u_i : W \to \mathbb{R}$. We also assume full comparability of personal or local value (cf. Sen, 1970/2017, p. 160). While we could formulate versions of bracketing without these assumptions, they make aggregation particularly straightforward, simplifying matters. This is also a good place to note that deep uncertainty not only puts pressure on the idea of precise probability as a good model of credence, but also on the assumption made here that we can know or precisely estimate the value for beneficiaries of outcomes we conceive of. However, incorporating additional imprecision along these lines is beyond the scope of this paper.

Now, let $U_{I'}(w) \triangleq \sum_{i \in I'} u_i(w)$ represent the 'total goodness' relative to a non-empty subset $I' \subseteq I$ of beneficiaries at a given $w \in W$. It is therefore assumed that goodness is additively separable across beneficiaries. Following our discussion in the previous section, the credences of the benefactor are represented by a set of probability measures $\mathbf{P}$, with a generic element $P : \wp(W) \to [0, 1]$. For a given proposition $e \in \wp(W)$, we will use the abbreviations $\mathbf{P}(e) \triangleq \{P(e) \mid P \in \mathbf{P}\}$ and $\mathbf{P}(e \mid a) \triangleq \{P(e \mid a) \mid P \in \mathbf{P}\}$. For a given probability measure P and any non-empty subset $I' \subseteq I$, define a signed measure $V^P_{I'} : \wp^*(W) \to \mathbb{R}$ as follows, provided $P(a) > 0$.

$$V^P_{I'}(a) \triangleq \sum_{w \in a} P(\{w\} \mid a) U_{I'}(w) \tag{1}$$

In line with the third claim from the previous section, we then define a relation $\preccurlyeq_{I'}$ on $\wp^*(W)$ representing subjective consequentialist ex ante betterness generalised to the imprecise setting relative to a non-empty subset $I' \subseteq I$ of beneficiaries as follows.

Definition 3.1 (Imprecise consequentialist ex ante betterness)**.** *For any $I' \subseteq I$, define a relation $\preccurlyeq_{I'}$ on $\wp^*(W)$, with asymmetric and symmetric parts $\prec_{I'}$ and $\sim_{I'}$, representing* imprecise consequentialist ex ante betterness *as follows.*

$$a \prec_{I'} b \iff \left(\forall P \in \mathbf{P}, V^P_{I'}(a) \leq V^P_{I'}(b)\right) \wedge \left(\exists P \in \mathbf{P}, V^P_{I'}(a) < V^P_{I'}(b)\right) \tag{2}$$

$$a \sim_{I'} b \iff \left(\forall P \in \mathbf{P}, V^P_{I'}(a) = V^P_{I'}(b)\right) \tag{3}$$

Equivalently, $a \preccurlyeq_{I'} b$ just in case $V^P_{I'}(a) \leq V^P_{I'}(b)$, for all probability measures P

in the credal set $\mathbf{P}$. We will henceforth abbreviate 'imprecise consequentialist ex ante betterness' as 'consequentialist betterness' and 'imprecise ex ante consequentialism' as 'consequentialism', and use phrases such as "b is *c-better* than a relative to I'" if $a \prec_{I'} b$, "b and a are equally *c-good*" if $a \sim_{I'} b$, et cetera. Note that $\preccurlyeq_{I'}$ is a preorder (reflexive and transitive). We will henceforth write $a \bowtie_{I'} b$ if and only if either $a \prec_{I'} b$, $b \prec_{I'} a$ or $a \sim_{I'} b$. Otherwise we write $a \not\bowtie_{I'} b$, and we say that it is *indeterminate* which prospect is at least as c-good as the other relative to I'. In other words, $a \not\bowtie_{I'} b$ if and only if there are probability measures $P, P' \in \mathbf{P}$ such that $V_{I'}^{P}(a) < V_{I'}^{P}(b)$ and $V_{I'}^{P'}(a) > V_{I'}^{P'}(b)$.

Moving on to choice, let a *betterness-based choice rule* be a partial function $C_{\trianglelefteq} : \wp^*(\wp^*(W)) \rightharpoonup \wp(\wp^*(W))$ which picks out a subset of permissible prospects from each non-empty menu of options in $\wp^*(W)$ (and only those menus) as a function of a binary betterness relation $\trianglelefteq$ on $\wp^*(W)$. We do not take a stance on when a proposition counts as an option, but see e.g. Jeffrey (1965/1983b, p. 84) for a popular proposal. For preorders, such as $\preccurlyeq_{I'}$ defined above, the appropriate consequentialist rule is arguably the following (Sen 1997, 2004; R. Bradley, 2017, pp. 160-1).

Definition 3.2 (Maximality)**.** *The set of permissible options from a menu A based on a betterness relation $\trianglelefteq$ according to* Maximality *is given by*

$$C_{\trianglelefteq}^{\mathtt{max}}(A) \triangleq \{a \in A \mid \neg\exists b \in A, a \triangleleft b\}. \tag{4}$$

In words: if there is no other available option which is ranked above a, then a is deemed permissible. To argue for something less permissive would be to be argue for a discrepancy between the goodness of prospects and permissible choice, which seems uncalled for, assuming consequentialism. However, note that when the relation in question is intransitive (and not just incomplete), we will have to look at alternative choice rules and refinements of Maximality (cf. Schwartz, 1972; R. Bradley, 2015). We return to this issue in §4.

	dep	$\neg dep$
AMF	$(10, 0, -100)$	$(10, 0, 50)$
$MAWF$	$(0, 1, -100)$	$(0, 1, 50)$

Table 1: The donation case

Finally, let us consider a stylised formalisation of Mogensen's (2021) donation case to get a feel for the framework. The choice problem is represented in Table 1, where dep stands for the event-proposition *that resources will be depleted*. Each pair of event- and act-propositions uniquely determines an element in W. There are three beneficiaries: the child you could save from malaria by donating to AMF, the child you could grant a wish to by donating to MAWF, and a beneficiary representing future generations (or perhaps the future itself, if indices represent suitable space-time regions). Suppose $\mathbf{P}(dep \,|\, AMF) = [0.3, 0.8]$ and $\mathbf{P}(dep \mid MAWF) = \mathbf{P}(dep) = 0.3$. We then have $V^P_{\{1,2,3\}}(MAWF) = 6$, for all $P \in \mathbf{P}$, and $\{V^P_{\{1,2,3\}}(AMF) \,|\, P \in \mathbf{P}\} = [-60, 15]$, meaning $AMF \not\bowtie_{\{1,2,3\}} MAWF$. And so $C^{\max}_{\preccurlyeq_{\{1,2,3\}}}(\{AMF, MAWF\}) = \{AMF, MAWF\}$, i.e., both options are permissible. While the specific numbers chosen here are largely unimportant, the structural features of this problem to notice are: $MAWF$ is determinately c-better for 2, AMF is determinately c-better for 1 *and* the group $\{1, 2\}$, but neither option is at least as c-good as the other for 3. The last-mentioned indeterminacy moreover swamps the calculation of total value, making it indeterminate which prospect is at least as c-good overall. Intuitively, many cases of cluelessness share a similar structure, whereby an 'easy' comparison is made 'hard' by taking into account deeply uncertain future effects.

4 On Super-Strong Pareto

Before precisely defining the two theories of bracketing, we set the stage by commenting on some relevant literature. In particular, we will look at 'Super-Strong Pareto' as a property of an overall ranking. That this sort of property should be satisfied has been rejected by multiple authors in various different contexts. But interestingly, bottom-up bracketing satisfies the property, and the arguments that can be mounted against Super-Strong Pareto also apply to top-down bracketing (even though the theory does not generally satisfy it). We respond to all but one of these arguments in this section.

Without further ado, consider the subtle and seldom recognised distinction between Strong Pareto and Super-Strong Pareto, using Hedden (2024), and Hedden and Muñoz' (2024, pp. 299-301) terminology (although stated in terms of abstract rankings and not betterness, for reasons that will become clear).

Definition 4.1 (Strong Pareto, SP)**.** *A ranking of prospects* $\trianglelefteq_I$*, with symmetric and asymmetric parts* $\approx_I$ *and* $\triangleleft_I$*, satisfies* Strong Pareto *just in case*

$$(\exists i \in I : a \triangleleft_{\{i\}} b) \wedge (\forall i \in I : a \trianglelefteq_{\{i\}} b) \implies (a \triangleleft_I b).$$

In words: if b is ranked above or co-ranked with a relative to every beneficiary, and there is some beneficiary for whom b is ranked above a, then b will also be ranked above a relative to the set of all beneficiaries. That e.g. an overall betterness ranking should satisfy Strong Pareto is widely accepted, at least conditional on overall value being reducible to personal value. But contrast Strong Pareto to the stronger property, Super-Strong Pareto:

Definition 4.2 (Super-Strong Pareto, SSP)**.** *A ranking of prospects* $\trianglelefteq_I$*, with symmetric and asymmetric parts* $\approx_I$ *and* $\triangleleft_I$*, satisfies* Super-Strong Pareto *just in case*

$$(\exists i \in I : a \triangleleft_{\{i\}} b) \wedge (\neg\exists i \in I : b \triangleleft_{\{i\}} a) \implies (a \triangleleft_I b).$$

That is, if b is *not ranked below* a for every beneficiary, and there is some beneficiary for whom b is strictly ranked above a, then b is also ranked above a overall. When $\trianglelefteq_{\{i\}}$ is complete for every $i \in I$, we see that $\trianglelefteq_I$ satisfies SP if and only if $\trianglelefteq_I$ satisfies SSP. But suppose, for instance, that we have a case with two beneficiaries $I = \{1, 2\}$ such that $a \triangleleft_{\{1\}} b$ but neither $a \trianglelefteq_{\{2\}} b$ nor $b \trianglelefteq_{\{2\}} a$. Since b is not ranked below a for the second beneficiary but ranked above a for the first, a ranking which satisfies SSP would rank b above a, whereas a ranking which only satisfies SP need not.

