Fixed-Sum Choiceworthiness Distribution in Maximizing Expected Choiceworthiness

Abstract

When choosing under moral uncertainty, one popular framework is to Maximize Expected Choiceworthiness (MEC). In MEC, many approaches to assign choiceworthiness and conduct intertheoretic comparison exist. Such approaches are faced with the challenge of accurately representing moral theories' judgements in assigning choiceworthiness to options and simultaneously enabling fair aggregation of assignments into a compound judgement. Fixed-Sum Choiceworthiness Distribution (FSCD) is a novel approach to choiceworthiness assignment and intertheoretic comparison. FSCD normalises moral theories' choiceworthiness assignments to a fixed sum of assigned choiceworthiness across options. By parallel to social choice theory's cumulative vote, it reflects ordinal theories and functions as a structural solution to intertheoretic comparability; by normalizing cardinal choiceworthiness assignments to an equal sum, it establishes equal say following the analogy of the moral parliament; and by variability in the range of choiceworthiness, it ensures sensitivity to stakes. This positions MEC with FSCD as a promising approach to moral uncertainty.

I. Moral Uncertainty and Maximizing Expected Choiceworthiness

To aggregate ethical theories' judgements when maximizing expected choiceworthiness, each theory should be required to distribute the same fixed sum of choiceworthiness over options. This is 'Fixed-Sum Choiceworthiness', a novel method to maximize expected choiceworthiness under moral uncertainty introduced in this paper.

1. Moral Uncertainty

In making a moral choice between options, we might consider two factors: What happens if we choose an option, and what moral value we ascribe to it. The former is obviously shrouded in uncertainty: We are not always sure what will happen due to our choices, and so we often employ decision-theoretical frameworks such as expected utility theory to give notice to all possible outcomes.

But the latter is likewise uncertain: Given disagreement around what the correct theory of morality is, we might be uncertain about what moral value options have: In the face of ethical dilemmas and subsisting moral disagreement, we might not be entirely confident that any one moral theory is entirely correct and instead put credence in many views. Hence, instead of following any one theory's action-guidance, we might likewise employ a framework to give notice to all plausible moral views.

2. Maximizing Expected Choiceworthiness

Recent years have seen increased work on such frameworks, such as by Lockhart (2000), Sepielli (2009) and Bostrom (2010). One particularly comprehensive framework mirrors the strategy of maximizing expected utility under empirical uncertainty - by maximizing expected choiceworthiness under moral uncertainty, in short MEC, first suggested by MacAskill (2014) and substantially developed by MacAskill (2016) and MacAskill et al. (2020, 2020b). To guide an agent's decision under moral uncertainty, MEC draws on two elements:

(a) An agent's distribution of credence over moral theories. For instance, an agent might think there is some merit to a consequentialist theory C and to a deontological theory D, and think them both equally plausible. The degree of this plausibility is the assigned credence; the set of such credences - adding up to 100% assigned to all theories at least somewhat plausible to the agent - is the credence distribution. For instance, the agent might have 50% credence in both C and D. (MacAskill 2014)

(b) Moral theories' assignment of choiceworthiness to options. In any given choice scenario, an agent can choose from a number of options; moral theories assign value in terms of choiceworthiness to these options.[1] (MacAskill 2014) For instance, in deciding whether to tell

[1] Choiceworthiness does not necessarily refer to the value of a causal consequence. It might just as well refer to the comparative choiceworthiness of different act-types (Colyvan et al. 2010), or a non-multiplicative model of assigning moral value (Tenenbaum 2017), or similar.

a white lie, a consequentialist view might assign choiceworthiness based on expected welfare – telling the lie produces more welfare than not telling it -, a deontological view might assign choiceworthiness in terms of (expected) compliance with duties – no duty compels the agent to lie, but there is a duty not to lie. The precise assignment of choiceworthiness is subject of extensive and controversial debate and will be discussed later; the approach suggested in this paper offers a novel way to conduct this assignment.

This information is processed as it would be in expected utility theory (EUT). In the parallel to EUT, choiceworthiness corresponds to cardinal utility and credence corresponds to likelihood of outcomes – s. MacAskill (2014) for elaborate discussion. For each option, an expected choiceworthiness value is determined by multiplying the credence the agent has in each moral theory which the choiceworthiness that theory assigns to the option, and then adding the results – as in the table below:

	C (50%)	D (50%)	Expected Choiceworthiness
Tell Lie	8	0	$50\% \times 8 + 50\% \times 0 = 4$
Don't Tell Lie	5	10	$50\% \times 5 + 50\% \times 15 = 7.5$

The aggregation suggests that the expected choiceworthiness of not telling the lie is higher and the agent should therefore choose not to tell the lie.

Two parallels motivate this framework: Its method of addressing moral uncertainty mirrors the prominent approach of maximizing expected utility under non-normative uncertainty; and its method of aggregating inputs of single theories into compound judgements often mirrors and draws from social choice theory. Thereby, MEC has emerged as a preeminent approach to moral uncertainty, receiving both ample support and further development (e.g. by Tarsney 2018, 2019; Trammell 2019; Riedener 2019; MacAskill et al. 2020b) and critical discussion (e.g. by Hedden 2016; Carr 2020, 2021; Newberry & Ord 2021).

II. Choiceworthiness Assignment: Requirements & Open Challenges

One central controversial feature of MEC is how choiceworthiness assignment is conducted – i.e. how we get from individual theories' moral judgement to aggregable numbers of choiceworthiness assigned to options. Several approaches for choiceworthiness assignment have been developed – and confronted with forceful criticism and open questions. This paper suggests a novel choiceworthiness assignment approach in light of these open questions. To motivate this approach, I will first develop two crucial requirements and a central open challenge for choiceworthiness assignment.

Choiceworthiness assignment approaches must provide theory representation and theory comparability:

1. Theory Representation

The requirement of theory representation holds that the choiceworthiness values must reflect the content of the theory's moral judgement[2]. Whatever moral judgement the theory makes, it must be integrated into the choiceworthiness distribution; otherwise, the aggregated guidance might disregard that judgement – for more expansive discussion. s. MacAskill et al. (2020, 58-60; 78-90). Most simply, this means that if a theory finds one option preferable over another, the former must be assigned more choiceworthiness. Furthermore, many theories make more complex statements: There might be indifference between options, some pairs of options might differ in choiceworthiness more than others, etc. Beyond these distinctions, I follow the literature in only considering theories with complete and transitive preferences over outcomes and lotteries.

Degrees of judgement complexities and the resulting requirements for choiceworthiness assignment can be understood as theories operating on different scale levels, roughly following distinctions drawn by MacAskill et al. (2020).

A theory that merely expresses not further quantified preferences between options operates on an ordinal scale. Kantian moral theory - at least in non-lottery choices - might be an example of this: Acting from duty is more desirable than doing nothing is more desirable than actively violating a duty; yet the theory makes no statements regarding (a) how the difference between inaction and duty violation compares to the difference between inaction and acting from duty or (b) that acting from one duty might be more desirable than acting from another. Accordingly, it makes no statements about intervals or levels of choiceworthiness. To accommodate ordinal theories, the choiceworthiness representation must assign more preferable options greater choiceworthiness than less preferable ones, and equally preferable options the same choiceworthiness: It must reflect statements such as $C_a > C_b$ and $C_a = C_c$.

