

Vive la Différence? Structural Diversity as a Challenge for Metanormative Theories

Christian J. Tarsney*

Penultimate version; final version forthcoming in *Ethics*.

Abstract

Decision-making under normative uncertainty requires an agent to aggregate the assessments of options given by rival normative theories into a single assessment that tells her what to do in light of her uncertainty. But what if the assessments of rival theories differ not just in their content but in their structure—e.g., some are merely ordinal while others are cardinal? This paper describes and evaluates three general approaches to this “problem of structural diversity”: *structural enrichment*, *structural depletion*, and *multi-stage aggregation*. All three approaches have notable drawbacks, but I tentatively defend multi-stage aggregation as least bad of the three.

1 The problem of structural diversity

How should an agent decide what to do when she is uncertain about basic normative principles—for instance, uncertain whether Kantianism or utilitarianism is the correct moral theory? This question has become the focus of a substantial and growing philosophical literature (initiated by Lockhart (2000), Ross (2006a,b), Guerrero (2007), and Sepielli (2009, 2010), among others). Most of the approaches to choice under normative uncertainty that have been proposed in this literature involve some form of *intertheoretic aggregation*—some method of combining assessments of options delivered by rival normative theories into a single assessment that tells the agent what to do in light of her uncertainty. The paradigm example of an intertheoretic aggregation rule is “maximize expected choiceworthiness” (MEC), which takes

*Global Priorities Institute, University of Oxford. Email: christian.tarsney@philosophy.ox.ac.uk

a probability-weighted sum of the choiceworthiness (value, rightness, etc.) of each option according to each normative theory, and instructs the agent to choose an option for which this sum is maximal.¹ But several other aggregation rules have been suggested in the literature, including “My Favorite Option” (MFO), which tells agents to choose the option that is most likely to be maximally choiceworthy or permissible,² various ordinal aggregation rules borrowed from voting theory (Nissan-Rozen, 2012; MacAskill, 2016), and stochastic dominance (Tarsney, 2018b).

Any approach to intertheoretic aggregation, however, faces the following challenge: Normative theories can disagree not just in the content but also in the *structure* of their assessments. For instance, some theories assess options on an interval or ratio scale, while others merely preorder options as more or less choiceworthy, while still others simply *classify* options as permissible or impermissible. It often seems, though, that the aggregation rules that are most appropriate for assessments with one kind of structure are clearly inappropriate—or simply inapplicable—to assessments with another kind of structure. For instance, MEC is inapplicable to theories that merely preorder options. On the other hand, MFO is *applicable* to theories that provide cardinal assessments, but seems clearly suboptimal, since it takes no account of that cardinal information. How, then, can we aggregate the assessments of theories with diverse structures, if no single aggregation rule is well-suited to all possible structures? I will call this the *problem of structural diversity*.

The problem of structural diversity has been recognized in the literature on normative uncertainty (e.g. by Sepielli (2010, pp. 38ff) and MacAskill (2014, pp. 117–122)) and, as we will see, several philosophers have implicitly or explicitly advanced solutions to it. These proposed solutions, however, have serious drawbacks, and

¹Expectational principles (stated variously in terms of *expected value*, *expected rightness*, or *expected choiceworthiness*) have been defended by Lockhart (2000), Ross (2006a), Sepielli (2010), MacAskill (2014), Riedener (2015), and MacAskill and Ord (forthcoming), among others.

²This rule is discussed (but not endorsed) by Lockhart (2000) and Gustafsson and Torpman (2014), among others.

have mainly tried to address one special case: the aggregation of theories with simple cardinal and ordinal structures. In this paper, I aim to offer a more general and systematic treatment that both demonstrates the scope of the problem and evaluates potential solutions more thoroughly than has so far been attempted.

In §§2–3, I introduce the problem of structural diversity more carefully, arguing in particular that the variety of possible structures a normative theory can exhibit—and hence the scale of the challenge for intertheoretic aggregation—is much greater than has so far been acknowledged. In §4, I survey a wide range of possible responses to the problem, of which I will focus on three: *structural enrichment*, *structural depletion*, and *multi-stage aggregation*. In §§5–7 I consider each of these approaches in turn. While all three have serious drawbacks, I conclude with a tentative defense of multi-stage aggregation as the least bad approach on offer.

2 What is a “structure”?

What is the “structure” of a normative theory—or, more importantly for our purposes, what does it mean to say that two theories have “the same structure” or “different structures”? It is surprisingly difficult to give a general, precise, and extensionally reasonable answer to this question. But in this section, I’ll try to say—albeit in a very rough-and-ready way—how we might understand these notions, in order to provide a framework and point of reference for later sections. Along the way, I’ll introduce some useful terminology and notation.

First, a few basic concepts: An *assessment* is some normative evaluation, ranking, scoring, etc, of a set of practical options. It is *normative* in the sense that it pertains to an agent’s reasons—e.g., it says *how much reason* an agent has to choose a given option, or whether she has more/less/equal/incomparable reason to choose one option as compared to another. I choose the term “assessment” (rather than, say, “ranking”) because it does not connote any particular structure, and hence does

not prejudge the question of what structures a normative theory can exhibit.

An *aggregation rule* is a function or procedure that takes as input two or more assessments of the same set of options, with probabilities or other weights attached, and delivers as output a single assessment of those options that represents a weighted combination of the input assessments. Aggregation rules are distinct from *decision* rules, in that the output of aggregation need not represent a final, all-things-considered prescription—it may incorporate only some of the decision-relevant considerations or possibilities in a given choice situation.

A *theory* is a maximal consistent set of assessments in terms of some normative concept, like *objective reasons*, *moral obligation*, or *utility*. Different kinds of theories correspond to different normative concepts. In this paper, our primary interest will be in *first-order normative* theories (hereafter simply “normative theories”) like Kantianism and utilitarianism. These theories assess options based on what we may call “empirical-belief-relative” reasons—that is, their assessments are sensitive to an agent’s empirical beliefs, but not (directly) sensitive to her normative beliefs. What makes them *theories* is that they leave no room for further assessments of the same kind—they purport to settle all questions concerning empirical-belief-relative reasons.³

Our primary focus will be on the process by which the assessments of normative theories are aggregated to generate the assessments of *metanormative* theories—i.e., theories of choice under first-order normative uncertainty. Metanormative theories like MFO or MEC assess options in terms of what we might call “fully subjective” reasons, which take account of the agent’s normative as well as empirical beliefs. I will assume that these metanormative assessments are *rational* assessments—that

³Though I use the word “theory,” I mean to include both “theoretical” and “anti-theoretical” normative worldviews. Even the most anti-theoretical picture of the normative world can be expressed as a maximal consistent set of assessments—those assessments will just be relatively piecemeal, difficult to axiomatize, perhaps involve widespread incomparability, etc.

is, that we are ultimately interested in the question of how to act *rationally* in the face of normative uncertainty.⁴

It will be useful to describe theories a little more formally. A normative theory T_i can be represented as an ordered quadruple $\langle \mathbf{O}, \sigma_i, I_i, f_i \rangle$. \mathbf{O} is the set of all possible practical options. Each option $O \in \mathbf{O}$ can be understood as a vector of non-normative properties, specified indexically in relation to an agent (e.g., “is believed [by the agent] to violate a promise”). σ_i is a set of normative properties and relations (e.g., “is more choiceworthy than”), and I_i is an interpretation function that specifies the extension of those properties and relations within \mathbf{O} . Finally, f_i is a choice function mapping each non-empty subset $S \subseteq \mathbf{O}$ (a *choice situation*) to a non-empty subset of itself (the *choice set* of theory T_i in choice situation S), in a way that supervenes on the σ_i -properties and relations of the options in S .⁵

We are now ready to give a serviceable, though lamentably imprecise, characterization of the sort of “structural diversity” that poses a challenge for metanormative theorizing. First, say that aggregation rule R *applies* to normative theory T_i iff, in all (or nearly all) choice situations, it is possible for R to yield a non-trivial output assessment given inputs assessments that include T_i ’s. For instance, MEC applies to simple cardinal theories since, under otherwise favorable conditions, it can take their assessments as input and generate a non-trivial output assessment (an assessment

⁴This framing of metanormative questions in terms of rationality is common in the literature on normative uncertainty, but not uncontroversial. For a defense, see Bykvist (2013). For the rival view that decision-making under normative (or at least moral) uncertainty is essentially a *moral* question, see Rosenthal (2019).

⁵I gloss over some difficult questions here: e.g., by what non-normative properties should possible options be individuated? And should σ_i include only “primitive” normative properties and relations, or should it include all formulae that can be constructed (with some specified set of logical tools) from those primitive properties and relations? I don’t have fully worked-out answers to these questions, but I set them aside since, as far as I can tell, they don’t impinge too much on the central questions of this paper.

that makes distinctions between the options being assessed). But it does not apply to a theory that merely preorders options, since a position in a preorder is not the sort of thing that can be multiplied by a probability, so attempting to apply MEC will yield an output assessment that is simply undefined for every option.

Next, say that aggregation rule R is *sensitive* to theory T_i if its output assessments are sensitive to the information provided by T_i 's input assessments. That is, if there are σ_i -properties or relations that systematically make no difference to R 's output assessments, then R is—at least to some extent—insensitive to T_i . So, for instance, MFO *applies* to theories that provide cardinal information, but is largely *insensitive* to those theories in that it is insensitive to the information they provide concerning *how much* a given option falls short of maximal choiceworthiness.

Finally, say that theory T_i is *covered by* aggregation rule R iff R applies and is sensitive to T_i . And say that two theories T_1 and T_2 have *relevantly similar structure*, relative to a set of aggregation rules \mathbf{R} , iff there is no aggregation rule in \mathbf{R} that covers T_1 and not T_2 , or vice versa.

