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

How should agents decide what to do when they're uncertain about basic normative principles? Most answers to this question involve some form of *intertheoretic value aggregation*—i.e., some way of combining the rankings of options given by rival normative theories into a single ranking that tells an agent what to do given her uncertainty. An important obstacle to any form of intertheoretic aggregation, however, is the *structural diversity* of moral and other first-order normative theories: The ranking of options that a first-order theory provides may take many structurally distinct forms, including ordinal and interval structures, various kinds of multidimensional, infinitary, and incomplete structures, and more. But it is typically far from obvious how to aggregate rankings with different structures. In this paper, I survey three general approaches to this problem and tentatively defend one of them as least-bad. I also explore the practical implications of this "problem of structural diversity" for decision-making under normative uncertainty and for various other questions in normative philosophy.

I begin by considering what structures a normative theory's assessment of options might have. What is conceptually essential to any normative theory, I argue, is the binary classification of options as *eligible* or *ineligible* for choice. If this is so, then a normative theory can in principle have *any* structure that is capable of inducing such a binary classification. This means that the range of possible normative structures is *extremely* broad—indeed, much broader than the recent literature on normative uncertainty has presupposed. The aggregation of structurally diverse theories therefore calls for a *general* method, one that can generalize beyond particular cases (e.g., cardinal and ordinal theories).

The first and simplest general approach to aggregating structurally diverse theories is *structural depletion*. This approach simply strips theories of all but some minimum, universal structure (e.g. binary or ordinal) for purposes of aggregation. It is is exemplified by the "My Favorite Option" view (discussed by Lockhart (2000), Gustafsson and Torpman (2014), and MacAskill (2014)), according to which normatively uncertain agents should simply choose the option with the greatest probability of being permissible. The major drawback of structural depletion is, of course, that we end up ignoring a great deal of seemingly relevant information—e.g., the relative degrees of choiceworthiness given by cardinal theories. Moreover, structural depletion theories must either draw an arbitrary and inelegant line between normative and empirical uncertainty or else embrace the extremely implausible implication that rational agents should be unresponsive even to cardinal degrees of risk generated by *empirical* uncertainty.

Structural enrichment approaches take the opposite tack, adding new structure to theories, e.g. by mapping ordinal rankings onto a cardinal scale. This approach is exemplified by the view developed in MacAskill (2014), which converts ordinal rankings into their cardinal Borda scores, normalizes those Borda scores with the value scales of cardinal theories by equalizing their respective variances, and then simply calculates expected value. The major drawback of such an approach is its apparent arbitrarity: There are a plethora of ways to map ordinal rankings onto a cardinal scale and more

generally to map any given sparse structure onto a richer structure, none of which is obviously privileged over the rest. In the context of a search for basic principles of rational choice, worries about arbitrarity are especially powerful, since given a genuinely arbitrary choice between decision-procedures, it is implausible that rationality should *require* us to adopt one rather than the other.

The third approach to combining structural diverse theories is *multi-stage aggregation*. On this approach, we begin by aggregating classes of identically-structured theories, then takes the result as input to one or more further stages of aggregation that combine larger classes of more distantly related theories. This results in an aggregation procedure that is sensitive to the information provided by more richly-structured theories (e.g., sensitive to cardinal degrees of choiceworthiness), without imposing alien structure on sparsely structured theories. A substantial drawback of multi-stage aggregation, however, pointed out by MacAskill (2014), is that it results in violations of a condition he calls *Updating Consistency*: Increasing your credence in a theory that regards option O as more choiceworthy than any alternative, while keeping the ratios between your credences in all other theories the same, can cause O to *drop* in the aggregate ranking.

Despite this cost, however, I argue (albeit tentatively) that multi-stage aggregation remains a more promising strategy than structural enrichment or structural depletion. In particular, I argue that proponents of both structural enrichment and structural depletion are already committed to a form of multi-stage aggregation, and to the resulting violations of Updating Consistency, in choice situations that involve *both* empirical and normative uncertainty. The only way to satisfy Updating Consistency is to adopt a one-stage decision procedure that aggregates over complete empirical + normative states of the world. But this results in an enormous misrepresentation of the first-order normative theories in which agents may have credence—e.g., replacing all normative theories with cardinal and expectational surrogates and thereby effectively eliminating the possibility not only of non-cardinal theories but also of theories that endorse non-expectational decision procedures like risk-weighed expected utility (Buchak, 2013) or maximin.

Though the *comparative* downside of multi-stage aggregation is therefore less than has been supposed, all three approaches undeniably have serious drawbacks. This "problem of structural diversity" needs more attention, both since it represents a serious challenge to the possibility of intertheoretic aggregation and since whatever approach we adopt will substantively constrain other aspects of our metanormative theories (e.g. ruling in or out views like My Favorite Option or the Borda method). It also has ramifications beyond metaethics, e.g. in the design of social choice mechanisms for group decision-making in the face of normative *disagreement*, or for aggregating the preferences of diverse individuals, some but not all of whom satisfy canonical axiom systems like von Neumann-Morgenstern that enable cardinal representation.

References

Buchak, L. (2013). *Risk and Rationality*. Oxford: Oxford University Press.

Gustafsson, J. E. and O. Torpman (2014). In defence of My Favourite Theory. *Pacific Philosophical Quarterly 95*(2), 159–174.

Lockhart, T. (2000). *Moral Uncertainty and Its Consequences*. New York: Oxford University Press.

MacAskill, W. (2014). *Normative Uncertainty*. Ph. D. thesis, University of Oxford.