Whereas ordinal theories make categorical, i.e. not further quantified, assessments, a theory with stronger and weaker preferences operates on a cardinal scale – i.e. an interval or ratio scale.[3]

A total consequentialist theory that desires to maximize Quality-Adjusted Life Years (QALYs) provides notions of intervals: Increasing overall welfare by 100 QALYs is better than increasing it by 50 QALYs, and so on. This can be reflected by an interval scale; to accommodate such preferences, the choiceworthiness representation must include information regarding the comparative size of intervals – i.e. statements of the kinds $(C_a - C_b) > (C_b - C_c)$ and $(C_a - C_b) = (C_b - C_d)$. Furthermore, it might quantify this difference, i.e. by statements of the kind $(C_a - C_b) = x(C_b - C_c)$, where the factor x expresses comparative interval size.

[2] This judgement could relate to action-guidance or axiology; s. Riedener (2015)
[3] Sometimes, 'cardinal scale' is used as a synonym for 'interval scale' instead. Herein, I use 'cardinal' to mean 'ratio or interval'.

Further assignment of values might be made based on ratio information: Generating 200 QALYs might be twice as choiceworthy as generating 100 QALYs. To accommodate ratio information, choiceworthiness representation must include information regarding the ratio y of choiceworthiness values – i.e. statements of the kind $C_a = yC_b$.

2. Theory Comparability

The requirement of theory comparability holds that choiceworthiness values assigned by moral theories must be comparable – and thereby aggregable – between different moral theories. There is a metaethical and a practical challenge to this issue of intertheoretic comparability:

Metaethically, it is not clear whether the moral preferences expressed as choiceworthiness are sufficiently comparable to be aggregable. A deontological claim that an action is in violation of a duty' and a consequentialist claim that an action causes a loss 50 QALYs seem different. Defenders might argue those are ultimately comparable claims, e.g. based on a universal-scale view by MacAskill et al. (2020), or on a constructivist approach by Riedener (2019) – but critics might take them to be type-different to an extent that prevents meaningful aggregation (s. e.g. Gracely 1996, Nissan-Rozen 2015).

Practically, no matter one's metaethical justification given the former problem, intertheoretical aggregation is challenging: Duties and QALYs might both track choiceworthiness – or not –, but they are not the same thing. We cannot, for instance, express the deontologist's choiceworthiness by counting duties and the consequentialist's choiceworthiness by counting QALYs; it is unclear how many QALYs correspond to what duty violation. Thus, with or without metaethical justification, choiceworthiness assignment must often undergo some homogenisation to ensure functional aggregation of individual theories' choiceworthiness assignments. Such homogenisation might employ mathematical normalization, such as MacAskill et al. (2020b)'s variance normalization; or voting approaches, as described e.g. in MacAskill (2016) or Tarsney (2019).

3. *Unequal Stakes & Equal Say*

In simultaneously ensuring representation and comparability, a critical tension has emerged: Between reflecting moral theories' unequal stakes in a decision on the one hand and ensuring equal say between theories on the other hand.

Sensitivity to varying stakes might be desirable: Sometimes, some relevant moral theories have more moral value to gain and lose from a choice situation than others. For a deontological view, only a minor duty violation might be at stake in a decision, while a consequentialist view might face a substantial delta in the consequences at stake – the comparative choiceworthiness assignment should be able to reflect this. Sepielli (2013) first explicitly advances this point, as a forceful response to Lockhart (2000)'s apparently stake-insensitive approach.

Simultaneously, however, theories should have equal say: No theory should have more say over the ultimate aggregate choice than another just because of some internal feature of the theory, i.e. without external justification from comparative stakes or credence assignments. For instance, equal say should prevent particularly fervent and zealous theories of morality to dominate any aggregation by virtue of their higher degree of conviction. To give one theory such unequal say would mean to bias the ultimate aggregation output in its favour, preventing a neutral judgement given uncertainty – s. MacAskill et al. (2020b) for in-depth discussion.

A tension arises because equal say is often at practical odds with stake sensitivity. On the one hand, too stringent of a rule on equal say might limit choiceworthiness assignment and prevent stake sensitivity: If each theory had to cast one vote on which option to choose, they might all have equal say – but no way to express higher or lower stakes. Inversely, too much leeway for stake sensitivity might threaten equal say: If a sufficiently dramatic theory was free to assign near-infinite choiceworthiness values without any normalization, that might objectionably sway the overall calculus in that theory's favour regardless of credences. This tension has been extensively discussed and leveraged by criticism of MEC (s. e.g. Lockhart 2000, Sepielli 2013, MacAskill et al. 2020, 84-87, Newberry & Ord 2021, 2).

In summary, to provide action-guidance under moral uncertainty based on the MEC framework, any approach to assigning choiceworthiness must enable:

a) Representation of moral theories' judgements
b) Comparability of choiceworthiness assignments between moral theories

To answer a weighty critique of MEC-based approaches, the choiceworthiness assignment approach should furthermore simultaneously provide:

c) A sensitivity to differences in stakes between theories
d) Equal say to ensure no arbitrary distortions of the aggregation

III. Introducing Fixed-Sum Choiceworthiness Distribution

I suggest an approach to choiceworthiness assignment that successfully navigates these desiderata. This approach is Fixed-Sum Choiceworthiness Distribution.

> *Fixed-Sum Choiceworthiness Distribution (FSCD)*: Every theory distributes the same, fixed amount of choiceworthiness s across the available options in a given choice scenario.

The principle of FSCD can be described in two convergent ways, each of which provide independent plausibility: as a voting method or as a data normalization method. Under both methods, any given

option is assigned the same transformed and thereby comparable choiceworthiness values C_x, which can then be multiplied by theory credences in the aggregative context of MEC.

1. FSCD as a Voting Method

FSCD can be described as a voting method: Theories provide a ranking of options and accordingly assign votes; votes correspond to choiceworthiness.

> *Voting Method:* For any given choice scenario under moral uncertainty, FSCD provides each theory with a fixed sum s of votes to be distributed freely across options:
>
> $$\{C | C \in [0, s]^X \wedge \sum_{A \in X} C(A) = s\}$$

This voting method expression draws on social choice theory: As pointed out by Lockhart (2000) and more recently MacAskill (2016), the aggregation of theories into a compound judgement in the context runs parallel to the aggregation of individual preferences into a social choice. Accordingly, different approaches to theory representation and aggregation run analogous to voting methods in social choice. To fulfil this requirement of MEC, theories' votes must be aggregable into a shared measure of choiceworthiness; the corresponding requirement in social choice is that votes be aggregable into scores for candidates. (MacAskill 2016, Pacuit 2019)

The social-choice-parallel to FSCD is cumulative voting, where each voter is asked to distribute a sum of votes across available candidates, with no further constraints (Felsenthal 1985). As demonstrated by Myerson (1995) and Pivato (2013), cumulative voting, by satisfying several basic conditions, is generally suitable to consolidate votes into an overall score – correspondingly, FSCD is suited to consolidate individual theories' input into an overall choiceworthiness measure for options, motivating use in MEC.

Other voting approaches also fulfil this requirement, but I believe cumulative vote and thereby FSCD is particularly promising: Voting approaches' comparative merit is often evaluated based on their respective failure modes and violations of intuitive principles (Pacuit 2019). I endorse FSCD because I believe cumulative vote's core failure modes of option individuation and tactical motivation can plausibly resolved in the context of moral uncertainty; I discuss this in depth later on.

2. FSCD as a Data Normalization Method

FSCD can also be described as a data normalization method: Quantifiable cardinal theories assign initially incomparable values to options, which are then mathematically normalized for comparison and aggregation.