At a first pass, it thus appears as if a theory of bracketing would satisfy SSP, since it seemingly amounts to disregarding indeterminacy. Although this is only true of bottom-up bracketing, as mentioned, the objections to SSP which we will consider are applicable to both kinds of bracketing. Note also that while we construe SP and SSP as properties of an overall ranking, they can of course be applied more widely—see Hedden (2024) for more. Also see Steele (2022) for an examination of more general principles similar to SSP, which concern any underlying set of criteria. Steele also points out the connection to cluelessness.

We will consider four arguments against the claim that SSP should be satisfied, three of them being based on the fact that a ranking satisfying SSP is not generally acyclic (cf. Hedden, 2024, p. 578).[9] They are:

(i) The semantic argument from cyclicity.

(ii) The synchronic choice argument from cyclicity.

(iii) The value-pump argument from cyclicity.

(iv) The argument from 'reasons-weighing'.

While we wait until § 6 to give a specific case where SSP and theories of bracketing lead to a cycle, we respond to the semantic argument, the synchronic choice argument and the argument from reasons-weighing here, and only take up also responding to the value-pump argument in §6.

First, the semantic argument states that a binary relation representing an axiology is generally acyclic, or even transitive, in virtue of the meaning of 'better than', and so any ranking which satisfies SSP does not represent an axiology. Broome (1991, pp. 11-2; 2004, pp. 50-63) for instance argues that the 'better than' relation is necessarily transitive. While e.g. Temkin (1987; 2012) and Rachels (1998) famously argue against this, we do not further discuss whether the semantic argument is convincing. Instead we simply note that we can understand our theories of bracketing in terms of rightness or choice-worthiness. In other words, if $a \lhd b$ according to a theory of bracketing, this could be understood as the claim that in a choice between a and b, b should be chosen; and otherwise, both a and b are permissible. In choice problems with more than two available alternatives, such a ranking should moreover only be understood as a tentative claim about choiceworthiness, one which could be overturned in light of how the rest of the alternatives in the menu are ranked (by Maximality or any suitable refinement of Maximality). In any case, to stay neutral on this issue, we will henceforth refer to the kind of binary relation representing theories of bracketing as a "ranking" on $\wp^*(W)$. Indeed, we have intentionally used such language in the previous paragraphs of this section.

[9] In varying contexts, going under different names, Super-Strong Pareto is also defined and pointed out to lead to cyclicity in e.g. Parfit (2011, p. 224), Temkin (2012, p. 429), Hare (2013, p. 176), MacAskill (2013), Askell (2018, pp. 244-6) and Nebel (2019).

Second, the synchronic choice argument states that when rankings are cyclic, it can be impossible to choose. In the context of the repugnant conclusion, Arrhenius et al. (2024) for instance write:

> The main problem with non-transitive value orderings in moral theory, however, is that such an ordering cannot form the basis of a satisfactory answer to the question of what one ought to choose. (*ibid*, pp. 17-8)

To see the issue, suppose we have three available options $\{a, b, c\}$ such that $a \lhd b \lhd c \lhd a$. Standard choice rules like Maximality would yield an empty set of permissible options, since there is no alternative which is not ranked below another. And so we are left with no guidance regarding what to do. However, this argument is unconvincing, as it only shows that we have to refine our choice rule. Indeed, Maximality is itself a refinement of *Optimality*, which deems an alternative permissible just in case it is ranked above or co-ranked with every other alternative, and is not suitable for mere preorders (R. Bradley, 2017, pp. 160-1). There is a large literature on choice rules appropriate for potentially cyclic rankings (see e.g. Schwartz 1972, 1990; Fishburn, 1977; Miller, 1980; Brandt et al., 2016), and we consider just one such refinement here (stated as a *ranking-based* rule, not a *betterness-based* one).

Definition 4.3 (Uncoveredness; Miller, 1980, pp. 72-4, Gustafsson, 2022, p. 6)**.** *The set of permissible options from a menu A based on a ranking $\unlhd$ according to* Uncoveredness *is given by*

$$C^{\text{unc}}_{\unlhd}(A) \triangleq \{a \in A \mid \neg\exists b \in A([a \lhd b] \wedge [\forall c \in A, (c \lhd a) \rightarrow (c \lhd b)])\}\,. \tag{5}$$

The idea is that an alternative a is impermissible whenever there is some $b \in A$ such that $a \lhd b$ and which also replicates all of a's advantages, in the sense that any c which is outranked by a must also be outranked by b. Intuitively, if every alternative ranked below a is also ranked below b, then b 'covers' a, making a impermissible. An alternative is then deemed permissible, on the other hand, precisely by not being covered in this sense. (It is straightforward to show that Uncoveredness coincides with Maximality whenever $\unlhd$ is a preorder.) So, if we for instance face a menu $A = \{a, b, c\}$ such that $a \lhd b \lhd c \lhd a$, every option is permissible. Moreover, if we are faced with the larger menu $A' = \{a, b, c, a^-\}$, where a^- is a 'mild souring' of a (meaning a prospect which is just slightly worse than a for everyone), the set of permissible options remains $\{a, b, c\}$ according to Uncoveredness, since a covers a^-. This fact resurfaces in §6.

Third, the value-pump argument roughly states that if a ranking of prospects is not generally acyclic, then that ranking can be turned into a value-pump in some sequential or dynamic decision problems (cf. Davidson et al., 1955, p. 145; Gustafsson, 2022, pp. 3-24), where this is a taken to be problematic. While more would have to be said in the direction of why exactly this is problematic, we simply note that we show how a natural generalisation of bracketing avoids value-pumps in §6.

Finally, we have the argument from 'reasons-weighing'. Upon noting that Super-Strong Pareto leads to cycles, Hedden investigates where satisfying Super-Strong Pareto supposedly goes wrong, writing the following (where we can read "a and b being on par" as "$a \not\bowtie_{I'} b$" in the present context).

> Super-Strong Pareto does not follow from the weighing model of reasons. For when [a] is on a par with [b] for i, the reasons having to do with i's welfare for choosing [a] are not exactly equally counterbalanced by the reasons having to do with i's welfare for choosing [b]. And so when we add in the reasons having to do with j's welfare, where [b] is better than [a] for j, those reasons don't necessarily tip the scales and break the tie. For there was no tie to be broken. (Hedden, 2024, p. 589)

This argument is however question-begging. There is no single 'weighing model of reasons' to fall back on. For instance, as we will see, on bottom-up bracketing, if it is indeterminate which prospect is at least as c-good for i, then i simply provides no reason for or against a or b. *That is the proposal*. Steele (2022, pp. 236-8) considers a similar argument in the context of the aforementioned more general principles, and also concludes that such an argument is indeed question-begging—at least in absence of "an alternative and even more compelling principle" (Steele, 2022, pp. 237-8). Nevertheless, this does not relieve us of the burden of justifying theories of bracketing. We still need to argue why bracketing constitutes a compelling way to construe betterness or choiceworthiness. And, in short, our argument is simply that theories of bracketing, especially top-down bracketing, handle cases of cluelessness in a satisfactory way, consistent with common-sense intuition.

5 Bracketing

We are now ready to precisely define the two aforementioned theories of bracketing. The aim is to capture the intuition that one e.g. ought to donate to AMF in the donation case, while still respecting determinate c-betterness. To make the latter notion more precise, consider the following definition.

Definition 5.1 (Extension; cf. R. Bradley, 2017, p. 235)**.** *A relation $R' \subseteq X \times X$ is an* extension *of $R \subseteq X \times X$ if and only if the following two conditions are satisfied:*

(i) $(x, y), (y, x) \in R$ *implies* $(x, y), (y, x) \in R'$*; and*

(ii) $(x, y) \in R$ *and* $(y, x) \notin R$ *implies* $(x, y) \in R'$ *and* $(y, x) \notin R'$*.*

In other words, we want the ranking of prospects representing a theory of bracketing to be an extension of $\preccurlyeq_I$. Or, put differently, a theory of bracketing should agree with consequentialism regarding two prospects a and b whenever $a \bowtie_I b$. As we will see, only top-down bracketing satisfies this condition, and hence we favour it over the bottom-up approach.

5.1 Bottom-up bracketing

The first theory of bracketing we consider, *bottom-up bracketing*, consists in ranking any two options based on the c-betterness relative to the subset of beneficiaries for whom c-betterness is respectively determinate. In other words, for any two prospects, we set aside those beneficiaries for whom we are clueless about which prospect is at least as c-good for them. ("Bottom-up" because we bracket on the level of individual beneficiaries after which we aggregate to reach an overall verdict.) Formally:

Definition 5.2 (Bottom-up bracketing)**.** *For $\mathfrak{B}(I, \{a, b\}) \triangleq \{i \in I \mid a \bowtie_{\{i\}} b\} \subseteq I$, define a relation $\preccurlyeq_I^{\mathsf{BU}}$ on $\wp^*(W)$, with asymmetric and symmetric parts $\prec_I^{\mathsf{BU}}$ and $\sim_I^{\mathsf{BU}}$,* *representing* bottom-up bracketing *as follows.*

$$a \prec_I^{\mathsf{BU}} b \iff a \prec_{\mathfrak{B}(I,\{a,b\})} b \tag{6}$$
$$a \sim_I^{\mathsf{BU}} b \iff a \sim_{\mathfrak{B}(I,\{a,b\})} b \tag{7}$$

If $\mathfrak{B}(I,\{a,b\}) = \varnothing$*, then* $a \not\bowtie_I^{\mathsf{BU}} b$*.*

Relative to individual beneficiaries, we moreover let bottom-up bracketing coincide with consequentialism, i.e., $\preccurlyeq_{\{i\}}^{\mathsf{BU}} := \preccurlyeq_{\{i\}}$, for all $i \in I$. Note then that bottom-up bracketing straightforwardly satisfies Super-Strong Pareto.