Though quite rough, this notion of structural similarity does what we want—it identifies the specific challenge that the structural diversity of normative theories poses to metanormative aggregation. In particular, it reflects the fact that whether a given form of structural diversity poses an obstacle to aggregation depends on which aggregation rules we find antecedently plausible in a metanormative context. The problem of structural diversity, we might say, is *the problem of how to aggregate the assessments of rival theories under normative uncertainty, given that most plausible aggregation rules cover only some of the theories in which an agent might have positive credence*.

3 What structures are possible?

Having said at least roughly what it *means* for normative theories to differ in their structure, the next question is: In what specific ways *can* normative theories differ in their structure? And in particular, in what ways do *plausible* normative theories differ in their structure, that affect the coverage of plausible aggregation rules?

In the literature on normative uncertainty, insofar as the problem of structural diversity has been acknowledged, the focus has been on the distinction between ordinal and cardinal theories—typically on the assumption that an expectational aggregation rule is appropriate for all cardinal theories, that ordinal theories are the major alternative, and that ordinal theories should either be required to satisfy axioms that allow them to be represented as cardinal (Ross, 2006b; Sepielli, 2010) or else aggregated by a distinct aggregation rule (MacAskill, 2016; Tarsney, 2019).⁶

But this dramatically understates the challenge of structural diversity. To see the scope of the challenge, it is worth briefly considering some of the potentially aggregation-relevant ways in which the structure of normative theories might differ. For instance...

- The cardinal structure of the real number line can be extended in various ways to accommodate the possibility of infinite or infinitesimal values. A

⁶The terms “ordinal” and “cardinal” are useful but imprecise. I use them as follows: A theory has *simple ordinal* structure if it *totally preorders* options in terms of relative choiceworthiness. A theory has *ordinal* structure more generally if it *preorders* options in terms of relative choiceworthiness *or some other binary normative relation*. A theory has *simple cardinal* structure if it assigns each option a degree of choiceworthiness from a one-dimensional preordered vector space of choiceworthiness values isomorphic to the real numbers. A *ratio-scale* theory is a simple cardinal theory whose choiceworthiness assignment is unique up to positive linear transformation. An *interval-scale* theory is a simple cardinal theory whose choiceworthiness assignment is unique only up to positive affine transformation. A *cardinal* theory more generally is any theory that assigns to options degrees of choiceworthiness *or some other normative property* that are arranged in any preordered vector space (which generalize the structure of the real numbers).

theory that recognizes the possibility of infinite value but does not distinguish degrees of infinite value has a value scale with the structure of the *extended* real number line (the real numbers together with the special elements $+\infty, -\infty$). Alternatively, if we want to distinguish degrees of infinite value, or allow for infinitesimal value, we might entertain a theory whose choiceworthiness scale has the structure of the hyperreal or surreal numbers.⁷

- A theory’s choiceworthiness scale might have more than one dimension. For instance, it might represent degrees of choiceworthiness by two-dimensional vectors, lexicographically ordered (i.e., $(x_1, y_1) \geq (x_2, y_2) \Leftrightarrow (x_1 > x_2 \vee (x_1 = x_2 \wedge (y_1 \geq y_2)))$).⁸ Such a two-dimensional structure could also be used to represent incomparable values, with the overall ordering being the intersection quasi-ordering of the total orderings given by the two dimensions (i.e., $(x_1, y_1) \geq (x_2, y_2) \Leftrightarrow (x_1 \geq x_2 \wedge y_1 \geq y_2)$). Either of these structures could be extended to arbitrarily many dimensions. And the various dimensions of a multidimensional choiceworthiness scale could themselves display various structures, e.g. ordinal, interval, ratio, extended reals, hyperreals, etc.
- At the other end of the scale of complexity, a theory could simply have what we might call *binary structure*: classifying options as permissible or impermissible, but making no normative distinctions among permissible options or among impermissible options. This could be the structure, for instance, of an extreme libertarian view according to which morality consists entirely of negative side constraints, which it is never permissible to violate, and such that all options

⁷Chen and Rubio (forthcoming) propose a normative theory with a surreal-valued utility function. Bostrom (2011) describes, without endorsing, a hyperreal-valued axiology, meant to enable utilitarian comparison of worlds with infinite populations. Several other approaches to infinite axiology (e.g. Vallentyne and Kagan (1997), Arntzenius (2014)) extend the utilitarian value scale beyond the real numbers in other, less canonical ways.

⁸This sort of theory is discussed by MacAskill (2014, pp. 47–50).

that do not violate side constraint are equally choiceworthy.

- Finally, philosophical expositions of “commonsense” morality should plausibly exhibit an *internal* diversity of structure, sometimes making precise interval- or ratio-scale comparisons (e.g., when our options involve providing a fixed material good to different numbers of people); sometimes making such comparisons only imprecisely (e.g., when comparing aesthetic values, or allocating health resources in a diverse population), perhaps with features like “parity” (Chang, 2002); and sometimes invoking constraint-like considerations that, even if non-absolute, cannot be treated as mere multiples of ordinary consequentialist considerations (e.g., the prohibition on punishing the innocent). It seems unlikely that commonsense morality, or any modest refinement thereof, can be represented by anything so straightforward as a total preorder or a real-valued choiceworthiness function on the set of all possible options.

Is there any order underlying this chaos of possible structures? Maybe not, but the following strikes me as a plausible hypothesis: What is essential to any normative theory is *binary* structure—more specifically, the identification of some set of options in each choice situation as permissible or “eligible for choice.”

Universal Binary Structure (UBS) The essential task of normative theories is to classify options in particular choice situations as permissible or impermissible. So a normative theory can have *any* structure that is capable of inducing such a binary classification—perhaps with the added constraint that at least one option in each choice situation is always permissible.

I already implicitly introduced this hypothesis in the last section, by allowing normative theories to recognize an arbitrary set of normative properties and relations, but requiring that every theory include a *choice function* mapping each choice situation to a choice set of permissible options.

UBS is motivated by the idea that the essential job of any normative theory is to guide agents' choices, i.e., to tell us what to do.⁹ If this is right, then nothing can count as a normative theory that does not make the binary distinction between permissible and impermissible options.¹⁰ But the diversity of possible structures explored above is explained by the many possible extensions of binary structure: any structure that can *induce* a binary classification of options (or, in the language of the last section, any set of normative properties and relations on which a binary classification of options can *supervene*) is a possible structure of a normative theory. And this means, more or less, any structure whatsoever.

The most obvious alternative to UBS is:

Universal Ordinal Structure (UOS) The essential task of normative theories is to *compare and rank* options in particular choice situations. So a normative theory can have any structure that is capable of inducing an ordinal ranking—that is, a preordering—of options in terms of choiceworthiness.

It's not entirely obvious that UBS and UOS are genuinely rival hypotheses. On the one hand, any theory with ordinal structure can induce a binary classification in which just the maximally choiceworthy options are permissible (at least for finite option sets). And on the other hand, any theory with binary structure can induce at least a degenerate, choice-set-dependent ordinal ranking (where, in each choice

⁹Of course, particular normative theories are not fully action-guiding in that they don't tell an agent what to do in light of her normative uncertainties. But the assessments given by normative theories play an essential *role* in action guidance, both as inputs to a final, fully action-guiding assessment, and insofar as an agent who is *certain* of a particular normative theory must be able to derive action guidance from its assessment alone.

¹⁰This claim is most plausible if “permissible” and “impermissible” are understood very thinly: e.g., an option *O* is permissible in situation *S* according to theory *T* iff it is possible for an agent who believes *T* with probability 1 to choose *O* in situation *S*, without thereby exhibiting akrasia or any other failure of practical rationality. Even normative theories that eschew traditional deontic notions plausibly define choice sets of options that are permissible in this thin sense.

situation, the permissible options are all equally choiceworthy, the impermissible options are all equally choiceworthy, and the permissible options are more choiceworthy than the impermissible options). But UOS is more demanding than UBS if it is read as requiring that normative theories uniquely define a choice-set-independent ranking of all possible options, or a choice-set-dependent ranking that can involve arbitrarily many ranks.

Adjudicating between UBS and (various precisifications of) UOS would require a lengthy discussion which I will not attempt here. As we will see, which hypothesis we accept has some bearing on questions we will discuss later—in particular, it tells us something important about what structural depletion views (and multi-stage aggregation views involving structural depletion) are committed to. So it is useful to have the two hypotheses on the table. But which hypothesis we accept will not make an essential difference to the arguments that follow.

Most importantly, whichever hypothesis we accept, we reach the conclusion that normative theories can display an enormous range of structures. If UBS is true, then they can display any extension of binary structure, which means, in effect, almost any structure imaginable. If UOS is true, the range of possible structures might be somewhat more constrained, but still goes far beyond simple ordinal, interval, and ratio scales (to include, e.g., the sort of multidimensional and infinitary structures described above). A general theory of rational choice under normative uncertainty must be able to accommodate this diversity. The balance of this paper will consider how that might be done.

4 Responses to structural diversity: a survey

In this section, I will briefly survey a long list of possible responses to the problem of structural diversity, before identifying three particularly interesting responses that we will explore in more depth in the following sections. The long list can

be divided into three categories: *Non-aggregation views* are theories of decision-making under normative uncertainty according to which we are never required to aggregate the assessments of multiple normative theories, and therefore don't need to worry about structural diversity as an obstacle to aggregation. *Aggregation without structural diversity* tries to legislate the problem away by allowing aggregation only within structurally homogeneous sets of theories. Finally, *aggregation with structural diversity* confronts the problem head on, looking for methods of combining rankings of options with arbitrarily diverse structures.