> *Normalization:* For any given choice scenario under moral uncertainty, FSCD normalizes theories' choiceworthiness value assignments to available outcomes to a fixed sum of

choiceworthiness s: For any given theory, let C_x be the normalized choiceworthiness of an option x; c_x the non-normalized moral value or choiceworthiness of x; and s the fixed sum as determined by FSCD. Only values of $c_x \geq 0$ are admissible.[4] Choiceworthiness values for a given theory are normalized as follows:

$$C_x = c_x \left(\frac{s}{\sum_{i=a}^{n} c_i} \right)$$

The basic idea of normalization holds that comparability between theories, particularly absent unifying metaethical information, might be established through normalizing the theories' choiceworthiness assignments to a mathematical feature. FSCD conducts such normalization to the sum of assigned choiceworthiness.

The intuition in favour of normalization generally flows from the challenges outlined in I, and is widely shared in the literature (e.g. MacAskill et al. 2020b, Lockhart 2000, Sepielli 2013 etc.).

The intuition in favour of normalization to the sum specifically can be captured via the analogy of the Moral Parliament (Bostrom 2009). According to this analogy, decision-making under moral uncertainty between different theories of morality can be conceptualised as a parliament concerned with making a moral decision; with every theory represented by a faction, and the relative sizes of factions determined by the agent's credence in theories. The challenge of moral uncertainty is to make decisions as this parliament would make them. One core feature of this analogy is the proportionality of credence and faction size, i.e. votes, in the parliament: How many votes a theory gets in determining the outcome is determined by the agent's credence in it. Normalisation to the sum specifically captures this intuition: If the sum is fixed, then the only factor determining vote count is credence. If choiceworthiness was normalized to any other variable, the sum of votes might accordingly vary. Then, some theories would get more votes than others for reasons other than comparative credence, which runs counter to the notion of a fair parliament.[5]

IV. Merits of Fixed-Sum Choiceworthiness Distribution

Having developed the approach, I will now revisit the desiderata developed above – theory representation, theory comparison, stake sensitivity and equal say - and demonstrate how each of them is successfully met by FSCD.

[4] This provision might apply to some hypothetical interval theories producing partly negative transformations; those may accommodate this provision via positive affine transformation.
[5] Newberry & Ord (2021) instead advance a game-theoretic interpretation. I don't claim that FSCD is the only valid formalisation of the parliamentary approach, just that it is motivated by similar intuition.

1. Theory Representation

Starting with the desideratum of theory representation and following the distinctions introduced above, I discuss how FSCD represents ordinal-level and interval-level moral theories.

a. Ordinal Theories

Take an ordinal moral theory that aims to assign choiceworthiness to options a through d. a is obligatory, b and c are permissible, and d is impermissible; therefore $a > b; a > c; a > d; b = c; b > d; c > d$. Such a ranking is easily captured via the voting method. However, the resulting assignment is underspecified: For $s = 100$, for instance, both $d = \{(a, 98), (b, 1), (c, 1), (d, 0)\}$ and $d' = \{(a, 26), (b, 25), (c, 25), (d, 24)\}$ are consistent with the theory's preferences. Thus, the choiceworthiness assignment of ordinal theories under FSCD is prima facie underspecified.

For some purposes, e.g. developing a maximally theory-representative model even at costs to applicability, such underspecification might be desirable. However, for other purposes, e.g. of practical application without much auxiliary research on cardinal representations of ordinal choiceworthiness assignment, FSCD needs to be more specific: If there are multiple equally permissible choiceworthiness assignments for ordinal theories, this ambiguity might likewise result in multiple equally valid aggregate moral judgements; and so MEC would not yield a comprehensive action guidance. Such underspecification can be avoided by introducing auxiliary specification. I suggest the Borda Assignment Provision:

> *Borda Assignment Provision*: (a) all choiceworthiness intervals should be equally large, and (b) the least choiceworthy outcome or outcomes should be assigned 0 choiceworthiness:
>
> (a) If $c_x > c_y \wedge c_y > c_z$, then $(c_x - c_y) = (c_y - c_z)$
> (b) If $\forall c_n : c_z \leq c_n$, then $c_z = 0$

It is possible to accept (a), but not (b). This would mean that the representation of ordinal theories remains underspecified, or multiply realisable, but might be attractive to reflect stakes in a decision (s. discussion in IV). Given space constraints, I however tentatively endorse accepting both parts of the provision to ensure the calculus yields one specific distribution and to build on the parallel to the Borda Count described below.

Then, one provision-compliant representation of the theories' preferences would be $d'' = \{(a, 50), (b, 25), (c, 25), (d, 0)\}$, another might be $d''' = \{(a, 5), (b, 2.5), (c, 2.5), (d, 0)\}$ - but all compliant representations would yield the same distribution if normalized to the same s; for instance, under $s = 50$, both would be transformed to $D = \{(A, 25), (B, 12.5), (C, 12.5), (D, 0)\}$.

Any ordinal theory with complete and transitive preferences can arrive at such a distribution only based on transitive complete preferences, FSCD and the provision: Suppose r is the ordinal rank of an option, with n being the highest-possible rank and 0 the lowest-possible rank. An option is ranked $r = n$ if no other options are preferable to it; it is ranked $r = n - 1$ if only options ranked $r = n$ are preferable to it; it is ranked $r = n - 2$ if only options ranked $r = n - 1$ and $r = n$ are preferable to it; and so on. Any transitive and complete ordinal ranking of options should thus be able to determine r for any option (s. e.g. Young 1974). It follows from accepting both elements of the provision that FSCD-normalized choiceworthiness for a given option y is then determined based on the normalization formula, with the rank functioning parallel to a pre-normalization choiceworthiness value.

$$C_y = r_y \left(\frac{s}{\sum_{i=a}^{n} r_i} \right)$$

The intuitive drive for this provision mirrors MacAskill's work on the suitability of the Borda Count as a suitable aggregation approach for non-intertheoretically-comparable ordinal theories (MacAskill 2014 73-87). Distributions under the provision mirror Borda Count's rank-based assignment of equidistant scores. Accordingly, aggregating two ordinal theories under FSCD given the provision amounts to conducting a Borda Count.[6] Following MacAskill et al.'s (2020, 59-63) general arguments in favour of Borda Count-based solutions, this should motivate FSCD and the provision.

This endorsement of the Borda Count for aggregating ordinal theories does not conflict with my previous argument in favour of cumulative vote – Borda as detailed above satisfies, if overspecifies, the demand of cumulative vote. It however fails to capture additional information provided by interval and ratio views (s. also MacAskill 2014, 89), which is why I – with MacAskill et al. (2020) - endorse a Borda-like method for ordinal theories, but reject it as an overarching approach to MEC.

Of course, when the results of the choiceworthiness distribution established via the Borda Assignment Provision are aggregated with other theories, the provision constitutes a potentially objectionable upscaling of a merely ordinal scale to a cardinal scale. For now, I simply aim to state the approach. I later discuss potential issues, including those arising from upscaling.

b. Cardinal Theories

Cardinal, i.e. ratio and interval, theories are easily captured via the normalization method; whatever choiceworthiness values the cardinal theory has assigned, they can simply be normalized and will then fulfil the conditions of FSCD. For instance, assume a prioritarian theory arrives, by some method of transformation and aggregation of welfare values, at a score-based ranking, wherein scores are assumed

[6] Specifically, a version of the Borda Count that allows for ties, i.e. multiple options at the same rank (Young 1974). This seems advisable in the context of moral uncertainty, where options might sometimes be equal in choiceworthiness. This might however lead to 'top-heaviness'– s. Objections.

to track choiceworthiness: $p = \{(a, 3), (b, 4), (c, 5), (d, 2)\}$. Normalising these values to $s = 50$, we get the choiceworthiness distribution $P = \{(a, 10), (b, 13.33), (c, 20), (d, 6{,}66)\}$. Or take a consequentialist theory under which choiceworthiness strictly corresponds to overall aggregate welfare brought about by an option: Given a welfare assignment of $c = \{(a, 2000), (b, 13000), (c, 3000), (d, 2000)\}$, and applying the transformation, we get $C = \{(a, 5), (b, 32.5), (c, 7.5), (d, 5)\}$.