Proposition 5.3. $\preccurlyeq_I^{\mathsf{BU}}$ *satisfies Super-Strong Pareto.*

Proof. Suppose there is an $i \in I$ such that $a \prec_{\{i\}}^{\mathsf{BU}} b$ and no $i \in I$ such that $b \prec_{\{i\}}^{\mathsf{BU}} a$. This means that there is an $i \in I$ such that $a \prec_{\{i\}} b$ and no $i \in I$ such that $b \prec_{\{i\}} a$. Then for all $i \in \mathfrak{B}(I,\{a,b\})$, it is either the case that $a \prec_{\{i\}} b$ or $a \sim_{\{i\}} b$, which entails that $a \prec_{\mathfrak{B}(I,\{a,b\})} b$. In turn, by (6), $a \prec_I^{\mathsf{BU}} b$. □

	e	$\neg e$
a	$(1,0)$	$(0,0)$
b	$(0,0)$	$(0,1)$

Table 2: Determinacy is not upwards hereditary

Also note that even if $\mathfrak{B}(I,\{a,b\})$ is non-empty, it can be the case that $a \not\bowtie_I^{\mathsf{BU}} b$. This is because determinate c-betterness is not generally 'upwards hereditary', in that even if it is determinate which prospect is at least as c-good for each beneficiary, it can still be indeterminate which option is at least as c-good relative to the whole set of beneficiaries. Consider the case illustrated in Table 2, for instance, where we have two alternatives $\{a,b\}$, beneficiaries $\{1,2\}$ and $\mathbf{P}(e \mid a) = \mathbf{P}(e \mid b) = \mathbf{P}(e) = [0,1]$. Note that $b \prec_{\{1\}} a$ and $a \prec_{\{2\}} b$, meaning $\mathfrak{B}(\{1,2\},\{a,b\}) = \{1,2\}$. Yet it is the case that $a \not\bowtie_{\{1,2\}}^{\mathsf{BU}} b$, since $a \not\bowtie_{\{1,2\}} b$.

Let us now apply bottom-up bracketing to the donation case (Table 1). Recall that we had three beneficiaries such that

$$MAWF \prec_{\{1\}} AMF,$$
$$AMF \prec_{\{2\}} MAWF, \text{and}$$
$$AMF \not\prec_{\{3\}} MAWF,$$

meaning $\mathfrak{B}(\{1,2,3\},\{AMF, MAWF\}) = \{1,2\}$. Relative to $\{1,2\}$, AMF is c-better than $MAWF$, and so $MAWF \prec^{\mathsf{BU}}_{\{1,2,3\}} AMF$. In other words, bottom-up bracketing achieves the intuitive verdict. The theory correctly brackets out the index representing future beneficiaries, which was the source of the indeterminacy.

While arguably prima facie compelling, bottom-up bracketing is *not* an extension of consequentialism. We can see this by noting that bottom-up bracketing violates statewise dominance (which is not true of consequentialism, of course). Take, for example, the choice problem represented in Table 3, where we have two options a and b, three beneficiaries $\{1,2,3\}$, and a credal set $\mathbf{P}$ such that $\mathbf{P}(e \mid a) = \mathbf{P}(e \mid b) = \mathbf{P}(e) = [0,1]$; meaning e and $\neg e$ can really be thought of as (non-atomic) states.[10] First notice that b is statewise dominant. That is, conditional on both e and $\neg e$ respectively, b is c-better than a. Hence $a \prec_{\{1,2,3\}} b$. However, neither 2 nor 3 are in $\mathfrak{B}(\{1,2,3\},\{a,b\})$. Because, for instance, relative to the measure $P \in \mathbf{P}$ such that $P(e) = 1$, a would be c-best for the second beneficiary and b would c-best for the third, whereas relative to the measure $P' \in \mathbf{P}$ such that $P'(e) = 0$, for instance, b would be c-best for the second beneficiary and a would c-best for the third. So it is indeterminate what is c-best for each of 2 and 3. On the other hand, a is c-best for 1 on each measure in $\mathbf{P}$, and so $\mathfrak{B}(\{1,2,3\},\{a,b\}) = \{1\}$. As a result, $b \prec^{\mathsf{BU}}_{\{1,2,3\}} a$, despite the fact that $a \prec_{\{1,2,3\}} b$.[11]

[10]See Kollin (ms) for discussion of an analogous case in the context of multidimensional welfare, where it is used to argue against the 'fixed-population person-affecting restriction'.

[11]This observation suggests a two-step approach, drawing on Herlitz (2018; 2022). Namely, that bottom-up bracketing should be defined as: "if c-betterness is determinate, then the theory agrees with consequentialism, otherwise bracket out the beneficiaries relative to whom c-betterness is indeterminate". However, this procedure appears quite inelegant, and more importantly, does not solve the underlying problem. For we can write down a slight modification of the case in Table 3, where we add a fourth beneficiary relative to whom c-betterness is indeterminate and who also makes overall c-betterness indeterminate. Here the bracketing

	e	$\neg e$
a	$(1, 10, 0)$	$(1, 0, 10)$
b	$(0, 0, 20)$	$(0, 20, 0)$

Table 3: A case of statewise dominance

Note that this case is also a novel counterexample to the claim that Super-Strong Pareto should be satisfied, since a is ranked above b relative to 1, whereas a is not ranked below b relative to 2 and 3, and so *any* ranking satisfying Super-Strong Pareto would then also rank a over b overall—again, contrary to statewise dominance.

The issue is that it can be indeterminate which option is at least as c-good for every beneficiary in a set $I' \subseteq I$, but simultaneously *determinate* which option is at least as c-good relative to the set as a whole. This is for instance true of $\{2, 3\}$ in the case above. The observation that indeterminate c-betterness is not generally upwards hereditary in this sense is important and motivates our second theory of bracketing. (Recall that determinate c-betterness is also not upwards hereditary.)

5.2 Top-down bracketing

We now consider *top-down bracketing*, which we define relative to any non-empty $I' \subseteq I$ (this will be convenient later). The basic idea is that we look for the largest possible subsets of I' relative to which c-betterness is determinate, and then base our overall comparison on what is c-best relative to those subsets. ("Top-down" because we bracket on the level of subsets of beneficiaries.) To formally state the theory, we first define a *bracket-set*.

Definition 5.4 (Bracket-set)**.** *A* bracket-set *relative to a non-empty $I' \subseteq I$ and a pair of prospects $\{a, b\}$ is a non-empty subset $I'' \subseteq I'$ such that $a \bowtie_{I''} b$.*

We then define what it means for a bracket-set to be *maximal*. Informally, a bracket-set is maximal just in case we cannot enlarge it in any way without it no longer being a bracket-set.

intuition arguably favours b, since b is c-better for $\{2, 3\}$, yet the two-step procedure will still rank a above b.

Definition 5.5 (Maximal bracket-set)**.** *A bracket-set I'' relative to I' and $\{a, b\}$ is* maximal *just in case there is no $I''' \subseteq I'$ which is (i) a bracket-set relative to I' and $\{a, b\}$, and (ii) $I'' \subset I'''$.*

Top-down bracketing is then formally defined as follows.

Definition 5.6 (Top-down bracketing)**.** *For any non-empty $I' \subseteq I$, define a relation $\preccurlyeq_{I'}^{\mathsf{TD}}$ on $\wp^*(W)$, with asymmetric and symmetric parts $\prec_{I'}^{\mathsf{TD}}$ and $\sim_{I'}^{\mathsf{TD}}$, representing* top-down bracketing *as follows, where $\texttt{MB} = \{I''_1, \ldots, I''_k\}$ is the set of maximal bracket-sets relative to I' and $\{a, b\}$.*

$$a \prec_{I'}^{\mathsf{TD}} b \iff (\forall I'' \in \texttt{MB}, a \preccurlyeq_{I''} b) \wedge (\exists I'' \in \texttt{MB}, a \prec_{I''} b) \tag{8}$$

$$a \sim_{I'}^{\mathsf{TD}} b \iff (\forall I'' \in \texttt{MB}, a \sim_{I''} b) \tag{9}$$

If $\texttt{MB} = \varnothing$, then $a \not\bowtie_{I'}^{\mathsf{TD}} b$.

In other words, for each pair of prospects, top-down bracketing first asks us to find the set of maximal bracket-sets. Then, if b is at least as c-good as a relative to every maximal bracket-set and is c-better relative to at least one maximal bracket-set, then b is ranked above a by top-down bracketing; and if a and b are equally c-good relative to every maximal bracket-set, then a and b are co-ranked by top-down bracketing. There may be other reasonable ways of aggregating the verdicts we get from the set of maximal bracket-sets. Exploration of such alternatives would take us too far astray however, and in any case, we hold that an alternative approach should at least agree with the determinate comparisons top-down bracketing, as defined here, makes.

Note that $\preccurlyeq_{I'}^{\mathsf{TD}}$ is not necessarily complete, even if $\texttt{MB} \neq \varnothing$. For there are cases in which there are multiple maximal bracket-sets, but which yield conflicting rankings. For instance, recall the case represented in Table 2, where we had two prospects a and b, a set of beneficiaries $I = \{1, 2\}$ and $b \prec_{\{1\}} a$, $a \prec_{\{2\}} b$ and $a \not\bowtie_{\{1,2\}} b$. In other words, there are two maximal bracket-sets, $\{1\}$ and $\{2\}$, but which are in disagreement, and so $a \not\bowtie_{\{1,2\}}^{\mathsf{TD}} b$.

It is also easy to show that top-down bracketing is indeed an extension of consequentialism (contrary to bottom-up bracketing).

Proposition 5.7. *For any non-empty $I' \subseteq I$, $\preccurlyeq_{I'}^{\mathsf{TD}}$ is an extension of $\preccurlyeq_{I'}$.*

Proof. If $a \prec_{I'} b$, then there is a unique maximal bracket-set relative to I' and $\{a, b\}$, namely I'. And so by (8), $a \prec^{\mathsf{TD}}_{I'} b$. If $a \sim_{I'} b$, then I' is the unique maximal bracket-set again. By (9), $a \sim^{\mathsf{TD}}_{I'} b$, and we are done. □

We now want to make sure top-down bracketing satisfactorily deals with the donation case (Table 1). Recall that there are three non-empty subsets which are bracket-sets (i.e., sets relative to which c-betterness is determinate): $\{1\}, \{2\}$ and $\{1, 2\}$. Among these, $\{1, 2\}$ is uniquely maximal, and so $MAWF \prec^{\mathsf{TD}}_{\{1,2,3\}} AMF$, since $MAWF \prec_{\{1,2\}} AMF$. In other words, just as bottom-up bracketing, top-down bracketing correctly brackets out the index representing future generations. Yet, since top-down bracketing is an extension of consequentialism, it does not violate statewise dominance like bottom-up bracketing. In particular, recall from our case of statewise dominance (Table 3) that $a \prec_{\{1,2,3\}} b$, meaning $\{1, 2, 3\}$ is the unique maximal bracket-set, and so $a \prec^{\mathsf{TD}}_{\{1,2,3\}} b$.