Non-aggregation views

1. First-order normative externalism: According to this view, there is no interesting sense in which what an agent ought to do directly depends on her normative beliefs or uncertainties. Rather, an agent should simply act on the *true* normative theory, even if she has no way of knowing what the true normative theory is and even if her evidence strongly supports some other theory. Versions of first-order externalism have been defended by Weatherson (2014, 2019), Harman (2015), and Hedden (2016). But this view is subject to powerful (and in my view, decisive) objections—see for instance Sepielli (2016, 2018), Podgorski (forthcoming) , and MacAskill and Ord (forthcoming, pp. 3–6).
2. “My Favorite Theory” (MFT): In its simplest form, this view says that an agent should always act on the single normative theory to which she assigns greatest credence. Versions of MFT have been defended by Gracely (1996) and Gustafsson and Torpman (2014). But MFT is also subject to very powerful objections, at least in the forms in which it has so far been articulated—see for instance MacAskill and Ord (forthcoming, pp. 6–9).

Aggregation without structural diversity

3. My Favorite (Structurally Homogeneous) Class of Theories: The most powerful criticism of the version of MFT defended by Gustafsson and Torpman (2014) is that it individuates theories so finely that an agent may distribute her credence over hundreds or thousands of theories, and the single theory in which she has greatest credence may still only command some tiny portion of her credal distribution. Elsewhere, I have suggested that a more plausible version of MFT might instruct an agent to act on the *class* of mutually comparable theories to which she assigns greatest credence Tarsney (2017, pp. 215–7). If shared structure is a prerequisite for intertheoretic comparability, then this view would sidestep the problem of structural diversity: although it requires agents to aggregate the verdicts of multiple theories to decide what to do, it requires such aggregation only within classes of structurally uniform theories. But this view seems more than a little *ad hoc*, and has yet to find any advocates in the literature.

4. Epistemic structural externalism: On this view, there is some privileged structure such that agents are rationally required to have positive credence only in normative theories that possess that structure, or can be represented as possessing that structure. Someone might take this view, for instance, if they find the von Neumann–Morgenstern (VNM) axioms so compelling that they think it is irrational to have any positive credence in normative theories whose rankings of risky options don't satisfy those axioms. Since theories that satisfy the VNM axioms can be represented as maximizing the expectation of a cardinal choiceworthiness function, all theories that satisfy these axioms are, plausibly, similar enough in their structures to enable intertheoretic aggregation. Though to my knowledge this “structural externalist” view has not been explicitly advocated in the normative uncertainty literature, it could be seen as implicit in Ross (2006b, p. 763) or Sepielli (2010, pp. 173–180).

5. Practical structural externalism: On this variant of the structural externalist view, while one is rationally permitted to have positive credence in views that lack the privileged structure, one is rationally required to ignore that credence for purposes of decision-making, and to take as inputs to practical reasoning only one's credence in theories that possess the privileged structure. This view seems less plausible than its epistemic cousin, and to my knowledge has not been defended in the literature.

Finally, there are the views that confront structural diversity head on, devising methods for aggregating theories whose structures differ in aggregation-relevant ways. Since these approaches will be the focus of the coming sections, I introduce them only briefly for now.

Aggregation with structural diversity

6. Structural enrichment: To aggregate structurally diverse theories, we should *add* structure to the more sparsely structured theories, until their structure is relevantly similar to that of more richly structured theories. For instance, we might “cardinalize” ordinal theories by turning their ordinal rankings into Borda scores, allowing them to be aggregated with theories that already possess cardinal structure by an expectational aggregation rule (MacAskill, 2014).
7. Structural depletion: To aggregate structurally diverse theories, we should simply ignore all but some minimal, universal structure that all normative theories have in common (e.g., binary or ordinal structure), and use aggregation rules that apply to theories with only that minimal structure but that may be relatively insensitive to theories with richer structure. For instance, we might ignore all but the binary structure of each theory and aggregate by means of MFO, or ignore all but the ordinal structure of each theory and aggregate by means of some non-scoring voting method (Nissan-Rozen, 2012).

8. Multi-stage aggregation: To aggregate structurally diverse theories, we should first aggregate classes of structurally similar theories by aggregation rules that cover each class, then take the outputs of those aggregations as inputs to further stages of aggregation that combine dissimilar structures by means of enrichment and/or depletion.

This list of options is certainly not exhaustive. Among other things, it leaves out various possible combinations of (1)–(8), like a combination of (6) and (7) that regiments all theories to an intermediate structure by enriching some and depleting others. But it includes what seem like the main contenders and, to my knowledge, all the answers to the problem of structural diversity that have been implicitly or explicitly forwarded in the literature.

In the rest of the paper, I will focus on options (6)–(8), the approaches that take the phenomenon of structural diversity seriously and devise aggregation rules to cope with it. My primary reason for this focus, of course, is that it’s where the interesting problem of structural diversity lies—if we adopt a view that makes the problem disappear by fiat, then there is not much left to say about structural diversity *per se*. But in addition, options (1)–(5) all have significant drawbacks. The debates over (1) and (2) (first-order externalism and MFT) are too large to enter into here, but the objections in the recent literature (cited above) seem quite powerful. (3)–(5) have yet to be considered in any depth, and so are harder to assess. (3) clearly inherits some, though not all, of the drawbacks of (2) (and feels less natural and parsimonious). There may be room for a plausible version of (4) (I am less optimistic about (5)), but it has yet to be developed. And at least *prima facie*, the vigorous and longstanding debate over structure-determining axiom systems like VNM counts against the view that doubting these axioms is proof of irrationality.

So, on the plausible though not incontestable assumption that rational decision-making should be sensitive to an agent’s uncertainties about basic normative questions, including uncertainties that cut across theories with relevantly different struc-

tures, let’s explore the three options that this assumption presents us with: structural enrichment, structural depletion, and multi-stage aggregation.

5 Structural enrichment

Structural enrichment makes theories structurally compatible, for purposes of aggregation, by imputing additional structure to certain theories. This approach seems to be implicit in Lockhart’s “Principle of Equity among Moral Theories,” which compares theories by representing every theory as assigning the same (cardinal) value v^+ to its most preferred option in a given choice situation, and the same (cardinal) value v^- to its least preferred option (Lockhart, 2000, p. 84). Lockhart, however, may simply have assumed that all the theories being compared had cardinal structure to begin with. The idea of structural enrichment is more explicit in MacAskill (2014). MacAskill recognizes the existence of merely ordinal theories, and suggests that the right way to aggregate them is by a kind of weighted Borda count: Very roughly, we should represent each theory T_i as assigning an integer score to each option O_j equal to the number of options T_i ranks below O_j , minus the number of options it ranks above O_j , and then aggregate over theories by taking a probability-weighted sum of these scores. By extension, MacAskill suggests that the right way to aggregate ordinal and cardinal theories with one another is to represent the ordinal theories by their (cardinal) Borda scores, and then normalize all of the theories at their variance—that is, rescale the choiceworthiness assignments of the cardinal theories, and the Borda scores derived from the ordinal theories, so that the variance (average squared distance from the mean) of each theory’s assignment is equal to the same, standard value (e.g., 1).¹¹ We can then aggregate the cardinal and ordi-

¹¹For MacAskill’s defense of the Borda rule, see MacAskill (2016). For his defense of variance normalization, see MacAskill (2014, Ch. 3). It should be noted that MacAskill (2016) focuses on cases where an agent has credence only in ordinal theories, so it is only MacAskill (2014) that defends the Borda method as a form of structural enrichment and a response to the problem of

nal theories together by simply maximizing expected “choiceworthiness”—or, more strictly, the expectation of a variable than in some cases (with respect to the cardinal theories) represents choiceworthiness itself, and in other cases (with respect to the ordinal theories) represents Borda scores derived from a choiceworthiness ranking.

There are at least three reasons to be skeptical of this approach. First, any form of structural enrichment seems to involve an element of arbitrariness. For instance, there are infinitely many ways to “cardinalize” the assessments of an ordinal theory. And if the theory is genuinely ordinal, then, tautologically, nothing in the content of the theory itself tells us which cardinalization to choose—they are all equally compatible with the information the theory does provide, namely its ordinal ranking.¹² More generally, beyond the special case of ordinal and cardinal theories, *any* form of “enrichment” involves adding information to a theory that isn’t part of the theory itself—i.e., making things up. Pulling information out of the aether in this way seems by its nature an arbitrary activity, and not something that rationality can require us to do in any particular way.

That said, the charge of arbitrariness is hard to make stick. All that is needed to escape it is some appropriate justification for doing things one way rather than any other. And purported justifications come cheap—it is possible to produce *some* argument in favor of just about any norm. So whether a norm can be convicted of arbitrariness just depends on whether its purported justifications count as appropriate, admissible, or sufficient. And whether a particular justification meets this threshold is always up for philosophical debate. So I don’t want to rest too much argumentative weight on the charge of arbitrariness.

The version of structural enrichment that seems to have the best claim to being

structural diversity.