A hypothetical cardinal, but purely interval[7], theory of morality might describe the same choice situation in different, equally valid choiceworthiness assignment. This apparent underspecification serves two important purposes: It allows FSCD to avoid values of $c_x/C_x < 0$, which would otherwise allow choiceworthiness assignments for single options well beyond s; and it allows such theories to express their understanding of a theories' stakes via choosing transformations with higher or lower ranges – extended discussion of this point follows. Beyond that, such rare theories are free to choose whatever valid transformation. I discuss why this causes no difficulties from tactical motivation later.

2. Theory Comparison

The challenge of intertheoretic comparison requires a justification for and a practical approach to aggregating different theories. Two major branches of answers to this challenge have emerged:

The first branch conducts metaethical investigation and derives an approach to representation and comparison from it by making claims about theories' value functions and how they compare; e.g. on the basis of common grounds between theories (Ross 2006, Sepielli 2009); or the assumption of a universal scale (MacAskill 2014). Two principal reasons for taking on this high metaethical burden have been proposed: Because one would want to make a substantive metaethical point on shared features between moral views (Ross 2006, MacAskill 2014, Riedener 2019); or because one thinks all other approaches to comparability lead to even greater practical difficulties (MacAskill et al. 2020). Neither of these reasons affect to my approach: I only aim to suggest a practical approach to provide action-guidance under normative uncertainty without any attached metaethical claims, and I believe the following discussion demonstrates that the absence of metaethical comparability does not impose practical issues.

Accordingly, I follow a second class of approaches that takes a more modest metaethical view and remains agnostic about metaethical comparability. Views of this class aim to instead design systems that can effectively address disagreement even given potential incomparability. In this, they draw from a parallel to social choice theory, where voting methods are often used to enable aggregation where one is agnostic or sceptical about interpersonal comparability. Lockhart's (2020) Principle of Equity among

[7] I take this to be a hypothetical edge case – I, with MacAskill et al. (2020) assume most cardinal theories can provide some information on global or local zero-points or ratios, rendering them not purely interval-scaled.

Moral Theories PEMT or MacAskill et al.'s Borda Count for incomparable ordinal theories (MacAskill et al. 2020, 73-75) are examples of this class.

Likewise, FSCD falls under this class. As demonstrated in its description as a voting method, FSCD fulfils the structural features that social choice theory requires to make aggregate choices from directly incomparable preferences. FSCD does not claim that the moral judgements of theories are comparable any more than a democratic vote claims that the political preferences of one citizen are directly comparable to another. The single-issue zealot voting for their party from deep ideological conviction and the political candidate's loyal family member voting for them don't make substantially comparable judgements – but the universal democratic vote enables aggregation of their preferences relating to who should get a seat in parliament. Similarly, the Kantian view and the utilitarian view might make judgements with differing metaethical contents – but FSCD, as a voting method, provides a shared aggregation method relating to what option the agent should choose under moral uncertainty.

Practically, then, the choiceworthiness values assigned as detailed above and weighted by the credence distribution can be aggregated through MEC. Here, the merit in the convergent expression of FSCD as a voting method and a normalization method becomes apparent: As observed by MacAskill et al. (2020), ordinal theories fit voting methods well, while cardinal theories are better represented via normalization methods. Since FSCD's voting and normalization method are convergent - i.e. yield identical choiceworthiness assignments from identical moral judgements - choiceworthiness as determined by either method can be directly aggregated with choiceworthiness determined by the other.[8]

For example, revisiting the theories described throughout this section and assuming credences of $O: 25\%; P: 40\%; C: 35\%$, we can simply conduct an aggregate judgement:

	O (25%)	P (40%)	C (35%)	Aggregate
a	25	10	5	$\Sigma_a = 12$
b	12.5	13.33	32.5	$\Sigma_b = 19.83$
c	12.5	20	7.5	$\Sigma_c = 13.75$
d	0	6.66	5	$\Sigma_d = 4.42$

3. Stake Sensitivity

In FSCD, stakes are reflected in the range of choiceworthiness. An example should elucidate this claim: Take two utilitarian views – an average view A and a total view T. A decision must be made between bringing about a world with one trillion people at 100 units of well-being each, and a world

[8] This is notably opposed to MacAskill et al. (2020)'s approach, the non-convergence within which requires conducting a separate meta-aggregation.

with one hundred people at 101 units of well-being each. T prefers the first world, given that it includes 100 trillion units of well-being as opposed to 10,100, whereas A prefers the second since it includes an average well-being of 101 as opposed to 100. Given the large difference between worlds under T and the comparably small difference under A, it seems like T has higher stakes in the decision than A.

I believe that the notion of stakes motivating this intuition is based on differences in value between outcomes: The more a theory has to gain or lose from a decision, the higher the stakes might be. We might think of a low-stakes-situation as a choice scenario in which all assigned choiceworthiness values are very close to each other so that the difference between making the right and the wrong choice in the eyes of the theory isn't particularly impactful, whereas a high-stakes-situation would be one with at least one option's choiceworthiness substantially diverging from another – i.e. the theory would have a lot to gain or lose. This notion of stakes is accordingly expressed in the range of choiceworthiness, i.e. the difference between the highest and lowest value.

FSCD reflects this understanding of stakes. Take the previous example: Understanding the total and average well-being respectively to determine T's and A's choiceworthiness distribution, a normalization to $s = 100$ yields the following results: $T = \{(a, 99.99), (b, 0.01)\}$, $A = \{(a, 50.25), (b, 49.75)\}$. FSCD correctly tracks T as higher stakes, i.e. having a much wider range of assigned choiceworthiness, than A. This difference in stakes will influence the outcome: Given even comparatively low credence in T, the massive range in T's assigned choiceworthiness will be enough to offset the comparatively miniscule range in A when determining the aggregate judgement.

Two often-discussed specific elements of stake sensitivity deserve further elucidation at this point: FSCD's treatment of absolutist theories and of drastically different stakes.

a. Absolutism

Many theories that contain (near) unlimited sensitivity to stakes struggle with absolutist moral theories, i.e. with theories that consider one option so superior to another that the difference cannot be quantified by means of finite choiceworthiness assignment – for instance, a theory might hold that it is wrong to murder, no matter what other moral goods are involved, and so the not-murdering would always be infinitely choiceworthy (Colyvan et al. 2010). Many accounts struggle with absolutism and, similarly, with almost-absolutist fanatical views, because in the interest of theory representation, they might be bound to assign choiceworthiness values high enough to swamp the entire calculus and render the aggregate verdict insensitive to all but fanatical theories. (Jackson & Smith 2006, MacAskill et al. 2020, 151-155)

FSCD has no issue with absolutist theories – by its provision of equalised sums, the most extreme assignment available is a spread of $(n, 0)$, wherein the lexicographical difference between the options is reflected by one option being infinitely as choiceworthy as another. Any preference beyond this

relative assignment, i.e. to want not only the ratio to be infinite, but also the distance in absolute terms to be higher than that assigned by other theories, seem related to preferences over the aggregate outcome – such preferences are tactical and should be dismissed in the context of MEC, as discussed later in detail. More complex absolutist views, e.g. ones including multiple tiers, are relegated to the provision of equal intervals: Since they can make no statements about the comparative size of the intervals they assign, they are treated like any other ordinal theory unable to do so. Absolutist theories slot into the FSCD framework with no irregularity and no obvious counterintuitive effects on representation or aggregation.

b. Drastically Different Stakes

Beyond absolutist edge cases, theories might simply diverge in decision-specific stakes so much that it is implausible to normalize their choiceworthiness assignments. For instance, FSCD might capture the aggregation of a deontological and an aggregate consequentialist view in low-stakes situation accurately. But when considering decisions with huge far-future ramifications, such as reducing existential risk, the stakes of some deontological theories might remain similar while the consequentialist stakes drastically increase.