Before moving on to further discussion of intransitivity, we forestall a potential objection to top-down bracketing. (Or in any case, point out an interesting feature of the theory.) Consider the following property which states that whenever a prospect b is co-ranked with or ranked above a prospect a relative to two non-overlapping sets of beneficiaries, b is also co-ranked with or ranked above a relative to the merged set of beneficiaries.

Definition 5.8 (Monotonicity). *A family of rankings* $\{\trianglelefteq_{I'} \mid I' \in \wp^*(I)\}$ *satisfies* Monotonicity *if and only if*

$$(I' \cap I'' = \varnothing) \wedge (a \trianglelefteq_{I'} b) \wedge (a \trianglelefteq_{I''} b) \implies a \trianglelefteq_{I' \cup I''} b.$$

Monotonicity is trivially satisfied by consequentialism. It also follows from $\mathfrak{B}(I', \{a, b\}) \cup \mathfrak{B}(I'', \{a, b\}) = \mathfrak{B}(I' \cup I'', \{a, b\})$ that bottom-up bracketing generalised to any non-empty $I' \subseteq I$ satisfies Monotonicity. But while prima facie plausible, top-down bracketing does not satisfy Monotonicity. To see this, consider a slight variant of the previous case of statewise dominance, illustrated in Table 4, where we have simply duplicated the first beneficiary. As before, $\mathbf{P}(e \mid a) = \mathbf{P}(e \mid b) = \mathbf{P}(e) = [0, 1]$. Now compare a and b relative to $\{1, 2\}$ and $\{3, 4\}$, respectively. Relative to the first group, $\{1\}$ is the unique maximal bracket-set, meaning $b \prec^{\mathsf{TD}}_{\{1,2\}} a$. Relative the second group, $\{4\}$ is the unique maximal bracket-set, and so $b \prec^{\mathsf{TD}}_{\{3,4\}} a$. Yet, $a \prec^{\mathsf{TD}}_{\{1,2,3,4\}} b$, since b is determinately c-better than a relative to the set of all beneficiaries. As before, the issue is of course that indeterminate c-betterness is not upwards hereditary; whenever they are in the same set, 2 and 3 together create a strong reason in favour of b.[12] So, while Monotonicity may have appeared plausible, there is natural and justified reason for why the property is not satisfied by top-down bracketing.[13]

	e	$\neg e$
a	$(1, 10, 0, 1)$	$(1, 0, 10, 1)$
b	$(0, 0, 20, 0)$	$(0, 20, 0, 0)$

Table 4: Another case of statewise dominance

6 Dynamic choice and cyclicity

Recall the objections from cyclicity to rankings satisfying SSP (such as $\preccurlyeq^{\mathsf{BU}}_{I}$) from § 4: the semantic argument, the synchronic choice argument and the value-pump argument. As mentioned, while top-down bracketing does not satisfy SSP, $\preccurlyeq^{\mathsf{TD}}_{I}$ is also not generally acyclic. And so these arguments not only apply to bottom-up bracketing, but also top-down bracketing. We already responded to the first two arguments, but not the third. This argument states that if a betterness or choiceworthiness ranking is not generally acyclic, then it can be turned into a value-pump in some dynamic decision problems; i.e., assuming some way of choosing dynamically based on a ranking, it is recommended that we trade a prospect for a strictly worse one, even though it is possible to just stick with the initial prospect at no cost. We respond to this argument in this section.

Here is the plan. We first provide a case where both bottom-up and top-down bracketing lead to a cycle and then spell out the value-pump argument in more detail (§ 6.1), provide some background on dynamic choice when the

[12] If we construe top-down bracketing as an axiology, one might think of some subsets as Moorean 'organic unities' in that "the value of such a whole bears no regular proportion to the sum of the values of its parts" (Moore, 1903, p. 30).

[13] We thank Caspar Oesterheld for discussion that led to this paragraph.

underlying ranking is a mere preorder (§ 6.2), and then respond to the aforementioned argument by showing how a natural way of generalising bracketing to the dynamic setting avoids value-pumps (§ 6.3). Throughout this section, our focus will be on top-down bracketing as it is the theory we favour.

6.1 The value-pump argument

Consider the prospects $\{a, b, c\}$ represented in Table 5 (loosely adapted from MacAskill, 2013, pp. 512-3, in a different context), where we have a set of beneficiaries $I = \{1, 2\}$ and once again a credal set $\mathbf{P}$ such that $\mathbf{P}(e \mid x) = \mathbf{P}(e) = [0, 1]$, for all $x \in \{a, b, c\}$. To apply bottom-up bracketing, we first need to specify the sets, $\mathfrak{B}(\{1, 2\}, \cdot)$. They are:

$$\mathfrak{B}(\{1,2\},\{a,b\}) = \{1\},$$
$$\mathfrak{B}(\{1,2\},\{a,c\}) = \{1,2\}, \text{ and}$$
$$\mathfrak{B}(\{1,2\},\{c,b\}) = \{1\}.$$

	e	$\neg e$
a	$(2,2)$	$(2,10)$
b	$(1,20)$	$(1,2)$
c	$(0,10)$	$(0,20)$

Table 5: Bottom-up and top-down bracketing both produce cycles

Correspondingly, we have the ranking

$$b \prec^{\mathsf{BU}}_{\{1,2\}} a,$$
$$a \prec^{\mathsf{BU}}_{\{1,2\}} c, \text{ and}$$
$$c \prec^{\mathsf{BU}}_{\{1,2\}} b,$$

since $b \prec_{\{1\}} a$, $a \prec_{\{1,2\}} c$ and $c \prec_{\{1\}} b$. In other words, a cycle. It is easy to see that top-down bracketing leads to the same cycle, since the maximal bracket-sets relative to each pair of alternatives coincides with the $\mathfrak{B}(\{1,2\}, \cdot)$

sets above, meaning we also have

$$b \prec^{\mathsf{TD}}_{\{1,2\}} a,$$
$$a \prec^{\mathsf{TD}}_{\{1,2\}} c, \text{ and}$$
$$c \prec^{\mathsf{TD}}_{\{1,2\}} b.$$

Notice that the source of the cyclicity in both cases is that different sets of beneficiaries are bracketed out for different comparisons of prospects.[14] If we understand both kinds of bracketing as axiologies, this appears to be an instance of Temkin's (2012) idea that betterness can be intransitive because betterness is 'essentially comparative' meaning "one must directly compare two alternatives in order to determine their relative ranking" (Temkin, 2012, p. 304). In other words, different comparisons may bring different aspects into focus, resulting in intransitivity.

To show how cycles like these lead to alleged trouble in dynamic or sequential contexts, consider the choice problem illustrated by the means of a decision tree in Figure 1—the 'upfront value-pump' (Gustafsson, 2022, p. 12). The prospects a, b and c are from Table 5, and a^- is a mild souring of a. We are to imagine that we start out with a and are offered a series of trades. For instance, at n_0 we can choose to trade a for the mild souring of a, or stick with a. At n_1, we would have the choice to trade a for b, or reject the offer and stick with a. And so on.

We can see that if we employ standard backwards induction or sophisticated choice (McClennen, 1990, p. 161), and the underlying ranking is either bottom-up or top-down bracketing, it is obligatory to trade a for a^- at n_0. Because first, since c is ranked above a, the former would be the uniquely permissible choice at n_2. This means that either trading for b or sticking with a and then choosing c would be permissible at n_1. Since b is ranked above c, the former would be uniquely permissible. By a similar argument, either choosing a^- or holding

[14] This observation suggests that, to avoid intransitivity, one could restrict top-down bracketing to cases where the set of maximal bracket-sets are identical for every comparison of prospects in the menu, and restrict bottom-up bracketing to cases where $\mathfrak{B}(I, \{a, b\}) = \mathfrak{B}(I, \{c, d\})$, for any pairs $\{a, b\}$ and $\{c, d\}$ in the set of available options. However, this response may appear ad-hoc and would limit the scope of both bottom-up and top-down bracketing, so we do not explore it further.

on to a and then trading for b is permissible at n_0. The outcome of the former plan is ranked above the outcome of the latter plan, and so trading a for a^- is obligatory at n_0. In other words, when combined with backwards induction, both bottom-up and top-down bracketing can be turned into value-pumps, recommending trading a prospect for a strictly worse.

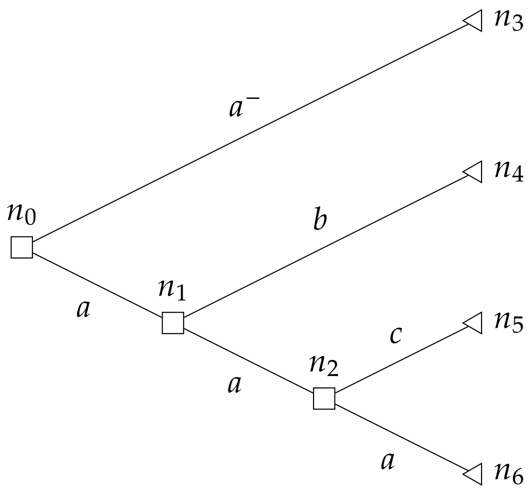

Figure 1: The upfront value-pump (cf. Gustafsson, 2022, p. 12)

Traditionally, there are two ways to reject the soundness of the value-pump argument. First, one could argue against (the implicit premise above) that trading for a^- is problematic. We set aside this kind of response here, but see e.g. Levi (2002). Second, one could dispute backwards induction, providing an argument for why a should not, in fact, be traded for a^- in the upfront value-pump. What motivates backwards induction is the so-called dynamic separability principle, which roughly states that one ought to be strictly forward-looking when evaluating plans, ignoring the larger dynamic context (McClennen, 1990, p. 122). This principle appears especially plausible in the context of broadly consequentialist dynamic choice. But some reject dynamic separability. For instance, proponents of resolute choice (McClennen, 1990; Gauthier, 1997) give up dynamic separability for so-called normal-form/extensive-form coincidence, which states that one should evaluate the set of plans at the root node as if one could implement them with one choice, disregarding evaluations at later nodes (McClennen, 1990, p. 115). Combined with Uncoveredness (as the way of choosing synchronically), we see that $\{ab, aac, aaa\}$ would be the set of permissible plans at n_0, which is sufficient in showing how resoluteness avoids the upfront value-pump. However, while we may have reasons to reject dynamic separability, resolute choice is questionable. In short, it is unclear why we should view dynamic choice problems as if they are normalised at the root, and then act according to the corresponding verdict throughout the decision problem. It seems that this would either amount to then choosing contrary to the underlying ranking at later points in the tree, or having the intention to implement a given plan somehow alter the ranking itself (cf. Buchak, 2013, p. 176-7). It is not clear why either mode of planning would be justified on a broadly consequentialist approach to dynamic choice. See Gustafsson (2022, pp. 66-74) for an extended critique of resolute choice, albeit in the context of preference-based choice.