¹²Formally, say that a function $f: \mathbf{O} \mapsto \mathbb{R}$ represents a preorder \succsim on \mathbf{O} just in case $\forall(O_i, O_j \in \mathbf{O})((O_i \succsim O_j \Rightarrow f(O_i) \geq f(O_j)) \wedge (O_i \succ O_j \Rightarrow f(O_i) > f(O_j)))$. For any preorder \succsim , there are infinitely (indeed, uncountably) many functions that represent it in this sense, and \succsim itself gives us no guidance as to which function we should prefer.

uniquely justified—and hence the best shot at beating the charge of arbitrariness—is MacAskill’s method of Borda scoring plus variance normalization. Borda scores seem like a particularly natural method of cardinalizing ordinal theories, since they treat each “step” in a theory’s ranking equivalently (equalizing the distance between adjacent ranks). More formally, converting ordinal ranks into cardinal values and maximizing the expectation of those values amounts to aggregation by means of a *scoring rule*, and along with its intuitive naturalness, Borda is the only scoring rule that satisfies certain desirable formal criteria—e.g., never selecting a Condorcet loser, never ranking a Condorcet winner last, never ranking a Condorcet loser above a Condorcet winner, and satisfying a certain attractive weakening of Arrow’s Independence of Irrelevant Alternatives.¹³

Variance normalization likewise has some uniquely attractive properties as a normalization method: It gives each theory a “say” in decision-making proportionate to its probability, in the sense that it equalizes the distance of each theory’s vector of choiceworthiness values from that of the uniform theory (the theory that treats all options as equally choiceworthy). And it equalizes the expected choiceworthiness, from each theory’s perspective, of having its assessments included in the aggregation process with a given unit of probability weight.¹⁴

But if the variance-normalized Borda rule *is* the most natural version of structural enrichment (uniquely overcoming the charge of arbitrariness), then this is bad

¹³For a summary of these properties, see MacAskill (2016) and citations therein (especially Saari (1990)). The Borda rule could also be defended via the axiomatization in Young (1974).

Another key step in MacAskill’s defense of the Borda rule is that only scoring rules can satisfy an appealing property he calls Updating Consistency (MacAskill, 2014, p. 75; MacAskill, 2016, p. 990). But as we will see in §7 below, once we consider the full context of decision-making under normative uncertainty—including interactions between empirical and normative uncertainties and intertheoretic comparability between cardinal theories—MacAskill’s view as a whole turns out to violate Updating Consistency.

¹⁴For precise statements of these criteria, see MacAskill (2014, pp. 110–116).

news for structural enrichment, since MacAskill’s method faces serious difficulties. These difficulties constitute a second objection to structural enrichment, since they problematize the most appealing version of the view. I have expressed my reservations about the Borda rule elsewhere (Tarsney, 2019), and will not belabor them here, but in short: (1) The Borda rule is extremely sensitive to the availability of irrelevant alternatives, in ways that seem plausible only if we interpret Borda scores as an imperfect proxy for hidden cardinal information, and not in the context of genuinely ordinal theories. (2) MacAskill’s method requires that we introduce a measure on \mathbf{O} , but it looks extremely difficult to define such a measure in a principled and general way without yielding wildly implausible practical conclusions.

Variance normalization has yet to be closely studied, but it also seems to have some significant drawbacks. First, a theory’s variance (either the variance of a cardinal theory’s choiceworthiness function, or the variance of the Borda scores representing an ordinal theory’s choiceworthiness ranking) may not be well-defined, even after the imposition of a measure, and there is no obvious way of generalizing variance normalization to theories that don’t have a defined variance. Second, variance normalization introduces another source of extreme sensitivity to irrelevant alternatives. For instance, suppose you are faced with a million options, $O_1, O_2, \dots, O_{1,000,000}$, and divide your credence between ten merely-ordinal normative theories, T_1, T_2, \dots, T_{10} . All ten theories agree that O_{1-100} are the only real candidates—all theories regard these hundred options as strictly better than all the rest. But while theories T_{1-9} give finely discriminating rankings even of the obviously-bad options—not treating any pair of options $O_{101-1,000,000}$ as exactly equal in choiceworthiness— T_{10} is simply indifferent among all the bad options $O_{101-1,000,000}$. Under variance normalization, this indifference among bad options gives T_{10} near-dictatorial power determine which of the more reasonable options you will choose: Even if T_{10} commands a mere one percent of your credence, the fact that T_{10} is indifferent among $O_{101-1,000,000}$ while T_{1-9} are not *guarantees* that the variance-normalized Borda rule will select T_{10} ’s

most preferred option.¹⁵

There is an important complication here: the distinction between “broad” and “narrow” metanormative theories (MacAskill, 2014, pp. 123–5). A *broad* theory takes as input each normative theory’s assessment of the set of all possible practical options, \mathbf{O} , and gives as its output a single metanormative assessment of the options in \mathbf{O} that applies to all choice situations. A *narrow* theory takes as input each normative theory’s assessment of the options *in a particular choice situation*, and gives as its output an option-set-relative metanormative assessment that applies only to that choice situation. Though many of the same formal worries apply to both broad and narrow versions of the variance-normalized Borda rule, the force of those worries varies considerably depending on which version we consider. For instance, sensitivity to irrelevant alternatives (as in the million-option case above) is a powerful objection to the narrow view, but arguably unproblematic for the broad view. On the other hand, the possibility of undefined variance seems remote on the narrow view, but pressing perhaps to the point of inevitability on the broad view.

There is much more to be said about MacAskill’s approach and other possible versions of structural enrichment. But there is at least some reason to suspect that structural enrichment is leading us down a dead end: The most plausible version of the approach—and the version that has the best claim to rebutting the charge

¹⁵A similar problem arises with cardinal theories: Suppose you are faced with ten options, O_{1-10} , and divide your credence between ten normative theories, T_{1-10} . O_{1-9} are all reasonable options, while O_{10} is clearly terrible—every theory considers it the least choiceworthy option. T_{1-5} regard O_{10} as merely very bad—its shortfall relative to the best option is ~ 10 times greater than that of the second worst option. But T_{6-10} regard O_{10} as *abominable*—its shortfall relative to the best option is $\sim 10^{10}$ times greater than than of the second worst option. Variance normalization implies, in this case, that T_{6-10} should get almost no say in determining which of your reasonable options, O_{1-9} , you choose, simply because of this disagreement about exactly how bad the obviously bad option O_{10} is. As in the ordinal case, variance normalization seems to make our choices implausibly sensitive to clearly irrelevant considerations.

of arbitrariness—faces a thicket of technical challenges, and it is at best unclear if they can all be overcome.

A third and final objection to structural enrichment is that it’s hard to generalize. As I emphasized in §3, the two cases that have received the most attention to date—simple cardinal and simple ordinal theories—are just two particularly salient members of a much larger set of possible structures, and a solution to the problem of structural diversity must have something to say not just about these two structures, but about structural diversity in general. But it’s hard to imagine any general principle of structural enrichment (e.g., a generalization of Borda counting/variance normalization) that would tell us what to do with multidimensional theories, various infinite value structures, theories involving intransitivities or option-set dependence, etc. Really following through on the structural enrichment approach may require a hodgepodge of *ad hoc* devices to handle all the aggregation-relevant ways in which one theory’s structure can outstrip another’s.

These objections are far from conclusive—it’s possible that the variance-normalized Borda rule, or some other form of structural enrichment, can be developed in a way that mitigates the worries I’ve described. And as we will see, the alternatives to structural enrichment have problems of their own. But these worries reinforce a more basic suspicion that structural enrichment simply does not take theories at face value. In imposing alien structure on theories, and inventing information where none exists to populate that structure, it seems like an *ad hoc* solution to the problem of structural diversity that does not take seriously agents’ credence in theories with genuinely sparse structure. So it is worth considering alternative approaches.

6 Structural depletion

The second approach, *structural depletion*, takes the opposite tack. Rather than “leveling up” sparsely structured theories to make them compatible with more richly

structured theories, it instead “levels down” the more enriched theories. More precisely, structural depletion does not edit the structure or content of theories in order to make a preferred aggregation rule applicable, but rather adopts an aggregation rule that applies to every theory from the outset, because it is insensitive to all but the minimal, universal structure that all theories have in common.

What aggregation rules are available to the proponent of structural depletion depends on what structure we take to be universal. If UBS is correct (i.e., if the only universal structural feature of normative theories is the binary classification of options as permissible or impermissible), then structural depletion demands an aggregation rule that takes only binary information as its input. This limits the proponent of structural diversity, more or less, to MFO: An agent should aggregate the assessments of her normative theories simply by summing up the probabilities that each option is permissible, and choosing an option for which this sum is maximal.¹⁶ If UOS is correct (i.e., if the minimal structure all normative theories have in common is ordinal rather than binary), then more aggregation rules are available—in particular, various non-scoring (i.e., non-cardinalizing) voting rules borrowed from social choice theory (Nissan-Rozen, 2012; Tarsney, 2019).¹⁷

¹⁶Other options open to the proponent of structural depletion, given UBS, include (i) a “threshold” view, on which an agent may choose any option that has a probability greater than t of being permissible (or, if no option exceeds threshold t , then an option with maximal probability) or (ii) a “satisficing” view, on which an agent may choose any option whose probability of being permissible is close enough to maximal. These views might be motivated by the intuitive over-demandingness of MFO in cases where several options carry only a very small risk of being impermissible.

¹⁷MFO can itself be regarded as a probability-weighted voting rule, somewhere in between simple plurality and approval voting: Each theory gives a yes/no vote on each option, and the option with the greatest weighted sum of “yes” votes wins. Each theory *can in principle* vote “yes” on any number of options (as with approval voting), but only *does* vote yes for options it regards as permissible. Insofar as most theories will regard only maximally choiceworthy options as permissible, and will judge in most choice situations that only one option is maximally choiceworthy (both non-obvious assumptions), MFO might in practice be more like simple plurality voting (where

As with structural enrichment, I will raise three interrelated worries about structural depletion—worries that apply to both binary and ordinal versions of the approach.

First, structural depletion throws away what seems like eminently decision-relevant information—in particular, it mishandles cases where cardinal intertheoretic choiceworthiness comparisons are possible. Suppose an agent, Alice, is faced with two options, O_1 and O_2 . She is almost sure that some version of maximizing consequentialism is true. Conditional on that assumption, she is sure that hedonic experience has moral value, but unsure whether there are also non-derivative *aesthetic* values. She also has non-zero credence in a non-consequentialist theory that provides merely ordinal information. She thus divides her credence between three normative theories: a hedonistic and a pluralistic version of consequentialism, and a version of non-consequentialism.