As argued above, I believe that the range-based approach can make considerable stake-based differences between theories. In situations where this range-based stake sensitivity does not suffice, I believe that drastically different stakes could be better addressed by ways of credence: Insofar as the agent believes that the consequentialist view is of drastically higher relative importance to their judgement in the context of population ethics than in the context of small-scale personal decisions, this agent's credences might just diverge between types of decisions. For instance, an agent might believe that aggregate consequentialism is a more plausible view when it comes to large-scale axiology, but that some deontological view is a more plausible action-guiding principle for low-stakes everyday decisions. No requirements obligate the agent to maintain a consistent credence distribution for all levels of decision – and if one theory provides convincing reasons that it might be able to specifically address large-scale axiological issues well, it would plausibly enjoy more credence in such cases.

Insofar as neither range-based stakes nor credences can capture the intuition behind the objection, I concede: FSCD's ability to give very-high-stakes views very high amounts of say over the outcome is limited. This ensures it tracks the notion of equal say captured by the parliamentary analogy and saves it from absolutism and fanaticism worries. To my mind, this is a sensible balance to strike given the tension between equal say and stake sensitivity - but insofar as a sensitivity to stakes beyond what I described above is desired, it might motivate some to reject FSCD.

4. Equal Say

Because every theory receives the same amount of choiceworthiness to distribute, i.e. the same number of votes, FSCD establishes equal say. I believe the obvious democratic intuition supports this claim: In the context of a democracy, everyone receiving one vote to cast as they like constitutes democratic equal say. FSCD builds on that principle – it only gives its parallel to voters, i.e. moral theories, greater freedom in how to distribute this vote by allowing for its fine-grained distribution across all options. But the parallel holds – if 'one person, one vote' constitutes equal say, it seems to follow that 'one theory, *s* votes' similarly does. This tracks the above elaboration on FSCD's motivation flowing from the Moral Parliament.

Hence, FSCD can accommodate varying stakes while still maintaining a plausible account of equal say and thereby provides the MEC framework with a strong response to one of its most salient challenges.

V. Objections

I will examine the most striking past and potential objections to my approach[9], ordered as they relate to the four desiderata of theory representation, theory comparability, unequal stakes and equal say.

2. Objections to Theory Representation

a. Cardinal Upscaling for Ordinal Theories

The first objection relates to the representation of ordinal theories on the cardinal scale of FSCD. The general version of this objection is to contend that upscalings of theories are generally impermissible – that the information added by overspecifying the intervals misrepresents the theory's moral judgements (Tarsney 2019).

i. Firstly, as I have suggested above, no ordinal theory is able to leverage a constructive objection against its upscaled cardinal representation: While the ordinal theory might have a purely negative objection against its representation – i.e. it might contend that the representation makes claims about interval size the theory has not made -, it cannot claim that it would prefer another cardinal representation. Accordingly, the objection should be read not as an objection to the specific upscaling, but to upscaling in general.

One might read this defence as implying that any cardinal representation of an ordinal theory is accurate as long as it maintains the ordinal ordering. Therefore, there might be no reason to follow the provision of equal intervals. Firstly, I believe that, absent any information on interval size, equal interval size is a reasonable prior: The option space of

[9] Owing to the scope of this paper, not all technical questions put forward in MacAskill (2014; 2016) and MacAskill et al. (2020; 2020b) can be comprehensively addressed. I focus on those that strike me as strongest most relevant to the suggested application area of practical action-guidance via a voting approach.

comparative interval sizes presumably includes all negative and all positive differences, which plausibly averages to no difference. By definition of the ordinal views, no information to the contrary exists, so no information would motivate a divergence from the presumptive neutral average of cardinalisation options. In that sense, suggesting unequal intervals would add another element of arbitrariness.

Furthermore, the motivation for the Borda Assignment Provision does not have to be that the provision is overall maximally accurate theory representation. It is sufficient that the provision provides the most accurate theory representation when upscaling ordinal theories; and then to further motivate the provision by parallel to the Borda Rule as described above. Then, the objection could only hold that upscaling is generally unadvisable. But to avoid upscaling while maintaining comparability, one would have to suggest underspecification, which I have addressed above; or downscaling of cardinal views, for that, s. below.

ii. When attempting to compare cardinal and ordinal theories, the salient alternative to upscaling the ordinal numbers would be downscaling the cardinal theories. Ordinal theories might support such downscaling: An ordinal theory that ranks $a > b > c$ is going to assign values of $O = \{(a, 2), (b, 1), (c, 0)\}$ no matter if we operate under Borda or FSCD with $s = 3$. The only difference to more restrictive approaches is how free the cardinal theories are to assign choiceworthiness; under pure Borda, they are constrained to similarly distributing values of 2, 1 and 0, whereas under FSCD, other distributions are possible. This is relevant to the outcome of the aggregation process; the increased flexibility afforded to cardinal views give ordinal theories comparatively less influence. But that seems justified: Cardinal choiceworthiness assessments are genuinely more specific and are therefore reflected by more specific assignments; this is what makes them cardinal theories. So in this case, cardinal theories object on constructive grounds – they could have a more accurate choiceworthiness assignment under cardinal representation. The ordinal preference for cardinal downscaling, however, does not flow from concerns of representation, but from the ordinal preference to have a stronger voice in the final aggregation. However, as discussed later, such preferences are tactical preferences not applicable to FSCD. Hence, cardinal theories' valid objection to downscaling should be sustained.

In conclusion: The upscaling of ordinal theories per Borda provision ascribes interval-level information to ordinal scales. However, this upscaling is not particularly concerning because (a) no constructive objection from theory representation against it can be made and (b) the objection from theory comparison is an inadmissible tactical objection.

3. Objections to Theory Representation

As mentioned above, choosing a voting approach is oftentimes a decision of which concessions to make. I argue that cumulative vote's failure modes are specifically less concerning in the context of moral uncertainty, thus reducing the downsides from these implicit concessions and further motivating FSCD. This specifically relates to two objections to cumulative vote applied to FSCD:

a. Tactical Motivation

An overarching problem with voting-based approaches is that voters might engage in tactical voting, i.e. assign votes not maximally reflective of their preferences, but maximally conducive to bringing about their preferred outcome (MacAskill 2014, 71; Felsenthal 1985). This might apply to FSCD: Driven by a preference over the outcome of the calculus, theories might assign choiceworthiness more radically and less representative of their comparative preferences (e.g. assign all choiceworthiness to their favoured option so that it is chosen).

However, I believe this objection arises from a misunderstanding of theories' agency: A theory is not a conscious agent with tactic preferences over outcomes. It seems absurd to think that, within MEC theories would have an incentive to misrepresent their preferences. Theories, as beliefs held by the individual, are no dissociated separate entity from the individual that has an interest in their fair aggregation. The decision-maker uses MEC to achieve as accurate of a compound moral judgement as possible and assigns choiceworthiness to that goal - so no adversarial interests motivate single moral theories to exaggerate or lie. So tactical-voting-based objections to FSCD fail.

b. Choice Setup and Option Individuation

Another objection holds that FSCD is sensitive to the specific setup of the choice situation, e.g. to the introduction of additional options or individuation of existing options (MacAskill et al. 2020, 96-98). This runs parallel to worries of violation of 'independence of irrelevant alternatives' or a sensitivity to 'cloning attacks' in social choice (Tideman 1987).