Yet another approach is Ahmed's (2017) self-regulation procedure, which roughly requires that one first identify all possible final outcomes deemed acceptable and then, at each choice node, choose an option that keeps at least one such acceptable outcome reachable, if possible. While Ahmed's approach avoids the upfront value-pump, we do not further expand on the approach here, and refer interested readers to Gustafsson (2022, pp. 18-9) for discussion and a case in which self-regulation arguably falls short.

We will instead opt for a slightly different kind of response to the value-pump argument. In short, we will argue that there is an aspect of imprecise consequentialist dynamic choice which should play a role when we generalise theories of bracketing to the dynamic setting, which does not get picked up when we simply combine, say, $\preccurlyeq_I^{\mathsf{TD}}$ with some standard dynamic choice rule. Specifying this aspect is what we turn to now.

6.2 Wise choice

Before moving on to generalising top-down bracketing to the dynamic setting, and showing how said generalisation avoids value-pumps, we first need to have a brief look at so-called *wise choice* developed by Rabinowicz (1995; 1997; 2020). Wise choice is an alternative to standard backwards induction and sophisticated choice (as well as resolute choice), and may, in general terms, be stated as follows (Rabinowicz, 2020, p. 534): a plan is permissible just in case (i) it is performable or feasible, i.e., there is never a reason to deviate from it once embarked upon, and (ii) it results in an outcome (or more generally, an uncertain prospect) which is no worse than the result of any other feasible plan. Wise choice amounts to a minor violation of dynamic separability, but never to choosing contrary to the underlying ranking at any node. We hence deem this violation acceptable. Although Rabinowicz' target is quite a lot wider, wise choice is especially important for dynamic choice under incompleteness, as it blocks the value-pump arguments for completeness without reducing to resoluteness. (We will consider an example in a moment.) Indeed, similar ideas are also developed and argued for in Chang (2005), Williams (2014), Bader (2019), S. Bradley (ms) and Petersen (ms) specifically in the context of parity, indeterminacy, incommensurability, incomplete preferences, et cetera. Wise choice also resembles forwards induction (Gustafsson, 2022, p. 30).

We hence contend that dynamic choice based on $\preccurlyeq_I$ should be wise. Our claim is then that, when generalising top-down bracketing to the dynamic setting, not only should the c-goodness of plans matter, but in line with wise choice, it should also matter whether a given plan is feasible with respect to c-betterness. It is beyond the scope of the paper to give a more formal definition of feasibility, but the rough idea we run with here is that a plan π is feasible if and only if, at all nodes which can be reached by the plan, the plan is never deemed c-worse than some other plan which also reaches that same node and which itself has not been deemed unfeasible. (Note that the definition is recursive, but it is the kind of recursion which eventually terminates since we are working with finite decision trees.) We will call such a plan "c-feasible". Let $\Pi(T, n_0)$ denote the set of all plans in a decision tree T at the root node n_0, and denote the set of c-feasible plans by $\mathcal{F}_{\preccurlyeq_I}(\Pi(T, n_0))$. Consequentialist wise choice then consists in moreover ruling out plans in $\mathcal{F}_{\preccurlyeq_I}(\Pi(T, n_0))$ which result in a prospect which is c-worse than that of some other plan in $\mathcal{F}_{\preccurlyeq_I}(\Pi(T, n_0))$.

To illustrate how wise choice works, and how it differs from merely proceeding by backwards induction based on $\preccurlyeq_I$, consider the decision tree in Figure 2, where $a^- \prec_I a \not\bowtie_I b \not\bowtie_I a^-$. This decision problem, 'the single-souring value-pump' (cf. Gustafsson, 2022, p. 26), is often used in arguments for complete preference or betterness relations. First, suppose we decide to implement the plan of simply choosing a at n_0. This plan is clearly c-feasible, because no other plan results in a c-better prospect. Second, we also see that the plan bb is c-feasible. For $a \not\bowtie_I b$ and $a^- \not\bowtie_I b$, and so at both nodes, bb does not prescribe a choice which is c-worse than some other option. Finally, what about the plan ba^-? Since just choosing a is c-feasible and $a^- \prec_I a$, there is a reason to deviate from the plan at n_0, i.e., ba^- is not c-feasible. We thus end up with the set $\mathcal{F}_{\preccurlyeq_I}(\Pi(T, n_0)) = \{bb, a\}$. Within this set, neither plan results in a prospect which is c-worse than some other, and so the set of overall permissible plans according to wise choice based on $\preccurlyeq_I$ is $\{bb, a\}$. Imprecise consequentialism hence avoids the single-souring value-pump. Note again that this is achieved without ever choosing contrary to $\preccurlyeq_I$.

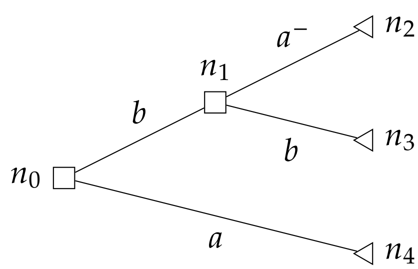

Figure 2: The single-souring value-pump (cf. Gustafsson, 2022, p. 26)

Suppose we instead employed backwards induction based on $\preccurlyeq_I$. At n_0, we have to reason about what would be chosen at n_1 and compare the corresponding prospect with a. Being completely forward-looking, one would plausibly choose arbitrarily at n_1, and so trading for b initially would, in effect, get us the prospect $a^- \cup b$. If c-betterness is in turn indeterminate betweeen $a^- \cup b$ and a, trading for b would be permissible at n_0. Choosing in such a way could then result in ending up with a^-, and hence being value-pumped. See Gustafsson (2022, pp. 34-7) for cases in which incompleteness and backwards induction appears even more problematic.

6.3 Dynamic bracketing

We now turn to generalising top-down bracketing to the dynamic setting. Recall from the last subsection that both c-betterness *and* wise choice c-feasibility should play into the definition. We will need the following piece of notation going forward. For a given decision tree T, generate a consequentialist betterness relation on the set of c-feasible plans, $\mathcal{F}_{\preccurlyeq_I}(\Pi(T, n_0))$, as follows: $\pi \preccurlyeq_I^* \pi'$ if and only if $O(\pi) \preccurlyeq_I O(\pi')$, where $O(\pi)$ denotes the prospect which the plan $\pi \in \mathcal{F}_{\preccurlyeq_I}(\Pi(T, n_0))$ would result in. We can then define a *dynamic bracket-set* as follows (this time relative to I, as opposed to any non-empty subset of I).

Definition 6.1 (Dynamic bracket-set)**.** *A* dynamic bracket-set *in a decision tree T relative to a pair $\{\pi, \pi'\} \subseteq \mathcal{F}_{\preccurlyeq_I}(\Pi(T, n_0))$ is a non-empty subset $I' \subseteq I$ such that $\pi \bowtie_{I'}^* \pi'$.*

A maximal dynamic bracket-set is defined just like before, i.e., as a dynamic bracket-set which cannot be enlarged in any way without no longer being a dynamic bracket-set. We then propose the following generalisation of top-down bracketing.

Definition 6.2 (Dynamic top-down bracketing)**.** *For any decision tree T, define a relation $\preccurlyeq_I^{\mathsf{DTD}}$ on $\mathcal{F}_{\preccurlyeq_I}(\Pi(T, n_0))$, with asymmetric and symmetric parts $\prec_I^{\mathsf{DTD}}$ and $\sim_I^{\mathsf{DTD}}$, representing* dynamic top-down bracketing *as follows, where $\mathtt{MB}^* = \{I'_1, \ldots, I'_k\}$ is the set of maximal dynamic bracket-sets relative to $\{\pi, \pi'\} \subseteq \mathcal{F}_{\preccurlyeq_I}(\Pi(T, n_0))$.*

$$\pi \prec_I^{\mathsf{DTD}} \pi' \iff (\forall I' \in \mathtt{MB}^*, \pi \preccurlyeq_{I'}^* \pi') \wedge (\exists I' \in \mathtt{MB}^*, \pi \prec_{I'}^* \pi') \tag{10}$$

$$\pi \sim_I^{\mathsf{DTD}} \pi' \iff (\forall I' \in \mathtt{MB}^*, \pi \sim_{I'}^* \pi') \tag{11}$$

If $\mathtt{MB}^* = \varnothing$, then $a \not\bowtie_I^{\mathsf{DTD}} b$.

The set of overall permissible plans in a given decision problem is then given by any suitable ways of choosing among $\mathcal{F}_{\preccurlyeq_I}(\Pi(T, n_0))$ based on $\preccurlyeq_I^{\mathsf{DTD}}$. We settle for Uncoveredness here (Definition 4.3). The definition is analogous to our synchronic version of top-down bracketing, where we can readily verify that $\preccurlyeq_I^{\mathsf{DTD}}$ is an extension of $\preccurlyeq_I^*$.