As I have argued elsewhere (Tarsney, 2018a), it is natural in a case like this to think that, insofar as the two consequentialist theories are in complete agreement about hedonic value and only disagree about the existence of aesthetic value, their cardinal choiceworthiness scales should be normalized at the value of a hedon—i.e., by treating the theories as assigning the same value to a given unit of hedonic experience.¹⁸ And if we spell out the content of the theories in a sufficiently fine-grained way, we can make a compelling theoretical case for this normalization as well (see in particular Appendix B of Tarsney (2017)).

Let's suppose that, given this normalization, the judgments of the three theories can be represented as in Table 1. The notable feature of the case is that, while the two consequentialist theories are equally probable and disagree about which option is more choiceworthy, the pluralistic theory T_2 regards the choice as much higher—

each voter is allowed only a single “yes” vote).

¹⁸Similar claims about this sort of case have been made by Ross (2006b, pp. 764–5) and MacAskill (2014, p. 134).

Theory	Credence	Assessment
T_1 (hedonism)	.49	$CW(O_1) = 0, CW(O_2) = 1$
T_2 (pluralism)	.49	$CW(O_2) = 0, CW(O_1) = 100$
T_3 (non-consequentialism)	.02	$O_2 \succ O_1$

Table 1: Alice’s choice, before depletion

Theory	Credence	Assessment (ordinal/binary)
T_1 (hedonism)	.49	$O_2 \succ O_1 / f_1(\{O_1, O_2\}) = \{O_2\}$
T_2 (pluralism)	.49	$O_1 \succ O_2 / f_2(\{O_1, O_2\}) = \{O_1\}$
T_3 (non-consequentialism)	.02	$O_2 \succ O_1 / f_1(\{O_1, O_2\}) = \{O_2\}$

Table 2: Alice’s choice, after depletion

stakes than the hedonistic theory O_1 . Given this information, it seems clear that Alice should choose O_1 : She is nearly certain that some consequentialist normative theory is correct, she is able to make intertheoretic comparisons between the consequentialist theories in which she has positive credence, and given those comparisons, O_1 has vastly greater expected choiceworthiness than O_2 (conditional on some version of consequentialism being true).

But any version of structural depletion—whether binary or ordinal—is committed to choosing O_2 . If we strip away everything but the binary/ordinal information, we are left with the information given in Table 2. Here we have nothing left to go on besides the facts that (i) O_2 is more likely than O_1 to be permissible, (ii) O_2 is more likely than O_1 to be maximally choiceworthy, and (iii) it is more likely that $O_2 \succ O_1$ than that $O_1 \succ O_2$. All of these facts, of course, favor O_2 over O_1 . By ignoring cardinal information for the sake of simplifying the aggregation process, therefore, we are forced to ignore an extremely compelling reason to choose O_1 .

Second, because it ignores cardinal information, structural depletion is vulnerable to diachronic inefficiencies. Consider, for instance, the following sequences of choice situations. Each situation involves two options, assessed by three equiprobable theories T_{1-3} , each with simple cardinal structure. (For simplicity, let $T_n(O_i)$

designate the cardinal value that T_n assigns to O_i .)

- $S_1 = \{O_1, O_2\}$: $T_1/T_2(O_1) = 1$, $T_1/T_2(O_2) = 0$, $T_3(O_1) = 0$, $T_3(O_2) = 10$
- $S_2 = \{O_3, O_4\}$: $T_2/T_3(O_3) = 1$, $T_2/T_3(O_4) = 0$, $T_1(O_3) = 0$, $T_1(O_4) = 10$
- $S_3 = \{O_5, O_6\}$: $T_3/T_1(O_5) = 1$, $T_3/T_1(O_6) = 0$, $T_2(O_5) = 0$, $T_2(O_6) = 10$

Any version of structural depletion seems committed to choosing options O_1 in S_1 , O_3 in S_2 , and O_5 in S_3 . O_1 , for instance, is the option most likely to be permissible in S_1 , most likely to be maximally choiceworthy, and more likely to be more choiceworthy than O_2 than to be less choiceworthy, so it defeats O_2 according to any ordinal social choice method that satisfies minimal neutrality criteria. And the case for O_3 and O_5 is exactly the same. This sequence of choices, however, yields total payoffs of 2/2/2 according to theories $T_1/T_2/T_3$, respectively, whereas choosing options O_2 , O_4 , and O_6 would have yielded payoffs of 10/10/10 according to $T_1/T_2/T_3$, respectively. That is, structural depletion commits the agent to a sequence of choices that is worse than an available alternative sequence, according to *every* theory in which she has positive credence.¹⁹

Two points are worth noting: First, this objection does not depend on any assumption of intertheoretic comparability. We can rescale the choiceworthiness assignments of each theory however we want and the point remains, that choosing $O_1/O_3/O_5$ is certainly worse than choosing $O_2/O_4/O_6$.²⁰ Second, this is a “forced” inefficiency: Structural depletion does not merely *permit* but *requires* the agent to choose the certainly-suboptimal sequence of options. This makes for a much more compelling objection, *inter alia* because there is no hope of avoiding the problem by

¹⁹A related diachronic inefficiency objection to MFO is described by Gustafsson and Torpman (2014, pp. 165–6).

²⁰The argument *does* depend on the stipulation that T_{1-3} all regard the choiceworthiness of a sequence of options as the sum of the choiceworthiness of the individual options making up the sequence. Theories need not evaluate sequences of options in this way, but it’s clearly possible.

supplementing a structural depletion approach with principles of diachronic rationality that let us avoid sure losses—any such principle would contradict, not extend, a metanormative view based on structural depletion.²¹

Both structural enrichment and (as we will see) multi-stage aggregation approaches can avoid this sort of inefficiency, by taking account of the cardinal information provided by T_{1-3} . This does not necessarily mean that these approaches will recommend the sequence $O_2/O_4/O_6$. Rather, their recommendations will depend on how we normalize T_{1-3} . If, for instance, the correct normalization preserved the prenormalized choiceworthiness values for T_{1-2} given above, but multiplied T_3 's values by 100, then an expectational aggregation rule would recommend the sequence $O_2/O_3/O_5$. But however we normalize the theories, we will not end up choosing the inefficient sequence $O_1/O_3/O_5$.

Third and finally, structural depletion seems committed to drawing a fundamental line between normative and empirical uncertainties. A thoroughgoing structural depletion view of decision-making under *uncertainty* in general might claim, for instance, that one should always choose an option with a maximal probability of being maximally choiceworthy *in a fully fact-relative sense*, i.e., in a sense that takes no account of one's normative *or* empirical beliefs. But this decision theory would imply, for instance, that if I suspect someone has poisoned my coffee, I should go

²¹The only version of structural depletion that might be able to avoid this objection, as far as I can see, is an approach that (i) takes ordinal rather than binary information as input and (ii) adopts the “broad” approach of imposing a single metanormative ranking on the set of all possible options and instructing agents to choose, in a given choice situation, an option that is maximal among their available options in that ranking. On this approach, how any given social choice method would rank, for instance, options O_1 and O_2 would depend on assumptions about the content of the set of all possible options (and, since this set is presumably infinite, would depend on a choice of measure on the set). This makes it very hard to figure out what such a view would recommend in any given choice situation, including S_{1-3} above—which is, of course, a significant objection in its own right to such a view. In any event, no one in the literature has tried to develop a view of this sort, so I set it aside as a possibility for future research.

ahead and drink so long as the probability that my cup contains deadly poison is a mere .49, rather than a prohibitive .51. No one, to my knowledge, is prepared to defend such a view. So the proponent of structural diversity must instead advocate a view on which we account for cardinal information when dealing with empirical uncertainty (e.g., to compute the expected choiceworthiness of an option according to a given normative theory), but then discard that information as soon as we are dealing with normative uncertainty—even when our normative uncertainties concern theories that provide us with usable cardinal information, as in the case of Alice. While there is lively debate over just how continuous normative uncertainty is with empirical uncertainty (see for instance Weatherson (2014, 2019), Tarsney (2017, pp. 57–66)), treating the two kinds of uncertainty differently when we are not forced to is at least *prima facie* inelegant and undermotivated.²²

²²Most of the arguments that have been offered in the literature for making a fundamental distinction between normative and empirical uncertainty are meant to support the conclusion that an agent’s normative beliefs are simply irrelevant to what she subjectively ought to do (e.g. Weatherson (2014, 2019), Harman (2015), Hedden (2016)). These arguments offer scant support for the view that agents ought to account for their normative uncertainties, but in a more structurally impoverished way than they account for their empirical uncertainties.

One possible basis for the kind of distinction drawn by structural depletion views is the claim that intertheoretic comparisons are generally impossible (Gracely, 1996; Gustafsson and Torpman, 2014; Nissan-Rozen, 2015; Hedden, 2016). But even if true, this claim does not force us to accept structural depletion and ignore cardinal information under normative uncertainty. We could instead, for instance, adopt a statistical normalization method (like variance normalization) or a bargaining-theoretic approach to normative uncertainty (Greaves and Cotton-Barratt, unpublished), both of which preserve sensitivity to cardinal information without presupposing intertheoretic comparability.

7 Multi-stage aggregation

This leaves us with the approach I have called *multi-stage aggregation*. On this approach, we first aggregate classes of structurally similar theories by aggregation rules appropriate to each class, then take these aggregations as inputs to further stages of aggregation that combine dissimilar structures by means of enrichment and/or depletion.²³

It's useful to consider an example. Suppose that Betty divides her credence between four theories: the two consequentialist theories described in the last section (hedonistic utilitarianism and a pluralistic theory that recognizes both hedonic and aesthetic value) and two ordinally structured deontological theories that disagree about whether one ought to kill innocent threats in other-defense, but are otherwise in complete agreement. Betty must decide whether to kill Carl: Carl was innocently driving his truck down a city street when his brakes failed, having been sabotaged by some nefarious malefactor. His truck is now careening unstoppably toward five innocent people trapped in a narrow alleyway. The only way to save the five is for Betty, perched on a nearby rooftop, to destroy Carl's truck with a bazooka. In the truck with Carl, however, are a dozen priceless works of art and the only score of a newly rediscovered Beethoven symphony. If Betty opens fire, both Carl and his truckload of aesthetic goods will be destroyed, but five innocents lives will be saved.