The most striking concern here is option individuation: In S_1 the choice is between a: telling a lie and b: not telling a lie; in S_2, it is between a': telling a lie today, b': not telling a lie, and c': telling a lie tomorrow– the latter would force the consequentialist in favour of lying, but not the deontologists refusing to lie, to split up their assigned choiceworthiness between a' and c'. So S_1 might, given some credence distributions, go in favour of the consequentialist view, while S_2 might not.

This objection is weaker in MEC than in social choice: As with tactical voting, in MEC, there are no incentives to asymmetrically manipulate the choice setup to distort the calculus. However, choice setup still might matter, e.g. when individuation occurs as an attempt to describe options as accurately as possible, so the objection warrants an answer.

MacAskill et al. (2020, 99) provide this answer: Instead of simply assigning votes or choiceworthiness values to options as they are arbitrarily individuated and described, the assignment is instead tied to a probability measure assigned to each option. Conceptualising all available courses of action as a two-dimensional space, an option's measure describes the fraction of possibility space taken up by that option. It is the default assumption that all options in a well set-up choice situation occupy equal parts of the possibility space, i.e. are assigned the same measure.[10] Assuming that a and b each take up $1/2$ of the space of possibilities and are therefore both assigned a measure of 0.5, individuation can be tracked: If an option a is split into a' and c' like in moving from S_1 to S_2, the probability measure of the option space taken up doesn't change. So a's measure is split up and distributed across a' and c', who (given equal probability) both receive a probability measure of 0.25. (MacAskill et al. 2020, 96-101)

Beyond MacAskill et al.'s (2020) application to variance and Borda Count, I argue the measure approach can be applied to FSCD:

i. Individuation

Applying the measure-based approach to FSCD, the relative probability measures determine the extent to which the choiceworthiness assigned to an option counts towards the fixed sum s. For instance, since c' and a' only have half the probability measure of b', their choiceworthiness likewise only counts half as much towards s. This allows the consequentialist theory to assign just as much choiceworthiness to each of a' and c' as it would to a. Thus, if an option's portion of probability space is individuated, the new option(s) retain the same choiceworthiness value, safeguarding the calculus against such modification of the choice setup. In technical terms, the normalization formula would be amended as follows, wherein m_x is the option's measure:

$$C_x = c_x \left(\frac{s}{\sum_{i=a}^{n} c_i} \right) \times 1/m_x$$

In cases in which individuation seems obviously called for, this approach likewise establishes resistance to objectionable compound choice setups: Take individuating 'saving someone from torture' into 'saving them in two years' or 'saving them now', with both being options considered by the agent. Such a case, following MacAskill et al.'s analysis (2020, 97), would be best described not by considering the new options as specifically individuated, but by considering the old, general option as a compound and thus assigning the same measure to the two new options as to other non-compound options.

[10] The term 'probability measure' refers to the mathematical makeup of the measure, not to the probability that an option might be chosen.

ii. General Setup

Furthermore, the measure also provides a theory-impartial criterion for a well-setup choice scenario without omissions or irrelevant alternatives that don't arise from individuation: Omitting options within the possibility space is impermissible, and so is adding options outside the possibility space – i.e. a choice ought to precisely include all options considered by the agent. Accordingly, no potentially 'irrelevant' alternatives can arbitrarily be added and no relevant alternatives can be removed to the detriment of an accurate result.

I believe this argument demonstrates that the measure-based approach is a generally suitable approach to choice-setup-concerns, and that it is readily applicable to FSCD. To completely satisfy the objection, it requires further formal development beyond this paper; for now, and in aiming to generally develop and motivate FSCD, its brief introduction and application should plausibly demonstrate the availability of an effective solution. Furthermore, if the measure-based response fails, then other approaches within MEC are likewise vulnerable to the choice-setup-objection. My claim that FSCD is a particularly promising approach *within* MEC would thus still be unaffected.

4. Objections to FSCD's Stake Sensitivity

Criticism might hold that FSCD would be stake insensitive for ordinal theories or for amplified theories.

a. Ordinal Theory Stake Insensitivity

Given the Borda Assignment Provision, ordinal theories' choiceworthiness representation ranges are not influenced by stakes, but rather entirely by the number of options and ranks. Accordingly, FSCD might be objectionably stake-insensitive in representing ordinal theories.

However, the notion of situation-specific stakes does not coherently apply to ordinal views. Instead, I believe stake intuitions on ordinal views reveal misguided assumptions of cardinality. Imagine a potential high-stakes situation for an ordinal theory. For instance, one might think that a decision D (between a: violating one duty and b: violating a hundred duties) is higher stakes than decision E (between a: violating one duty and c: violating ten duties). However, these intuitions about stakes are intuitions about intervals – they carry the assumption that the difference between a and b is greater than the difference between a and c. Insofar as the moral theory makes such distinctions between choiceworthiness intervals, it seems to no longer be truly ordinal, but instead able to quantify intervals – i.e. at least partly cardinal. But if the theory is an interval theory, the ordinal-specific specifying provisions that remove stake-sensitivity no longer apply.

Even if a specific framework to accommodate stake-sensitive ordinal theories were required, FSCD could still accommodate this by applying only the first part of the Borda Assignment Provision– i.e. enforce equal intervals, but not $C = 0$ for the lowest option(s). Then, the size of the constant

choiceworthiness interval between ranks as implied by the first part is up to the moral theory; therefore, range could vary based on stakes - and stake-sensitivity thus be expressed. I take this to be a good solution for this hypothetical edge case, but not the first choice for the overall approach, since if no specific notion of stakes informs the size of the constant interval, the choiceworthiness assignment for ordinal views would remain underspecified.

b. Amplified Theory Stake Insensitivity

Moral views might be amplified versions of other moral views. For instance, one consequentialist view C might be hedonic utilitarianism, which considers the overall compound welfare of all humans relevant, while another view C' might consider overall consequences to human welfare just as relevant, but additionally assign additional value to the welfare of the agent's kin. In decisions exclusively regarding the agent's kin, C's and C''s judgement would not diverge in relative choiceworthiness assignment; yet C' has more expansive reasons for their assignment than C given the additional kinship-based reasons - here, C' is an amplified version of C (MacAskill et al. 2020, 125-130). In their most extensive point of criticism to approaches like FSCD, MacAskill et al. argue that in such cases, an amplified theory should be allowed to distribute more choiceworthiness than its non-amplified version.

I however believe amplified theories can instead be addressed by more precise theory individuation. As it is described above, C' is a compound theory that includes both the hedonic utilitarian view and a motivation to care about welfare arising from kinship. We might therefore understand C' as follows: It consists of one function simply linearly mapping human welfare to choiceworthiness, i.e. C; and additionally, a different function D mapping human welfare to choiceworthiness that is entirely dependent on degree of kinship. Or, ex negativo, D is C' without the belief in overall hedonic utilitarianism. Adding C and D results in a choiceworthiness assignment that assigns just as much choiceworthiness based on overall welfare as C; and that additionally assigns some choiceworthiness based on kinship-amplified welfare. This fits the initial description of C'.

So, the credence distribution of the agent could be described as having some credence in C and some credence in D instead of as having some credence in C and some credence in C'. And for some distribution of credences (depending on the exact relative strength of moral reasons from the components in C'), individuating the moral theories at play into C and D instead of C and C' would yield the same compound results. But in this case, two distinct views with no amplificatory relationship, and thus no dilemma, exist. If one individuates theories so that additive compound theories like the speculative C' are separated into non-compound components, amplified theories are no issue.