Let us now see how dynamic top-down bracketing deals with the upfront value-pump (Figure 1) before showing that the theory is generally not value-pumpable. Recall that we have a set of beneficiaries $I = \{1, 2\}$ and a set of final prospects $\{a, a^-, b, c\}$ such that

$$
\begin{array}{lll}
a^- \prec_{\{1\}} a, & a^- \prec_{\{2\}} a, & a^- \prec_I a, \\
c \prec_{\{1\}} a, & a \prec_{\{2\}} c, & a \prec_I c, \\
b \prec_{\{1\}} a, & a \not\bowtie_{\{2\}} b, & a \not\bowtie_I b, \\
c \prec_{\{1\}} b, & b \not\bowtie_{\{2\}} c, & b \not\bowtie_I c.
\end{array}
$$

By similar reasoning as before, we find that the set of c-feasible plans is $\mathcal{F}_{\preccurlyeq_I}(\Pi(T, n_0)) = \{ab, aac\}$. For aaa requires that we choose contrary to $\preccurlyeq_I$ at n_2. Similarly, since a^- is c-worse than c and aac is c-feasible, choosing a^- at n_0 would be to choose contrary to $\preccurlyeq_I$. Relative to the pair $\{ab, aac\}$, we immediately see that the unique maximal bracket-set is $\{1\}$, relative to which $c \prec^*_{\{1\}} b$. And so ab is the uniquely permissible plan according to dynamic top-down bracketing and we avoid the upfront value-pump.

Dynamic top-down bracketing is also not generally value-pumpable. For the set of c-feasible plans $\mathcal{F}_{\preccurlyeq_I}(\Pi(T, n_0))$ will never contain a plan which prescribes choices which lead to being value-pumped, and since dynamic top-down bracketing roughly consists in making a choice within this set, it trivially follows that top-down bracketing cannot be turned into a value-pump. Note that we could also generalise bottom-up bracketing in an analogous way, and show how the theory is not value-pumpable, *mutatis mutandis*.

At this point, one might object that our generalisation does, at least in part, require resoluteness. For dynamic top-down bracketing sometimes entails choosing contrary to $\preccurlyeq_I^{\mathsf{TD}}$. Consider, for instance, the choice at n_0 in the upfront value-pump. Note that $b \prec_I^{\mathsf{TD}} a^-$, and yet the uniquely permissible plan according to dynamic top-down bracketing is ab. In other words, we are asked to choose contrary to $\preccurlyeq_I^{\mathsf{TD}}$, the ranking which we spent the previous parts of this paper defending. But while this may this appear problematic, the idea here is just that c-feasibility takes precedence over the rankings of final prospects made by top-down bracketing. In this case, choosing a^- is not c-feasible, and so that is why this plan is ruled out. The motivation for having c-feasibility playing this role is similar to the motivation for preferring top-down bracketing to bottom-up bracketing in the first place. Namely, that top-down bracketing extends consequentialism. Similarly, if a plan is not c-feasible, then it should not be chosen by a dynamic generalisation of top-down bracketing.

As a last remark: we are not confident that what we have presented is the most plausible generalisation of bracketing. For one, we suspect there may be problem cases involving the gathering of information that could change the set of maximal dynamic bracket-sets. Other rules might handle such cases better. In particular, the benefactor might have the option to choose in such a way that they would obtain information that would lead to some I' being a maximal dynamic bracket-sets under their updated credences, despite not being so under their current ones. And our approach may not always recommend gathering such information, even when it seems intuitively required. Still, as long as they are consistent with wise choice, modifications to the generalisations we have presented here will avoid the value-pump objection.

7 From cluelessness to neartermism

Before concluding, we briefly discuss the implications of bracketing for *longtermism*, roughly the view that what matters most is how our actions affect the far future (Greaves and MacAskill, 2021; MacAskill, 2022). In brief, although it is not always true that we are more clueless about the long-term consequences of our actions than the more immediate ones, it should be uncontroversial that this is generally the case. Bracketing thus poses a direct challenge to longtermism. This section is devoted to laying out this argument in a bit more detail.

Greaves and MacAskill (2021) informally define *axiological strong longtermism* to be the thesis that far-future effects are the most important determinant of the value of our options. The idea is that if an option is best or near-best, then it is so in virtue of the effects it would have on the far future. They also define deontic strong longtermism to be the thesis that far-future effects are the most important determinant of what we ought to do. Assuming consequentialism, axiological strong longtermism entails deontic strong longtermism. (See Greaves and MacAskill, 2021, pp. 3, 26, for more precise definitions of axiological and deontic strong longtermism.) Strong longtermism seems especially well-motivated on precise, orthodox expected value theory (without discounting), because the vastness of the future means that even very small probabilities of enormous future value can swamp expected values.

However, it is easy to see how these forms of strong longtermism can, in combination with imprecise ex ante consequentialism, fail to be action-guiding. If our credences about the far-future consequences of our actions are sufficiently imprecise, strong longtermism will have it that it is indeterminate what option is overall best or what we ought to do. But the challenge from bracketing goes further, as it seems to push in favour of a 'neartermist' view on which options are best or most choiceworthy in virtue of their more immediate effects in the near-term. Again, this is because we are generally not clueless about what is c-good for (some non-empty subset of) beneficiaries in the near future, whereas we are generally clueless about what is c-good for future beneficiaries, and so the latter set should be bracketed out.

Greaves and MacAskill (2021, pp. 22-3) do consider cluelessness and imprecision as a challenge to strong longtermism. They write:

> While we do not take a stand on whether or not any imprecision of valuation is either rationally permissible or rationally required [...], we don't ourselves think that any plausible degree of imprecision in the case at hand will undermine the argument for strong longtermism. For example, we don't think any reasonable representor even contains a probability function according to which efforts to mitigate AI risk save only 0.001 lives per $100 in expectation. This does seem less clear, however, than the claim that this is not a reasonable precise credence function. (*ibid*, p. 23)

However, we take this to be an unjustified amount of confidence, failing to appreciate the extent of our ignorance of what the far future holds, and the long-term effects of our actions. Longtermists have sought to meet the epistemic challenge by pointing to the possibility of 'lock-in' events whose value we would not be clueless about, as they render the long-run future thereafter highly predictable (e.g., Tarsney 2023). Candidates include human extinction (bad) and the creation of a benevolent superintelligence designed to guarantee a flourishing civilization for the indefinite future (good). But even if it is true that we are not clueless about the value of lock-in events, this is not enough to show we are not clueless about the value of the actions that are available to us.

This is because, as Friederich (2025) argues, the causal pathways relating our actions to lock-in events are highly complex, and our efforts to navigate towards better lock-in states may backfire. Reducing nuclear weapons capacities could undermine deterrence and thus increase the probability of first strikes; research on novel pathogens aimed at preventing pandemics may enable the creation of bioweapons; and even drawing attention to catastrophic risks may backfire by raising the salience to bad actors of opportunities to cause harm (Friederich 2025, pp. 455-6). And our beliefs about the many pathways by which our actions could cause or prevent lock-in events should be so imprecise as to leave us clueless about our actions' long-term consequences.

The epistemic challenge is further exacerbated by our knowledge that we are unaware (Steele and Stefánsson, 2021) of various value-relevant ways the long-term future could turn out, and of many causal pathways by which our actions might affect lock-in events. While Greaves and MacAskill (2021, pp. 20-1) consider how conscious unawareness might be taken to be a challenge to strong longtermism, they appear to not recognise the argument that unawareness should lead to great imprecision (Roussos, 2021). In particular, introducing a 'catchall proposition'—roughly standing for all the possibilities we are unaware of—seems to not only put significant pressure on the idea of precise probability assignments, but also on small credal sets. As Henderson et al. (2010, p. 190) put it (in the context of the problem of new theories in the philosophy of science): "there is no particularly principled way to decide how much initial probability should be assigned to the catchall". Wenmackers and Romeijn (2016, p. 1235) (also in the context of the philosophy of science) even go as far as to suggest that we should give up probabilism altogether, assigning neither precise nor imprecise credences to the catchall. Although this approach arguably goes too far, it nevertheless suggests that large credal sets offer a more plausible accommodation of unawareness.[15]

In summary, the dialectic we have described in this section runs as follows. Initially, compelled by precise, orthodox expected value theory, it seems that the consequentialist ex ante value of our actions is dominated by their effects

[15] de Canson (2024, pp. 7, 14) also suggests that imprecise credences might be appropriate for the catchall. Also see R. Bradley (2017, p. 255) and Steele and Stefánsson (2021, p. 105) for discussion of the interaction between conscious unawareness and credal imprecision.

on the far future, and so strong longtermism naturally follows. However, once we recognise the predicament of cluelessness and adopt imprecise credences as a more appropriate representation of our uncertainty, strong longtermism leaves us without action guidance. Moreover, theories of bracketing that set aside consequences about which we are clueless, point in the opposite direction: toward focusing on the more immediate effects of our actions where determinate comparisons can be made.

8 Concluding remarks

In summary, we have formulated two alternatives to orthodox consequentialism motivated by the intuition that we should bracket out those consequences of our actions which we are clueless about. We favoured top-down bracketing over bottom-up bracketing on the grounds that the latter does not extend orthodox consequentialism and violates statewise dominance. Four objections to both kinds of bracketing were addressed—three being based on the fact that both theories do not generally rank options acyclically. First, in response to the semantic argument from cyclicity, we argued that bottom-up and top-down bracketing could both be understood in terms of choiceworthiness. Second, we argued that we simply have to refine our choice rules in light of cyclicity to reach reasonable verdicts in synchronic choice problems. Third, we objected that Hedden's argument from the weighing of reasons (adapted to our setting) simply assumes the orthodox consequentialist view. Finally, in response to the value-pump argument, we showed that a natural generalisation of bracketing is not susceptible to value-pumps.

To conclude the paper, we now want to point to some ways of generalising the ideas developed here. First, one could plausibly generalise theories of bracketing to reasons aggregation and reasons-based choice (cf. Dietrich and List, 2011; Lord and Maguire, 2016). Indeed, following e.g. Wedgwood (2022), there are reasons to think that the kind of formalism we employed in this paper is suitable for such a project. The imprecise probabilities from our setting may then correspond to more general imprecise weights on reasons. One could also work with some other dimension of prospects rather than beneficiaries, such as states or events. For instance, a bottom-up theory of statewise bracketing might recommend that one set aside states conditional on which one is clueless, or of which one is clueless.[16] Finally, the ideas here could plausibly also be applied to the problem of how you ought to act when you are uncertain about which moral theory is correct (MacAskill et al., 2020), as well as to the aggregation problem for value pluralists (Hedden and Muñoz, 2024).

