

DISSERTATION ABSTRACT

ESSAYS ON AGGREGATION IN DELIBERATION AND INQUIRY

RUSH T. STEWART

1. INTRODUCTION

In *Essays on Aggregation in Deliberation and Inquiry*, I study models of aggregation for probabilities, full beliefs, and preferences. Aggregation matters for finding consensus, neutral positions, or agreement, which serves crucial functions in rational deliberation and inquiry for both individual agents and groups. While it may be obvious that neutrality or consensus matters in settings of groups of decision makers or inquirers, such notions are also crucial for single agents. A rational agent can have multiple goals or values that come into conflict or she may suspend judgment between various ways of evaluating events with respect to subjective probability, for example. Such cases require an account of neutrality. So an account of neutrality is important for decision theory and epistemology even when the focus is on individual agents. *Essays on Aggregation in Deliberation and Inquiry* is a sustained exploration of taking indeterminacy—in value, in probability—seriously in the context of aggregation. My dissertation comprises six papers, naturally grouped into three sections: probability, full belief, and preference.

2. PROBABILITY

In “Probabilistic Opinion Pooling with Imprecise Probabilities” (a version of which is forthcoming in *The Journal of Philosophical Logic*), I argue that collective opinion is more properly represented by imprecise probabilities (IP) than by a numerically determinate probability function as on the standard account. I provide three arguments. First, if pooling is interpreted as reaching a *consensus* in probabilistic opinion, IP pooling is on firmer philosophical ground. IP models allow for suspending judgment between some number of probability distributions by not ruling them out for use in deliberation and inquiry, and reflect the common ground among the group concerning which probability distributions *are* ruled out. Such a consensus constitutes a neutral initial position from which to launch further inquiry. Precise pooling functions, on the other hand, do not allow for an analogous suspense of judgment, and may yield collective probabilistic opinions endorsed by none of the group members. Second, there are various IP pooling functions that each jointly satisfy more of the standard pooling axioms that have been discussed in the statistical literature than *any* precise pooling recipe. Third, in the IP setting, the tension between a pooling method’s being justified on epistemic or procedural grounds—reflected in the tension between satisfying certain formal epistemic and procedural constraints—dissipates, an artifact of the assumption of precision.

“Learning and Pooling, Pooling and Learning” explores which types of probabilistic updating or learning commute with the proposed IP pooling methods of the previous chapter. Commutativity of pooling and learning, which requires that the result of aggregating and then updating is the same as first updating and then aggregating, matters for avoiding diachronic Dutch books, for safeguarding against certain other manipulations of group opinion, and for views according to which groups can be agents. Positive results are stated for Bayesian conditionalization (and a generalization of it), imaging, and a certain parameterization of Jeffrey conditioning. This last observation is obtained with the help of a generalization of a characterization of precise externally Bayesian pooling operators due to Carl Wagner to the IP setting. These results strengthen the case that pooling should employ imprecise

probabilities since no precise pooling method is as versatile. These first two chapters relating to probability rehearse and expand the arguments and proposals presented in two first-authored papers with a fellow graduate student in the philosophy department at Columbia.

“Convergence” and “merging of opinions” results present a distinct and pedigreed notion of consensus that envisions a community of inquirers—with access to an idealized, infinite stream of data—updating on shared evidence to reach agreement. My third chapter examines such results in the setting of imprecise probability and aggregation, relating pooling to this other approach to consensus. In Chapter 2, I prove various commutativity results for learning and pooling in the IP setting. In Chapter 3, “Pooling and Learning, and Learning, and Learning...,” I extend some merging of opinions results for different probabilistic “learning” rules in the context of pooling with imprecise probabilities.

3. FULL BELIEF

In two other papers, I explore the limitations of belief revision theories for treating consensus. Belief revision theories deal with full beliefs on their own terms, without presupposing a bridge or reduction to degrees of belief. Given the role consensus is supposed to play in the social aspects of inquiry and deliberation, it is important that we can always identify a consensus as the basis of joint inquiry and deliberation. But in Chapter 4 (“Unanimous Consensus against AGM?”), I show that if we think of an agent revising her beliefs (genuinely or for the sake of the argument) to reach a consensus, then on the received view of belief revision, AGM belief revision theory, certain simple and compelling consensus positions are *not* always available.

That account of consensus, *shared agreement*, is a conservative account of consensus in full belief. In a sense, it is the analogue of the conciliatory position in the literature on peer disagreement for *sets* of belief. A shift to consensus as shared agreement suspends judgment on beliefs about which parties do not agree. Isaac Levi is among those who advocate it, and in Chapter 5, I present a constructive challenge to articulate motivations for that account in terms of Levi’s own theory of justified belief change: if the standard (epistemic) goals of revision do not apply to shifting to consensus *as he claims*, how is shifting to consensus justified? I go on to argue that consensus not only fails to motivate an affirmative answer to the question of *whether* to revise, but revising to consensus violates principles of *how* to revise (on Levi’s preferred account). In particular, postulates of *mild contraction* are generally violated in revising to consensus. I argue that the most promising response to the challenge is that, in general, a revision to consensus as shared agreement (when rational) should be thought of as a hypothetical revision, not a genuine one. Hypothetical or *for the sake of the argument* revision is not subject to the same constraints as genuine revision.

4. PREFERENCE

Social choice theory is famously beset by some very substantial limitative results such as Arrow’s Impossibility Theorem and Sen’s Liberal Paradox. Such results are sometimes thought to be devastating for the prospect of reasonable democratic collective choice. I propose a framework for consensus in preference that allows for *indeterminacy* and show that we can skirt both Arrow’s and Sen’s results. In other words, the framework allows more desiderata for preference aggregation to be jointly satisfied than determinate aggregation methods consistently can. Furthermore, the proposed account offers a way of resolving not only some important conceptual tensions in democratic theory between social choice theory and deliberative democracy, but also some tensions in welfare economics between consequentialist and procedural constraints on aggregation. Finally, the account admits philosophical motivation in terms of the same conception of neutrality between rival points of view that underwrites my approach to probability pooling and belief aggregation. I present the case in this chapter as a proof of concept for indeterminacy in the context of social choice.