An objection to my response might hold that C' could not be a compound theory, but instead a completely different view that incidentally overlaps with C – but draws on different reasons; for instance, C's moral stake in overall welfare could be motivated by divine command, and C''s stake in

overall welfare be an imperative of reason. In such a case, it is inaccurate to say C' consists in part of C, so my suggested individuation is impermissible.

However, in such a case, C' is no longer an amplified version of C. If C and C' are metaethically different, it is not clear that C' should be higher stakes than C. The difference in metaethical features might give rise to inconsistencies in stakes: For instance, it might be possible that C has higher stakes because disobeying the divine command is more serious than disobeying the law of reason. Hence, the relationship between C and C''s stakes then depends on other features of the theories – so MacAskill et al.'s contention that an amplified theory needs to be higher-stakes does not hold.

If amplified theories are compound theories, they can be addressed through precise theory individuation; if amplified theories differ metaethically, they don't carry worrisome stake implications.

5. *Objections to FSCD's Equal Say*

FSCD's notion of equal say might be subject to criticism in cases of mixed aggregation or of top-heaviness.

a. Mixed Aggregation

The Borda Assignment Provision might be understood to cause an unfair imbalance in cases of mixed aggregation between ordinal and non-ordinal theories: While non-ordinal theories might vary in stakes, ordinal theories will, by virtue of the provision, assign choiceworthiness at the maximum possible range. Following the range-based stake definition, this means ordinal theories always assign choiceworthiness as if operating under maximal stakes. Given equal credence, this might give comparatively more say to ordinal theories when cardinal theories have non-maximal stakes in the decision.

Specifically, this might lead the critic to suggest that, for the sake of intertheoretic fairness, cardinal theories should likewise be constrained to assign 0 choiceworthiness to their least choiceworthy option. I reject this suggestion: There is no good reason from theory representation to force cardinal views to assign 0 choiceworthiness to the lowest option as well - in fact, it is harmful to theory representation insofar as it limits range-based stake representation. There might, however, be good reason to force ordinal theories to do so – the fixed point ensures a unique choiceworthiness distribution and intuitively flows from the Borda Count.

I instead argue that this asymmetric treatment of ordinal and cardinal theories by FSCD is often justified and can furthermore be addressed through the credence distribution:

Stake-insensitive ordinal views seem to make no range-quantifiable difference between the stakes in different choice situations – and therefore genuinely do treat every situation as having equally high stakes. One could lower these fixed equally high stakes across the board, e.g. by assigning not 0, but a

higher number of choiceworthiness, to the lowest outcome – but that might likewise be considered unfair because then, ordinal theories would be assumed to have lower stakes than some high-stakes cardinal views. No single range or fixed lowest point seems to capture the nondescript stakes of a stake-insensitive ordinal view without causing fairness issues when compared to either higher- or lower-stakes cardinal views. No concrete assignment is exclusively correct – but all else equal, my proposition is motivated by practicality and the parallel to Borda.

Instead, the issue can be addressed at the level of the credence distribution: The issue at the heart of the unfairness picked up by the objection is the ordinal views' commitment to equal stakes – and the consequences for fair aggregation that follow from it. If the agent finds this implausible, this might simply mean they find theories that make no difference between choice scenarios less plausible than more nuanced theories - which can be addressed via credence: Based on this specific issue with stake-insensitive ordinal theories, agents might think them less plausible, assigning comparatively less weight to them than to cardinal theories. Thereby, the balance between views is adjusted to counteract the objectionable unfairness.

b. Top-Heaviness and Equal Say

MacAskill et al. (2020b) reject theories like FSCD because they argue they violate equal say when aggregating top-heavy and bottom-heavy theories. In the context of a given choice situation, a top-heavy theory refers to a theory that describes many options as fairly choiceworthy and only few as comparatively unchoiceworthy, whereas a bottom-heavy theory is the inverse, describing only very few of the options as choiceworthy. See below, where X is bottom-heavy and Y is top-heavy. (MacAskill et al. 2020b, 87-90)

	X	Y
a	1	33
b	1	32
c	1	32
d	97	3

In FSCD, such divergences can arise in the context of the Borda Assignment Provision: A theory that has many ties for a high r might be understood as top-heavy, whereas a theory that has many ties for a low r might be understood as bottom-heavy.

This might be concerning: It seems like X has a higher chance of getting its way than Y, because it assigns choiceworthiness in a way that is more likely to beat out a hypothetical third theory's

preferences – the option *X* deems most choiceworthy is described as almost three times as choiceworthy as *Y*'s best option. I do not find this contention particularly concerning for two reasons.

Firstly, the examples motivating this notion of equal say suppose that MEC generally intends to pick out exactly one advisable option. This is not the only relevant application; for example, imagine a situation in which an agent is using MEC to figure out a ranking of available options, but there are non-moral reasons for them to reject one of the options under many aggregation outcomes – it might be inconvenient, terribly expensive, etc. If these non-moral reasons disqualify the bottom-heavy theory's favourite option, the theory is left with close to no influence on the actual outcome yet if the same happens to the top-heavy theory, it is left with a large remainder of its original choiceworthiness assignment to influence the decision. In establishing the second-best option, *X* only assigns 1 unit of choiceworthiness, whereas *Y* assigns 32, and similar for every *n*th-best option. In more general terms: Whenever an agent uses MEC not to pick out a single (or correspondingly, very few) option, but instead to pick out several advisable options or to arrive at an overall ranking, the bottom-heaviness of a theory is less disadvantageous; and it might even be beneficial to be top-heavy instead. It is selective framing of MEC as exclusively concerned with the former that makes the top-heaviness issue appear unproportionally concerning.

Secondly, having less say in cases solely aimed at picking out one advisable option might be an accurate representation of what it means to be top-heavy. Being bottom-heavy might be taken to mean 'specifically wanting one thing', whereas being top-heavy might be taken to mean 'somewhat wanting many things' – this is what finding one option particularly choiceworthy as opposed to finding many options somewhat choiceworthy means. And it does not strike me as objectionable that, when aiming to pick out exactly one option, a theory that has a stronger positive, substantial opinion on what that one option should be – i.e. the bottom-heavy theory – should have more say. The bottom-heavy theory is less indifferent regarding what the best option is, so it assigns it votes accordingly and receives corresponding effective say.

In conclusion, the top-heaviness objection is mostly limited to cases where only one (or, to an extent, very few) options must be chosen; and in these cases, it seems consistent with a plausible reading of the practical content of bottom-heaviness to let bottom-heavy theories have more effective say. On this view, MacAskill et al. (2020b)'s objection fails.[11]

[11] Moreover, since their approach of variance normalization specifically withholds this additional effective say from bottom-heavy theories, the failure of this objection might be of concern to variance normalization's availability to ensure equal say – but this paper does not afford the space to make the comprehensive comparative case against variance normalization.

VI. Comparison

I conclude with a brief overview of how the described merits position FSCD in comparison to two prominent approaches. This is no exhaustive comparative discussion; this paper's aim is to independently motivate FSCD as one viable solution with unique merits, not conclusively prove its universal superiority. Accordingly, the following comparison merely aims to give an overview over the trade-offs involved in choosing between FSCD and some rival approaches.