Acknowledgements

We thank Gustav Alexandrie, Eric Olav Chen, Guillaume Corlouer, Lukas Finnveden, Aidan Goth, Alex Kastner, Andreas Mogensen, Caspar Oesterheld, Sami Petersen, Richard Pettigrew, Derek Shiller, Michael St. Jules, Christian Tarsney, Philip Trammell, and Timothy L. Williamson for valuable discussion and comments on a previous version of this paper.

References

Ahmed, A. (2017). Exploiting Cyclic Preference. *Mind*, *126*(504), 975–1022. https://doi.org/10.1093/mind/fzv218

Andersson, H., & Herlitz, A. (Eds.). (2022). *Value Incommensurability: Ethics, Risk, and Decision-Making*. Routledge. https://doi.org/https://doi.org/10.4324/9781003148012

Arrhenius, G., Ryberg, J., & Tännsjö, T. (2024). The Repugnant Conclusion. In E. N. Zalta & U. Nodelman (Eds.), *The Stanford Encyclopedia of Philosophy* (Winter 2024). Metaphysics Research Lab, Stanford University. https://plato.stanford.edu/archives/win2024/entries/repugnant-conclusion/

Askell, A. (2018). *Pareto Principles in Infinite Ethics* [Doctoral dissertation, New York University].

[16] Albeit of an entirely different sort, the discounting of tiny probabilities may be understood as a kind of statewise bracketing, where states are, very roughly, set aside if their associated probability is sufficiently small. Perhaps unsurprisingly then, variations of such discounting lead to violations of acyclicity and statewise dominance (Kosonen 2022, pp. 137-95, ms; Cibinel, 2023). Thanks to Eric Olav Chen for pointing out this connection.

Bader, R. M. (2019). Agent-Relative Prerogatives and Suboptimal Beneficence. In M. Timmons (Ed.), *Oxford Studies in Normative Ethics* (pp. 223–250, Vol. 9). Oxford University Press. https://doi.org/https://doi.org/10.1093/oso/9780198846253.003.0011

Blackorby, C., Bossert, W., & Donaldson, D. (2007). Variable-population extensions of social aggregation theorems. *Social Choice and Welfare*, *28*(4), 567–589. https://doi.org/10.1007/s00355-006-0188-x.

Bolker, E. D. (1966). Functions Resembling Quotients of Measures. *Transactions of the American Mathematical Society*, *2*, 292–312.

Bolker, E. D. (1967). A Simultaneous Axiomatization of Utility and Subjective Probability. *Philosophy of Science*, *34*(4), 333–340. https://doi.org/10.1086/288171

Boole, G. (1940). *Collected Logical Works* (Vol. 2). The Open Court Publishing Company.

Bradley, R. (2015). A Note on Incompleteness, Transitivity and Suzumura Consistency. In C. Binder, G. Codognato, M. Teschl, & Y. Xu (Eds.), *Individual and Collective Choice and Social Welfare* (1st, pp. 31–47). Springer. https://doi.org/https://doi.org/10.1007/978-3-662-46439-7

Bradley, R. (2017). *Decision Theory with a Human Face*. Cambridge University Press. https://doi.org/https://doi.org/10.1017/9780511760105

Bradley, R., & Drechsler, M. (2013). Types of Uncertainty. *Erkenntnis*, *79*(6), 1225–1248. https://doi.org/10.1007/s10670-013-9518-4

Bradley, S. (ms). *Pick the Sugar*.

Brandt, F., Brill, M., & Harrenstein, P. (2016). Tournament Solutions. In F. Brandt, V. Conitzer, U. Endriss, J. Lang, & A. D. Procaccia (Eds.), *Handbook of Computational Social Choice* (pp. 57–84). Cambridge University Press. https://doi.org/https://doi.org/10.1017/CBO9781107446984

Broome, J. (1990). Bolker-Jeffrey Expected Utility Theory and Axiomatic Utilitarianism. *The Review of Economic Studies*, *57*(3), 477–502. https://doi.org/https://doi.org/10.2307/2298025

Broome, J. (1991). *Weighing Goods: Equality, Uncertainty and Time*. Wiley-Blackwell. https://doi.org/10.1002/9781119451266

Broome, J. (1997). Is Incommensurability Vagueness? In R. Chang (Ed.), *Incommensurability, Incomparability, and Practical Reason*. Harvard.

Broome, J. (2004). *Weighing Lives*. Oxford University Press. https://doi.org/10.1093/019924376X.001.0001

Buchak, L. (2013). *Risk and Rationality*. Oxford University Press. https://doi.org/10.1093/acprof:oso/9780199672165.001.0001

Chang, R. (Ed.). (1997). *Incommensurability, Incomparability, and Practical Reason*. Harvard.

Chang, R. (2002). The Possibility of Parity. *Ethics*, *112*(4), 659–688. https://doi.org/10.1086/339673

Chang, R. (2005). Parity, Interval Value, and Choice. *Ethics*, *115*(2), 331–350. https://doi.org/10.1086/426307

Cibinel, P. (2023). A Dilemma for Nicolausian Discounting. *Analysis*, *83*(4), 662–672. https://doi.org/10.1093/analys/anac095

Davidson, D., McKinsey, J., & Suppes, P. (1955). Outlines of a Formal Theory of Value, I. *Philosophy of Science*, *22*(2), 140–160. https://doi.org/10.1086/287412

de Canson, C. (2024). The Nature of Awareness Growth. *Philosophical Review*, *133*(1), 1–32. https://doi.org/10.1215/00318108-10880419

Dietrich, F., & List, C. (2011). A Reason-Based Theory of Rational Choice. *Noûs*, *47*(1), 104–134. https://doi.org/10.1111/j.1468-0068.2011.00840.x

Ellsberg, D. (1961). Risk, Ambiguity, and the Savage Axioms. *Quarterly Journal of Economics*, 643–69.

Fishburn, P. C. (1977). Condorcet social choice functions. *SIAM Journal on Applied Mathematics*, *33*(3), 469–489. https://doi.org/10.1137/0133040

Friederich, S. (2025). Causation, Cluelessness, and the Long Term. *Ergo: An Open Access Journal of Philosophy*, *12*. https://doi.org/10.3998/ergo.7428

Gärdenfors, P., & Sahlin, N.-E. (1982). Unreliable probabilities, risk taking, and decision making. *Synthese*, *53*(3), 361–386. https://doi.org/10.1007/BF00486156

Gauthier, D. (1997). Resolute Choice and Rational Deliberation: A Critique and a Defense. *Noûs*, *31*(1), 1–25. https://doi.org/10.1111/0029-4624.00033

Greaves, H. (2016). Cluelessness. *Proceedings of the Aristotelian Society*, *116*(3), 311–339. https://doi.org/10.1093/arisoc/aow018

Greaves, H., & MacAskill, W. (2021). The Case for Strong Longtermism [Global Priorities Institute Working Paper, No. 5].

Griffin, J. (1986). *Well-Being: Its Meaning, Measurement, and Moral Importance*. Clarendon Press. https://doi.org/https://doi.org/10.1093/0198248431.001.0001

Gustafsson, J. E. (2022). *Money-Pump Arguments*. Cambridge University Press. https://doi.org/https://doi.org/10.1017/9781108754750

Hansson, S. O. (1996). Decision Making Under Great Uncertainty. *Philosophy of the Social Sciences*, *26*(3), 369–386. https://doi.org/10.1177/004839319602600304

Hare, C. (2013). *The Limits of Kindness*. Oxford University Press. https://doi.org/ttps://doi.org/10.1093/acprof:oso/9780199691999.001.0001

Harsanyi, J. C. (1955). Cardinal Welfare, Individualistic Ethics, and Interpersonal Comparisons of Utility. *Journal of Political Economy*, *63*(4), 309–321.

Hedden, B. (2024). Parity and Pareto. *Philosophy and Phenomenological Research*, *109*(2), 575–592. https://doi.org/10.1111/phpr.13056

Hedden, B., & Muñoz, D. (2024). Dimensions of Value. *Noûs*, *58*(2), 291–305. https://doi.org/10.1111/nous.12454

Helgeson, C. (2020). Structuring Decisions Under Deep Uncertainty. *Topoi*, *39*(2), 257–269. https://doi.org/10.1007/s11245-018-9584-y

Henderson, L., Goodman, N. D., Tenenbaum, J. B., & Woodward, J. F. (2010). The Structure and Dynamics of Scientific Theories: A Hierarchical Bayesian Perspective. *Philosophy of Science*, *77*(2), 172–200. https://doi.org/10.1086/651319

Herlitz, A. (2018). Nondeterminacy, Two-Step Models, and Justified Choice. *Ethics*, *129*(2), 284–308. https://doi.org/10.1086/700032

Herlitz, A. (2022). Nondeterminacy and Reasonable Choice. In H. Andersson & A. Herlitz (Eds.), *Value Incommensurability: Ethics, Risk, and Decision-Making* (pp. 143–161). Routledge. https://doi.org/10.4324/9781003148012-11

Hill, B. (2019). Confidence in Beliefs and Rational Decision Making. *Economics and Philosophy*, *35*(2), 223–258. https://doi.org/10.1017/s0266267118000214

Insua, D. R. (1992). On the Foundations of Decision Making Under Partial Information. *Theory and Decision*, *33*(1), 83–100. https://doi.org/10.1007/bf00133984

Jeffrey, R. (1983a). Bayesianism with a Human Face. In J. Earman (Ed.), *Testing Scientific Theories* (pp. 133–156). Minneapolis: University of Minnesota Press. https://hdl.handle.net/11299/185349

Jeffrey, R. (1983b). *The Logic of Decision* [Originally published in 1965]. University of Chicago Press.

Joyce, J. M. (2005). How Degrees of Belief Reflect Evidence. *Philosophical Perspectives*, *19*(1), 153–179. https://doi.org/10.1111/j.1520-8583.2005.00058.x

Joyce, J. M. (2010). A Defense of Imprecise Credences in Inference and Decision Making. *Philosophical Perspectives*, *24*(1), 281–323. https://doi.org/10.1111/j.1520-8583.2010.00194.x

Kagan, S. (1998). *Normative Ethics*. Routledge. https://doi.org/10.4324/9780429498657

Kollin, S. (ms). *The Greater Good that Benefits None*.