The death of an innocent person, let's assume, amounts to a loss of 20 hedons, while the loss of the artistic contents of Carl's truck would amount to a loss of 10 hedons.²⁴ However, their destruction would also amount to a loss of 200 aesthetons (a unit of aesthetic value that, on Betty's pluralistic theory, has the same value

²³This approach was first described—but, as we will see, not endorsed—by MacAskill (2014, p. 118), who calls it a “multi-step procedure.”

²⁴Suppose that, although people will derive enjoyment from these works, that enjoyment is highly substitutable, so that if Carl's truck is destroyed, the would-be appreciators of the artworks it contains would be able to derive nearly as much pleasure from other works.

Theory	Credence	O_1 (kill)	O_2 (don't kill)
T_1 (hedonism)	.3	-30	-100
T_2 (pluralism)	.3	-230	-100
T_3 (Kill IT's.)	.25	permissible	impermissible
T_4 (Don't kill IT's.)	.15	impermissible	permissible

Table 3: Betty's choice, Stage 1

as a hedon). So while Betty's hedonistic theory (T_1) supports killing Carl to save the five, her pluralistic theory (T_2) opposes it. Of Betty's deontological theories (T_{3-4}), meanwhile, one claims that killing Carl is permissible and doing nothing is impermissible, while the other claims that doing nothing is permissible and killing Carl is impermissible. Table 3 represents each theory's assessment of the situation.

Multi-stage aggregation could in principle proceed in either of two directions: (1) aggregating the most richly structured theories first, and depleting the result to facilitate further stages of aggregation, or (2) aggregating the most structurally impoverished theories first, and enriching the result to facilitate further stages of aggregation. I will focus on (1), which as far as I can see is the more promising approach. On this approach, we first aggregate T_1 and T_2 . These theories have the same simple cardinal structure and (let's assume) the aggregation rule appropriate to them is MEC. The expected choiceworthiness of O_1 , relative to the set of theories $\{T_1, T_2\}$, is $(.3 \times -30) + (.3 \times -230) = -78$, and the expected choiceworthiness of O_2 is $(.3 \times -100) + (.3 \times -100) = -60$.²⁵ So, as far as Betty's consequentialist theories are concerned, she should choose O_2 (i.e., not kill Carl).

But how do we aggregate these consequentialist theories with the deontological theories T_3 and T_4 ? What multi-stage aggregation commits us to is aggregating T_1 and T_2 with these less-structured theories *as a pre-aggregated unit*, rather than

²⁵It will make no difference to our final conclusions whether we use the agent's unconditional credences or her credences conditional on the class of theories being aggregated (that is, conditional on $T_1 \vee T_2$). For simplicity, I use the unconditional credences.

Theory	Credence	O_1 (kill)	O_2 (don't kill)
C_1 (consequentialism)	.6	impermissible	permissible
T_3 (Kill IT's.)	.25	permissible	impermissible
T_4 (Don't kill IT's.)	.15	impermissible	permissible

Table 4: Betty's choice, Stage 2

individually. For simplicity, let's assume that T_3 and T_4 have merely binary structure. In this case, we should take the pre-aggregated class of theories $C_1 = \{T_1, T_2\}$, along with the theories T_3 and T_4 , as inputs to a binary aggregation rule. Table 4 represents this second stage of this aggregation procedure. If, as I have claimed, the appropriate aggregation rule for binary-structured theories is MFO, then since Betty's aggregate credence in views according to which O_1 is permissible is .25, while her aggregate credence in views according to which O_2 is permissible is .75, we conclude that she should not kill Carl.

Note that the two-stage procedure we have just described is not equivalent to MFO *simpliciter*: According to MFO, Betty should choose an option such that the sum of the probabilities of *individual normative theories* that assess it as permissible is maximal. In Betty's case, this means choosing O_1 : The theories that prescribe O_1 , viz. T_1 and T_3 , command a combined credence of .55 while the theories that prescribe O_2 , viz. T_2 and T_4 , command a combined credence of only .45. Multi-stage aggregation instead selects O_2 because it is sensitive to the cardinal information provided by T_1 and T_2 : Although T_1 (which prescribes O_1) and T_2 (which prescribes O_2) are equally probable, the stakes are higher according to T_2 , and hence the pair of theories on balance support O_2 over O_1 . This information is ignored by single-stage MFO, but is reflected in the output assessments of a multi-stage procedure.

The *prima facie* appeal of multi-stage aggregation is that it's responsive to the information provided by more richly structured (e.g., cardinal) theories, without artificially imposing that structure on less-structured (e.g., ordinal) theories. In Betty's case, this means that the cardinal values assigned by T_1 and T_2 can make

a difference to our practical conclusions, without the need to implicitly or explicitly cardinalize the non-cardinal theories T_3 and T_4 . In other words, multi-stage aggregation avoids the first two objections to structural depletion—it accounts for cardinal information when it’s available, and thereby can avoid the sort of diachronic inefficiency described in the last section—while also avoiding the first objection to structural enrichment, by not arbitrarily inventing information to fill out structurally impoverished theories.²⁶

But like structural enrichment and depletion, multi-stage aggregation is open to serious objections. I’ll first consider a pair of closely related objections, and sketch a version of multi-stage aggregation intended to mitigate them. I will then introduce a third objection, which requires more extended discussion.

The first two objections are as follows: First, multi-stage aggregation introduces an unacceptable element of arbitrariness, in that its outputs depend on an “order of aggregation.” More precisely, multi-stage aggregation requires that we organize normative theories into something like a nested hierarchy, that tells us which theories or sets of theories should be aggregated with each other at each stage of the aggregation process; the output of multi-stage aggregation can depend on *which* nested hierarchy we choose; and—the objection goes—any choice of nested hierarchy is at least partially arbitrary. Second, for this reason and because of the variety of aggregation methods it potentially involves, multi-stage aggregation results in a decision procedure that is implausibly complex—in particular, too complex to represent a basic requirement of rationality. Unless it can be shown that the chaos of multi-stage

²⁶Importantly, structural enrichment and multi-stage approaches avoid the sort of inefficiency described above only if they make intertheoretic comparisons in a way that is consistent across choice situations. To do this, they must either (a) find some non-statistical basis for intertheoretic comparisons (like agreement on a certain category of value or class of choice situations), or (b), if they rely on statistical normalization methods, adopt a scope that is broader than single choice situations (e.g., equalizing the range or variance of each theory’s choiceworthiness function over a broad domain like *all possible options* or *all the options faced by a particular agent in her lifetime*).

aggregation—perhaps indefinitely many stages of aggregation, with potentially very different aggregation rules being applied to different classes of theories at the lower levels—has some underlying, unifying idea that explains and justifies it, we should reject the approach simply on grounds of parsimony.

These are serious objections, at least to certain forms of multi-stage aggregation. In fact, I initially inclined toward a multi-stage view that involved a potentially indefinite number of stages, but changed my mind after thinking through the pitfalls involved in defining such a complex order of aggregation. Fortunately, however, there is a more straightforward version of multi-stage aggregation that at least mitigates worries about arbitrariness and complexity. Here is a sketch of the view.

The Two-Stage View

- The basic motivating idea behind this view is that rational choice should be based on, and responsive to, all and only the decision-relevant information available to the agent. That is, we should not throw away relevant information but, equally, should not invent information where none exists.
- Respecting this imperative requires a multi-stage procedure because, roughly, the correct response to certain information is to apply an aggregation rule that is inapplicable when that information is not available. For instance, some theories assign cardinal and intertheoretically comparable degrees of choice-worthiness to options; the correct way of accounting for this information is to take a probability-weighted average; but this aggregation rule is inapplicable, e.g., to merely-ordinal theories. Any single aggregation rule, therefore, will either require us to ignore (or at least respond inadequately/incorrectly) to relevant information from more richly structured theories, or else to invent information in order to apply more structurally demanding aggregation rules to less-structured theories.
- The solution to this dilemma is as follows: First, associate each normative

theory with an optimal aggregation rule—a rule that represents the optimal response to the decision-relevant information provided by that theory.²⁷ Second, aggregate sets of theories that share the same optimal aggregation rule, using that rule. Finally, using whatever aggregation rule is optimal for theories with only the minimal structure essential to all normative theories (whether that is binary or ordinal), aggregate the theories for which this is the optimal aggregation rule, together with the results of the first stage of aggregation.

The two-stage view largely assuages worries about arbitrariness and complexity. First, the order of aggregation does not look particularly arbitrary: We aggregate sets of theories that share the same optimal aggregation rule, then aggregate all theories, regardless of their optimal aggregation rule. Second, the resulting decision procedure need not be particularly complex—its complexity will depend on how many different aggregation rules we think are optimal for at least some theories. And whatever complexity results is well motivated by the simple underlying idea that we should respond as best we can to the decision-relevant information provided by the theories in which we have positive credence.²⁸

²⁷What makes an aggregation rule optimal for a given theory is up for debate: It might depend simply on the structure of the theory, with the optimal aggregation rule representing the optimal response to the sort of information provided by theories with that structure. But it might also depend on the decision-theoretic content of the theory itself—e.g., the aggregation rule the theory endorses for responding to empirical uncertainty or, if theories are characterized richly enough to contain such information, the aggregation rule it endorses for dealing with *normative* uncertainty.