1. MacAskill, Ord and Bykvist's Approach

MacAskill et al. (2020) suggest a three-part model: Aggregation of ordinal theories via Borda Count, aggregation of fully comparable cardinal theories per unaltered MEC, and aggregation of non-comparable cardinal theories per variance normalization. Either the choiceworthiness values assigned by these respective approaches or the separate judgements of these three aggregation methods are then further aggregated into the compound judgement. Comprehensive piece-by-piece comparison to this extensively developed approach goes beyond the scope of this paper – what follows is a brief overview.

a. Metaethical Demandingness

MacAskill et al. (2020, 141-147) base their notions of comparability, particularly among closely related moral theories such as amplified theories, on the metaethical supposition of a universal scale of choiceworthiness. Therefore, their approach makes a much broader claim of actual intertheoretic comparison. FSCD makes no such claims – and is therefore unsuitable to non-action-guiding applications of moral uncertainty. However, FSCD is also not susceptible to the wealth of criticism leveraged at MacAskill et al.'s approach's metaethical grounding – such as those by Hedden (2016), Riedener (2019), Carr (2020, 2021), and shared by Tarsney (2018) and Aboodi (2022). I believe my response to MacAskill et al.'s contentions demonstrates that, for practical action-guiding applications, choiceworthiness assignment approaches can avoid the high burden of proving metaethical intertheoretic comparability and the numerous objections that come with it.

b. Complexity in Implementation

MacAskill et al.'s view requires both three different methods of aggregation and a method of meta-aggregation, either via variance normalization or via aggregating ordinal sub-outcomes (MacAskill et al. 2020, 108-11). FSCD is comparatively light-weight and applies a uniform pair of methods to all levels of theory without meta-aggregation. While this come at the cost of some edge-case precision, I believe FSCD's more straightforward nature lends it to greater ease of use for action-guiding applications.

c. Treatment of Equal Say and Stakes

MacAskill et al. (2020, 86-90; 2020b) take variance normalization to best capture equal say and apply it to some theory comparisons. FSCD's notion of equal say as equal votes is more constraining and more extensive. As discussed above, there is a trade-off here: FSCD is comparatively more readily able to handle fanaticism and absolutism, but less sensitive to potentially desirable tracking of particularly large stake differences. As argued above, I believe this trade-off favours FSCD.

2. Structuralism

Structural solutions assume intertheoretic comparability based on shared structural features; similarly, solutions I term quasi-structural aggregate theories absent metaethical intertheoretic comparability based on structural features. FSCD is such a quasi-structural solution; depending on interpretation, Lockhart's (2000) PEMT is either structural or quasi-structural as well.

I believe my argument implies there is general merit to quasi-structural approaches, and some of my responses to potential criticism at least partly apply to objections levelled against other quasi-structural as well. However, my arguments regarding the Moral Parliament, equal say and range-based stake representation are all specific to FSCD and don't lend themselves to other (quasi-)structural views – so insofar these arguments succeed in motivating my approach, they suggest that FSCD is comparatively preferable.

VII. Conclusion

Fixed-Sum Choiceworthiness Distribution prescribes and normalizes moral theories' choiceworthiness assignments to the sum of choiceworthiness distributed across options, thereby fulfilling the technical desiderata of an aggregation method in maximizing expected choiceworthiness. I argued that FSCD is able to express stake-sensitivity through the range of choiceworthiness, captures a relevant notion of equal say through providing equal votes, and successfully tracks common intuitions in challenging cases of theory aggregation. Accordingly, FSCD should be considered a viable approach to choiceworthiness assignment and intertheoretic comparison in MEC deserving of further consideration.

References

Aboodi, R. (2022). Normative Uncertainty without Unjustified Value Comparisons. *Journal of Ethics and Social Philosophy 21*(3).

Bostrom, N. (2009). *Moral Uncertainty: Towards a Solution?* https://www.overcomingbias.com/2009/01/moral-uncertainty-towards-a-solution.html

Carr, J. (2020). Normative uncertainty without theories. *Australasian Journal of Philosophy 98*(4), 747–762.

Carr, J. (2021). The hard problem of intertheoretic comparisons. *Philos Stud* 179 1401–1427.

Colyvan, M., Cox, D., & Steele, K. (2010). Modelling the Moral Dimension of Decisions. *Noûs 44*(3), 503-529.

Gracely, E. (1996). On the Noncomparability of Judgments Made by Different Ethical Theories. *Metaphilosophy 27*(3), 327-332.

Gustafsson, J., & Torpman, O. (2014). In Defence of My Favourite Theory. *Pacific Philosophical Quarterly 95*(2), 159-174.

Gustafsson, J. (2022). Second Thoughts about My Favourite Theory. *Pacific Philosophical Quarterly 103*(3), 448-470.

Felsenthal, D. (1985). Is cumulative voting really different from one-man, one-vote? *Electoral Studies 4*(2), 141-148.

Hedden, B. (2016). Does MITE Make Right?: On Decision-Making under Normative Uncertainty. *Oxford Studies in Metaethics 11*, 102-128.

Jackson, F. & Smith, M. (2006) Absolutist Moral Theories and Uncertainty. *The Journal of Philosophy 103*(6), 267-283.

Lockhart, T. (2000). *Moral Uncertainty and its Consequences*. Oxford University Press.

MacAskill, W. (2014). *Normative Uncertainty*. Dissertation. Oxford University.

MacAskill, W. (2016). Normative Uncertainty as a Voting Problem. *Mind 125*(500), 967-1004.

MacAskill, W., Bykvist, K. & Ord, T. (2020). *Moral Uncertainty*. Oxford University Press.

MacAskill, W., Cotton-Barratt, O. & Ord, T. (2020b). Statistical Normalisation Methods in Interpersonal and Intertheoretic Comparisons. *Journal of Philosophy 117*(2), 61-95.

Myerson, R. (1995). Analysis of Democratic Institutions: Structure, Conduct and Performance. *Journal of Economic Perspectives, 9*(1), 77-89.

Newberry, T., & Ord, T. (2021). The Parliamentary Approach to Moral Uncertainty. *FHI Technical Report 2021-2*. Oxford: Future of Humanity Institute.

Nissan-Rozen, I. (2015). Against Moral Hedging. *Economics and Philosophy 31*(3), 349–69.

Pacuit, E. (2019). Voting Methods. In: Zalta, E. (ed.) *The Stanford Encyclopedia of Philosophy* (Fall 2019 Edition).

Pivato, M. (2013). Social welfare with incomplete ordinal interpersonal comparisons. *Journal of Mathematical Economics 49*(5), 405-417.

Riedener, S. (2019). Constructivism about Intertheoretic Comparisons. *Utilitas 31*(3), 277-290.

Riedener, S. (2015). *Maximising Expected Value Under Axiological Uncertainty. An Axiomatic Approach*. Dissertation, Oxford University.

Ross, J. (2006) Rejecting Ethical Deflationism. *Ethics 116*(7), 742–768.

Sepielli, A. (2009). What to do when you don't know what to do. *Oxford Studies in Metaethics 4*, 5-28.

Sepielli, A. (2013). Moral Uncertainty and the Principle of Equity among Moral Theories. *Philosophy and Phenomenological Research* 86(3), 580-589.

Tarsney, C. (2018). Intertheoretic Value Comparison: A Modest Proposal. *Journal of Moral Philosophy 15*(3), 324-344.

Tarsney, C. (2019). Normative Uncertainty and Social Choice. *Mind 128*(512), 1285-1308.

Tenenbaum, S. (2017). Action, Deontology, and Risk: Against the Multiplicative Model. *Ethics 127* (3), 674-707.

Tideman, T. (1987). Independence of Clones as a Criterion for Voting Rules. *Social Choice and Welfare* 4(3), 185–206.

Young, H. (1974). An Axiomatization of Borda's Rule. *Journal of Economic Theory 9*(1), 43–52.