Konek, J. (2019). Comparative Probabilities. In R. Pettigrew & J. Weisberg (Eds.), *The Open Handbook of Formal Epistemology* (pp. 267–348). PhilPapers Foundation.

Koopman, B. O. (1940). The Axioms and Algebra of Intuitive Probability. *Annals of Mathematics*, *41*(2), 269–292. https://doi.org/https://doi.org/10.2307/1969003

Kosonen, P. (2022). *Tiny Probabilities of Vast Value* [Doctoral dissertation, Oxford University].

Kosonen, P. (ms). *How to Discount Small Probabilities*.

Lenman, J. (2000). Consequentialism and Cluelessness. *Philosophy and Public Affairs*, *29*(4), 342–370. https://doi.org/10.1111/j.1088-4963.2000.00342.x

Levi, I. (1974). On Indeterminate Probabilities. *Journal of Philosophy*, *71*(13), 391–418. https://doi.org/10.2307/2025161

Levi, I. (2002). Money Pumps and Diachronic Books. *Philosophy of Science*, *69*(S3), S235–S247. https://doi.org/https://doi.org/10.1086/341849

Lord, E., & Maguire, B. (Eds.). (2016). *Weighing Reasons*. Oxford University Press. https://doi.org/10.1093/acprof:oso/9780199315192.001.0001

MacAskill, W. (2013). The Infectiousness of Nihilism. *Ethics*, *123*(3), 508–520. https://doi.org/10.1086/669564

MacAskill, W. (2022). *What We Owe the Future*. Basic Books.

MacAskill, W., Bykvist, K., & Ord, T. (2020). *Moral Uncertainty*. https://doi.org/https://doi.org/10.1093/oso/9780198722274.001.0001

Mahtani, A. (2018). Imprecise Probabilities and Unstable Betting Behaviour. *Noûs*, *52*(1), 69–87. https://doi.org/10.1111/nous.12148

Mahtani, A. (2019). Imprecise Probabilities. In R. Pettigrew & J. Weisberg (Eds.), *The Open Handbook of Formal Epistemology* (pp. 107–130). PhilPapers Foundation.

Mahtani, A. (2020). The Dispositional Account of Credence. *Philosophical Studies*, *177*(3), 727–745. https://doi.org/10.1007/s11098-018-1203-7

McClennen, E. F. (1990). *Rationality and Dynamic Choice: Foundational Explorations*. Cambridge University Press. https://doi.org/https://doi.org/10.1017/CBO9780511983979

Meacham, C. J. G. (2012). Person-Affecting Views and Saturating Counterpart Relations. *Philosophical Studies*, *158*(2), 257–287. https://doi.org/10.1007/s11098-012-9884-9

Miller, N. R. (1980). A New Solution Set for Tournaments and Majority Voting: Further Graph-Theoretical Approaches to the Theory of Voting. *American Journal of Political Science*, *24*(1), 68–96. https://doi.org/https://doi.org/10.2307/2110925

Mogensen, A. (2021). Maximal Cluelessness. *Philosophical Quarterly*, *71*(1), 141–162. https://doi.org/10.1093/pq/pqaa021

Mogensen, A., & Thorstad, D. (2022). Tough enough? Robust satisficing as a decision norm for long-term policy analysis. *Synthese*, *200*(36). https://doi.org/10.1007/s11229-022-03566-5

Monton, B. (2019). How to Avoid Maximizing Expected Utility. *Philosophers' Imprint*, *19*. https://doi.org/http://hdl.handle.net/2027/spo.3521354.0019.018

Moore, G. E. (1903). *Principia Ethica* (T. Baldwin, Ed.). Dover Publications.

Nebel, J. M. (2019). Asymmetries in the Value of Existence. *Philosophical Perspectives*, *33*(1), 126–145. https://doi.org/10.1111/phpe.12128

Parfit, D. (1984). *Reasons and Persons*. Oxford University Press. https://doi.org/https://doi.org/10.1093/019824908X.001.0001

Parfit, D. (2011). *On What Matters: Volume Two*. https://doi.org/https://doi.org/10.1093/acprof:osobl/9780199572816.001.0001

Parfit, D. (2016). Can We Avoid the Repugnant Conclusion? *Theoria*, *82*(1), 110–127. https://doi.org/10.1111/theo.12097

Petersen, S. (ms). *Dynamic Choice without the Completeness Axiom*.

Rabinowicz, W. (1995). To Have One's Cake and Eat It, Too: Sequential Choice and Expected-Utility Violations. *The Journal of Philosophy*, *92*(11), 586–620. https://doi.org/10.2307/2941089

Rabinowicz, W. (1997). Wise Choice on Dynamic Decision-Making Without Independence. In E. Ejerhed & S. Lindström (Eds.), *Logic, Action and Cognition* (Vol. 2). Springer. https://doi.org/10.1007/978-94-011-5524-3_5

Rabinowicz, W. (2020). Between Sophistication and Resolution: Wise Choice. In K. Sylvan (Ed.), *Routledge Handbook of Practical Reason* (1st). Routledge. https://doi.org/10.4324/9780429266768-43

Rachels, S. (1998). Counterexamples to the Transitivity of Better Than. *Australasian Journal of Philosophy*, *76*(1), 71–83. https://doi.org/10.1080/00048409812348201

Ramsey, F. P. (1928). A Mathematical Theory of Saving. *The Economic Journal*, *38*(152), 543–559. https://doi.org/10.2307/2224098

Raz, J. (1986). *The Morality of Freedom*. Oxford University Press. https://doi.org/https://doi.org/10.1093/0198248075.001.0001

Roussos, J. (2021). Unawareness for Longtermists [Presentation at the 7th Oxford Workshop on Global Priorities Research]. https://joeroussos.org/wp-content/uploads/2021/11/210624-Roussos-GPI-Unawareness-and-longtermism.pdf

Schwartz, T. (1972). Rationality and the Myth of the Maximum. *Noûs*, *6*(2), 97–117. https://doi.org/10.2307/2216143

Schwartz, T. (1990). Cyclic Tournaments and Cooperative Majority Voting: A Solution. *Social Choice and Welfare*, *7*, 19–29. https://doi.org/10.1007/BF01832917

Sen, A. (1997). Maximization and the Act of Choice. *Econometrica*, *65*(4), 745–779. https://doi.org/10.2307/2171939

Sen, A. (2004). Incompleteness and Reasoned Choice. *Synthese*, *140*(1-2), 43–59. https://doi.org/10.1023/b:synt.0000029940.51537.b3

Sen, A. (2017). *Collective Choice and Social Welfare: An Expanded Edition* (Expanded) [Originally published in 1970]. Penguin Books. https://doi.org/https://doi.org/10.2307/j.ctv2sp3dqx

Singer, P. (2015). *The Most Good You Can Do: How Effective Altruism Is Changing Ideas About Living Ethically*. Yale University Press. https://doi.org/10.12987/9780300182415

Skyrms, B. (1980). Higher Order Degrees of Belief. In D. H. Mellor (Ed.), *Prospects for Pragmatism: Essays in Memory of F. P. Ramsey* (pp. 109–137). Cambridge University Press.

Steele, K. (2022). Incommensurability That Can(not) Be Ignored. In H. Andersson & A. Herlitz (Eds.), *Value Incommensurability: Ethics, Risk, and Decision-Making* (pp. 231–246). Routledge. https://doi.org/10.4324/9781003148012-16

Steele, K., & Stefánsson, H. O. (2021). *Beyond Uncertainty: Reasoning with Unknown Possibilities*. Cambridge University Press. https://doi.org/https://doi.org/10.1017/9781108582230

Stefánsson, H. O. (2017). What is "Real" in Probabilism? *Australasian Journal of Philosophy*, *95*(3), 573–587. https://doi.org/10.1080/00048402.2016.1224906

Tarsney, C. (2023). The epistemic challenge to longtermism. *Synthese*, *201*(195). https://doi.org/10.1007/s11229-023-04153-y

Temkin, L. S. (1987). Intransitivity and the Mere Addition Paradox. *Philosophy and Public Affairs*, *16*(2), 138–187.

Temkin, L. S. (2012). *Rethinking the Good: Moral Ideals and the Nature of Practical Reasoning*. Oxford University Press.

Thorstad, D., & Mogensen, A. (2020). Heuristics for Clueless Agents: How to Get Away with Ignoring What Matters Most in Ordinary Decision-Making [Global Priorities Institute Working Paper, No. 2].

Vallentyne, P., & Kagan, S. (1997). Infinite Value and Finitely Additive Value Theory. *Journal of Philosophy*, *94*(1), 5–26. https://doi.org/https://doi.org/10.2307/2941011

van Fraassen, B. (1990). Figures in a Probability Landscape. In J. Dunn & A. Gupta (Eds.), *Truth or Consequences: Essays in Honor of Nuel Belnap* (pp. 345–356). Kluwer Academic Publishers. https://doi.org/https://doi.org/10.1007/978-94-009-0681-5_21

Wedgwood, R. (2022). The Reasons Aggregation Theorem. *Oxford Studies in Normative Ethics*, *12*, 127–148. https://doi.org/https://doi.org/10.1093/oso/9780192868886.003.0007

Wenmackers, S., & Romeijn, J. (2016). New Theory About Old Evidence. A Framework for Open-Minded Bayesianism. *Synthese*, *193*(4). https://doi.org/10.1007/s11229-014-0632-x

White, R. (2009). Evidential Symmetry and Mushy Credence. *Oxford Studies in Epistemology*, *3*, 161–186.

Wilkinson, H. (2020). Infinite Aggregation: Expanded Addition. *Philosophical Studies*, *178*(6), 1917–1949. https://doi.org/10.1007/s11098-020-01516-w

Wilkinson, H. (2021). *Infinite Aggregation* [Doctoral dissertation, Australian National University]. https://doi.org/10.25911/KFA4-0C89

Williams, J. R. G. (2014). Decision-Making Under Indeterminacy. *Philosophers' Imprint*, *14*(4), 1–34. http://hdl.handle.net/2027/spo.3521354.0014.004