²⁸Indeed, the two-stage view is arguably *simpler* than other, putatively single-stage views in the literature. Most extant views, as we will see below, first aggregate over empirical possibilities conditional on each normative theory, before aggregating over normative theories. If the optimal aggregation rule for a normative theory is just the rule that it uses for aggregating over empirical possibilities, then the two-stage view lets us aggregate over both empirical and normative possibilities at the first stage of the two-stage procedure. This amounts to moving from a two-stage procedure that draws a line between normative and empirical uncertainties, to one that draws a line between uncertainties within, and those between, classes of normative theories that share the

But there is another, more pointed objection to the multi-stage approach. As MacAskill has pointed out (MacAskill, 2014, pp. 117–9), any version of multi-stage aggregation seems likely to violate an intuitively compelling axiom he calls Updating Consistency.

Updating Consistency For any choice situation S , theory T_i , and option O , if O is maximally choiceworthy of the options in S according to T_i , and O is maximal in the metanormative assessment of S , then if the agent increases her credence in T_i while keeping the ratios between her credences in all other theories the same, creating a new choice situation S' , O is still maximal in the metanormative assessment of S' .

MacAskill describes a case where any multi-stage aggregation procedure seems committed to violating Updating Consistency, which I have reproduced in Table 5. Suppose, in this case, that the cardinal theories T_2 and T_3 are intertheoretically comparable. If we first aggregate these two theories expectationally, we get a ranking in terms of expected choiceworthiness, $O_3 \succ O_2 \succ O_1$, which covers 5/9 of the agent’s credence. When we then aggregate this ordinal ranking with that given by T_1 , any reasonable method should conclude that O_3 is the preferred option.

But now suppose that the agent reduces her credence in T_2 to zero, and increases her credence in T_1 and T_3 to 4/7 and 3/7, respectively. Then, by the same sort of reasoning, we find that O_1 is the preferred option. So *decreasing* one’s credence in a theory that views O_1 as the most choiceworthy option, while keeping the ratios between one’s credences in all other theories the same, can cause O_1 to rise in the metanormative assessment.

Against this objection to multi-stage aggregation, however, there is a compelling response: A kind of multi-stage aggregation and—more to the point—the violations

same optimal aggregation rule. And it results in a simpler view insofar as it requires fewer separate aggregations—rather than separate aggregations for each normative theory, we aggregate at the first stage over potentially very large classes of normative theories.

Theory	Credence	Assessment
T_1	$4/9 \Rightarrow 4/7$	$O_1 \succ O_2 \succ O_3$
T_2	$2/9 \Rightarrow 0$	$CW(O_1) = 20, CW(O_2) = 10, CW(O_3) = 0$
T_3	$3/9 \Rightarrow 3/7$	$CW(O_1) = 0, CW(O_2) = 10, CW(O_3) = 20$

Table 5: Multi-stage aggregation violates Updating Consistency. (Arrows indicate updates.)

of Updating Consistency that come with it seem nearly unavoidable, once we remember that our overall theory of decision-making under uncertainty has to account for empirical as well as normative uncertainty.

Consider a variant of MacAskill’s case, described in Table 6. In this case the agent has three options, O_{1-3} , positive credence in a simple ordinal theory T_1 and a simple cardinal theory T_2 , and positive credence in two (relevantly distinct) empirical states of the world, S_1 and S_2 . T_1 gives the same assessment of the choice situation regardless of the state of the world, but according to T_2 , the choiceworthiness of O_1 and O_3 depends on the state of the world—specifically, if the world is in state S_1 then O_1 is the best option, but if the world is in state S_2 then O_3 is the best option.

In this situation, any plausible version of structural depletion and MacAskill’s own preferred version of structural enrichment both violate Updating Consistency: Before updating, both views select O_3 , while after the update—in which the agent reduces her credence in a possibility, $T_2 \wedge S_1$, according to which O_1 is maximally choiceworthy—both views select O_1 .²⁹

Consider structural depletion. Assuming we don’t adopt the extreme view that ignores differences in cardinal stakes across empirical as well as normative possibilities, structural depletion tells us to compute T_2 ’s empirical-belief-relative assessment

²⁹Cases like this require probabilistic dependence between normative and empirical beliefs—otherwise the ratio between the agent’s credences in $T_1 \wedge S_1$ and $T_1 \wedge S_2$, which I’ve left unspecified, could not be the same in both choice situations. But this seems entirely possible, albeit unusual. (For defense of the possibility of this sort of probabilistic dependence, see Podgorski (forthcoming).)

Theory	Credence	Assessment
T_1	$4/9 \Rightarrow 4/7$	$O_1 \succ O_2 \succ O_3$
$T_2 \wedge S_1$	$2/9 \Rightarrow 0$	$CW(O_1) = 20, CW(O_2) = 10, CW(O_3) = 0$
$T_2 \wedge S_2$	$3/9 \Rightarrow 3/7$	$CW(O_1) = 0, CW(O_2) = 10, CW(O_3) = 20$

Table 6: Structural depletion and the variance-normalized Borda rule both violate Updating Consistency, in cases involving both empirical and normative uncertainty.

and then aggregate it with T_1 's, using only the binary or ordinal information they supply. T_2 's empirical-belief-relative assessment supplies the ranking $O_3 \succ O_2 \succ O_1$. Before updating, this assessment commands greater total probability weight than T_1 's, making O_3 the preferred option. After the update, it commands less total probability weight, so O_1 is the preferred option.

With respect to the variance-normalized Borda rule, things are roughly the same. Once again, we begin by computing T_2 's empirical-belief-relative assessment. Both before and after updating, these expected values are equally spaced.³⁰ So when we convert T_1 's ranking to Borda scores and variance-normalize the Borda scores derived from T_1 with the empirical-belief-relative expectations of T_2 , we end up simply choosing the option preferred by the more probable normative theory. Before updating, this is O_3 (preferred by T_2); after updating, it is O_1 (preferred by T_1).³¹

Both structural depletion and the best-developed version of structural enrichment, therefore, imply that reducing your credence in one theory, $T_2 \wedge S_1$, while keeping the ratios between your credences in all other theories the same, *improves* the metanormative assessment from the perspective of that theory.

The way to escape these violations of Updating Consistency, for proponents of

³⁰Before updating, the probability-weighted sum of choiceworthiness values given by T_2 is $40/9$, $50/9$, and $60/9$ for $O_1/O_2/O_3$, respectively. After updating, the values are 0 , $30/7$, and $60/7$.

³¹Note that variance normalization alone is enough to generate violations of Updating Consistency in cases like this—for instance, if we replace T_1 with a cardinal theory according to which $CW(O_1) = 2$, $CW(O_2) = 1$, and $CW(O_3) = 0$. So abandoning Borda for another method of enrichment won't solve the problem.

either structural enrichment or structural depletion, is to endorse a genuinely single-stage aggregation rule for both empirical and normative uncertainty. This seems like a clear non-starter for structural depletion. For a proponent of structural enrichment, things are a bit more complicated. A natural move is to simply take a single expectation over normative-plus-empirical possibilities. This has the downside of ignoring the aggregation rules that theories themselves endorse for responding to empirical uncertainty—e.g., ignoring some theories’ non-neutral attitudes toward risk. If, say, a theory tells me that when I’m uncertain about the empirical state of the world, I should maximize a *risk-weighted* expectation of utility, value, or choice-worthiness (Buchak, 2013), using the risk function $r(x) = \sqrt{x}$, then the theory’s empirical-belief-relative assessments will reflect that risk aversion, but a simple expectation over normative-plus-empirical possibilities will not. (For an interesting attempt to mitigate this problem, see Dietrich and Jabarian (unpublished).)

But additionally, this move by itself will not allow us to respect Updating Consistency, unless we are able to make intertheoretic comparisons entirely without recourse to statistical methods like range or variance normalization. The option of using these statistical methods comprehensively to normalize the assessments given by each normative-plus-empirical possibility is, I assume, a non-starter: If, say, we variance-normalize the choiceworthiness values associated with the possibilities (i) “hedonistic utilitarianism is true, and it will rain tomorrow” and (ii) “hedonistic utilitarianism is true, and it won’t rain tomorrow,” we would generally conclude that the normative significance of a hedon according to hedonistic utilitarianism depends on whether or not it will rain tomorrow. So we must instead normalize assessments associated with the same normative theory in a more natural, “content-based” manner (e.g., assuming that hedonistic utilitarianism assigns the same value to a hedon regardless of the state of the world), reserving statistical methods like range or variance normalization for comparisons *between* normative theories.

Similarly, if we think that some normative theories are naturally comparable

Theory	Credence	Assessment
T_1	$3/10 \Rightarrow 0$	$CW(O_1) = 1, CW(O_2) = 0, CW(O_3) = -10^9$
T_2	$4/10 \Rightarrow 4/7$	$CW(O_1) = 1, CW(O_2) = 0, CW(O_3) = 0$
T_3	$3/10 \Rightarrow 3/7$	$CW(O_1) = 0, CW(O_2) = 1, CW(O_3) = 0$

Table 7: Combining statistical and content-based normalization methods results in violations of Updating Consistency.

(e.g., the hedonistic and pluralistic versions of consequentialism introduced in §6) while others are not (e.g., total and average utilitarianism), we will naturally want to normalize sets of comparable theories according to the true comparisons between them (e.g., assuming that hedonism and pluralism agree on the value of a hedon), while reserving statistical methods for normalizing incomparable theories. This means that we normalize *classes* of comparable theories with other, incomparable theories based on their collective variance (or some other statistical feature).³²

But such a mixture of normalization methods itself results in violations of Updating Consistency. Consider the case described in Table 7, involving three simple cardinal theories, where T_1 and T_2 are comparable with each other, but not with T_3 . (We can assume either that T_1 and T_2 are distinct normative theories like hedonism and pluralism, or that they represent conjunctions of the same normative theory with different empirical possibilities.)

Since T_1 enormously inflates the collective variance of $\{T_1, T_2\}$, it leaves these theories with almost no influence over the choice between O_1 and O_2 , allowing T_3 's preference for O_2 to win out. But if the agent reduces her credence in T_1 to zero while increasing her credence in the other two theories proportionately (as shown in the table), the situation is reduced to an aggregation of the symmetric rankings of T_2 and T_3 , and since T_2 is the more probable theory, O_1 is now the preferred option.³³ Notably, this means that MacAskill's view violates Updating

³²This is the approach endorsed by MacAskill (2014, pp. 119–121).

³³I assume here that we compute the collective variance of a set of mutually comparable theories

Consistency even with respect to purely normative uncertainty, before we consider interactions with empirical uncertainty. We get the same result if we use other statistical properties (e.g., range) to normalize $\{T_1, T_2\}$ with T_3 .

The only way I can see for a structural enrichment view to avoid violations of Updating Consistency, therefore, is to avoid any recourse to statistical methods like variance normalization, by claiming that *all* theories are comparable by more natural methods, just as closely related theories like hedonism and pluralism appear to be. But this is an ambitious and *prima facie* very implausible claim. Apart from apparent incomparabilities between simple cardinal theories like total and average utilitarianism, this approach would have to claim that there are natural, non-statistical bases for comparison between, say, the *Borda scores* derived from the rankings of an ordinal theory and the cardinal choiceworthiness values given by a simple cardinal theory. But it's very hard to see how we could make this sort of comparison by anything other than statistical methods. I conclude, therefore, that whatever the intuitive appeal of Updating Consistency, it is unlikely that any plausible metanormative theory will be able to satisfy it.

Thus, the most serious known drawback of multi-stage aggregation seems to be an endemic defect of metanormative theories, which cannot be plausibly avoided by structural enrichment or structural depletion. While I haven't given a comprehensive or even a full-throated defense of the multi-stage approach, therefore, it seems to me that the objections it faces are more tractable than the objections to structural enrichment or structural depletion, and that multi-stage aggregation is therefore the most promising approach to structural diversity.

in a probability-weighted manner, rather than giving equal weight to every possible theory in the set. The latter approach seems infeasible, since there will generally be infinitely many possible theories in any such set, and without probability weights, their collective variance will be undefined.

8 Conclusion

Developing an adequate theory of decision-making under normative uncertainty has proven to be an enormously challenging project. It is natural and appropriate, when we first confront such a challenge, to help ourselves to large simplifying assumptions that make the problem more tractable. Thus, the recent literature on normative uncertainty has mostly assumed either that normative theories all share the same structure, or that they exhibit only a narrow range of structures. But I have tried to show that important considerations come into view when we set aside these simplifying assumptions. The range of possible normative structures is much broader than has generally been acknowledged, and so the problem of structural diversity requires a *general* approach, not one that is merely tailored to the special case of simple ordinal and cardinal theories. Finding such a general approach is not easy—all the obvious candidates have serious drawbacks. And the choice is consequential: The fortunes of particular metanormative views, like My Favorite Option or the Borda rule, may turn on the success or failure of the more general approaches to structural diversity that they exemplify.

I have focused on three approaches to structural diversity—structural enrichment, structural depletion, and multi-stage aggregation—and argued that multi-stage aggregation is the least bad of the three. But even if this is right, it still leaves many questions unanswered. For instance, could a hybrid approach that combines two or more of the views surveyed in §4 improve on the approaches we have considered, capturing the advantages of the views it combines while avoiding their disadvantages? And how complete was the survey in the first place—are there other approaches to structural diversity, qualitatively distinct from those I’ve surveyed and from any hybridization thereof? Finally, to what extent do analogous problems arise in other domains that involve aggregation of assessments with potentially diverse structure (e.g., preference or belief aggregation in populations where some individuals but not others satisfy axioms that allow for unique cardinal rep-

resentations), and what can we learn from these analogies? I leave these questions, among others, for future research.

References

- Arntzenius, F. (2014). Utilitarianism, decision theory and eternity. *Philosophical Perspectives* 28(1), 31–58.
- Bostrom, N. (2011). Infinite ethics. *Analysis and Metaphysics* 10, 9–59.
- Buchak, L. (2013). *Risk and Rationality*. Oxford: Oxford University Press.
- Bykvist, K. (2013). Evaluative Uncertainty, Environmental Ethics, and Consequentialism. In R. I. Hiller, Avram and L. Kahn (Eds.), *Consequentialism and Environmental Ethics*. Routledge.
- Chang, R. (2002). The possibility of parity. *Ethics* 112(4), 659–688.
- Chen, E. K. and D. Rubio (forthcoming). Surreal decisions. *Philosophy and Phenomenological Research*. [URL: <https://doi.org/10.1111/phpr.12510>].
- Dietrich, F. and B. Jabarian. Decision under normative uncertainty. Unpublished manuscript, September 2018.
- Gracely, E. J. (1996). On the noncomparability of judgments made by different ethical theories. *Metaphilosophy* 27(3), 327–332.
- Greaves, H. and O. Cotton-Barratt. A bargaining-theoretic approach to moral uncertainty. Unpublished manuscript, November 2018.
- Guerrero, A. A. (2007). Don't know, don't kill: Moral ignorance, culpability, and caution. *Philosophical Studies* 136(1), 59–97.
- Gustafsson, J. E. and O. Torpman (2014). In defence of My Favourite Theory. *Pacific Philosophical Quarterly* 95(2), 159–174.

- Harman, E. (2015). The irrelevance of moral uncertainty. In R. Shafer-Landau (Ed.), *Oxford Studies in Metaethics*, Volume 10. Oxford: Oxford University Press.
- Hedden, B. (2016). Does MITE make right? On decision-making under normative uncertainty. In R. Shafer-Landau (Ed.), *Oxford Studies in Metaethics*, Volume 11. Oxford: Oxford University Press.
- Lockhart, T. (2000). *Moral Uncertainty and Its Consequences*. New York: Oxford University Press.
- MacAskill, W. (2014). *Normative Uncertainty*. Ph. D. thesis, University of Oxford. [URL: https://www.academia.edu/8473546/Normative_Uncertainty].
- MacAskill, W. (2016). Normative uncertainty as a voting problem. *Mind* 125(500), 967–1004.
- MacAskill, W. and T. Ord (forthcoming). Why maximize expected choice-worthiness? *Nous*. [URL: <https://doi.org/10.1111/nous.12264>].
- Nissan-Rozen, I. (2012). Doing the best one can: A new justification for the use of lotteries. *Erasmus Journal for Philosophy and Economics* 5(1), 45–72.
- Nissan-Rozen, I. (2015). Against moral hedging. *Economics and Philosophy* 31(3), 1–21.
- Podgorski, A. (forthcoming). Normative uncertainty and the dependence problem. *Mind*. [URL: <http://www.abelardpodgorski.com/research.html>].
- Riedener, S. (2015). *Maximising Expected Value Under Axiological Uncertainty: An Axiomatic Approach*. Ph. D. thesis, St. John's College, University of Oxford. [URL: <https://uzh.academia.edu/StefanRiedener>].
- Rosenthal, C. (2019). *Ethics for Fallible People*. Ph. D. thesis, New York University.

- Ross, J. (2006a). *Acceptance and Practical Reason*. Ph. D. thesis, Rutgers University – New Brunswick. [URL: <https://philpapers.org/rec/ROSAAP-2>].
- Ross, J. (2006b). Rejecting ethical deflationism. *Ethics* 116(4), 742–768.
- Saari, D. G. (1990). The Borda dictionary. *Social Choice and Welfare* 7(4), 279–317.
- Sepielli, A. (2009). What to do when you don't know what to do. In R. Shafer-Landau (Ed.), *Oxford Studies in Metaethics*, Volume 4, pp. 5–28. Oxford: Oxford University Press.
- Sepielli, A. (2010). *'Along an Imperfectly-Lighted Path': Practical Rationality and Normative Uncertainty*. Ph. D. thesis, Rutgers University Graduate School - New Brunswick. [URL: <https://rucore.libraries.rutgers.edu/rutgers-lib/26567/>].
- Sepielli, A. (2016). Moral uncertainty and fetishistic motivation. *Philosophical Studies* 173(11), 2951–2968.
- Sepielli, A. (2018). How moral uncertainty can be both true and interesting. In M. Timmons (Ed.), *Oxford Studies in Normative Ethics*, Volume 7. Oxford: Oxford University Press.
- Tarsney, C. (2017). *Rationality and Moral Risk: A Moderate Defense of Hedging*. Ph. D. thesis, University of Maryland, College Park. [URL: <https://drum.lib.umd.edu/handle/1903/19981>].
- Tarsney, C. (2018a). Intertheoretic value comparison: A modest proposal. *Journal of Moral Philosophy* 15(3), 324–344.
- Tarsney, C. (2018b). Moral uncertainty for deontologists. *Ethical Theory and Moral Practice* 21(3), 505–520.
- Tarsney, C. (2019). Normative uncertainty and social choice. *Mind* 128(512), 1285–1308.

- Vallentyne, P. and S. Kagan (1997). Infinite value and finitely additive value theory. *The Journal of Philosophy* 94(1), 5–26.
- Weatherson, B. (2014). Running risks morally. *Philosophical Studies* 167(1), 141–163.
- Weatherson, B. (2019). *Normative Externalism*. Oxford: Oxford University Press.
- Young, H. P. (1974). An axiomatization of Borda’s rule. *Journal of Economic Theory* 9(1), 43–52.