

## **Abstracts**

### **5th CPH LU Workshop on Social Epistemology**

**6-7 DEC 2012, Lund**

**<http://www.fil.lu.se>**

#### **Aviezer Tucker (University of Austin, Texas): Modeling the generation of knowledge from multiple testimonies**

Bayesian and other formal epistemologists have known at least since Laplace's pioneering application of Bayesian probability to the epistemology of testimony that when multiple independent testimonies, even of low individual reliability, share their propositional content, and the prior probability of the truth of this content is sufficiently low, they generate highly probable beliefs, knowledge. Though Bovens & Hartman and Olsson do not refer to Laplace in their books, their formal approaches are sophisticated versions of his basic insight. Still, articulating the relevant concepts of independence, testimony, reliability, and the precise relevant relation between the testimonies has proven more difficult than their formal representations and modeling because it is necessary to clearly conceptualize:

1. What are the relations between the multiple testimonies that group them together to generate knowledge?
2. What is the meaning of the independence of testimonies that is necessary for the generation of knowledge from multiple testimonies?
3. How can and do we determine the reliabilities of testimonies?

I start by recapitulating the achievements of the formal analysis of the inference of knowledge from multiple testimonies, most notably in the recent books of Bovens & Hartmann and Olsson and the articles by Shogenji. Then, I proceed to analyze critically the existing literature about the relevant relations between the units of testimony, their reliability, and independence. I agree with Olsson's criticisms of Bayesian coherentism; I argue then against conditional and causal interpretations of independence, as well as against frequentist and Markov conditional interpretations. I agree with Olsson's criticism of the endogenous interpretation of reliability. Instead, I advocate the interpretations of testimonial reliability as the ratio of information that is preserved at the end of a transmission process to the information that was generated and transmitted at the beginning of the process. I analyze the generation of knowledge from multiple testimonies in terms of tracing information flows. Then, I present a new alternative, a three stages modular model of knowledge generation from multiple testimonies that I argue fits the actual veritistic best practices of institutionally embedded professionals who infer knowledge from multiple testimonies, like detectives, historians, intelligence analysts, and investigative journalists. Finally, I hint at some of the larger implications of the analysis I present here for broader issues in social epistemology, and the nature of social knowledge.

## **Gustavo Cevolani (University of Bologna): Belief merging and truth approximation**

In this paper, we investigate the problem of truth approximation via belief merging, i.e., we ask whether, and under what conditions, a group of agents merging together their beliefs makes progress toward the truth about the underlying domain.

We answer this question by proving some formal results on how different operators of belief merging perform with respect to the task of truth approximation, construed, in Popperian terms, as increasing verisimilitude or truthlikeness.

Belief merging can be conceived, like traditional AGM-style belief revision, as a kind of belief change (Konieczny 2011). Theories of belief change aim at specifying how the beliefs or theories of an ideal agent should change in response to new, possibly conflicting, information (Hansson 2011). The problem of belief merging is how different rational agents, who trust each other, should merge or combine their own beliefs, which may be mutually incompatible, in order to agree on a collectively accepted theory.

Theories of verisimilitude specify what does it mean for a given theory to be closer to the whole truth about a target domain than another theory (Oddie 2008). An interesting question is then under what conditions belief change increase the verisimilitude of our theories, i.e., promote cognitive progress about the domain. This question was first raised by Niiniluoto (1999) with respect to AGM belief revision, and has been recently discussed by a number of other philosophers of science (Kuipers and Schurz 2011). In this paper, we extend this kind of analysis to the case of belief merging.

We start by focusing on only two agents, who combine theories  $T_1$  and  $T_2$  into a merged theory  $T_1 \circ T_2$ .

First, we prove some negative results, mirroring the ones obtained by Niiniluoto and others for AGM belief revision and showing that, in general, belief merging does not track truth approximation. More specifically, we prove that if  $T_1$  and  $T_2$  are both true, then  $T_1 \circ T_2$  is true and more verisimilar than  $T_1$  and  $T_2$ . If however either  $T_1$  or  $T_2$  is false, then  $T_1 \circ T_2$  may well be less verisimilar than both theories, even if both  $T_1$  and  $T_2$  are approximately true or highly verisimilar.

Second, we provide some reasons to be more optimistic about the possibility of approaching the truth via belief merging. In particular, we prove the following theorem: if, in a suitably defined sense, "the truth" is "between"  $T_1$  and  $T_2$ , then  $T_1 \circ T_2$  is true, even if  $T_1$  and  $T_2$  are false and quite distant from the truth.

Third, we generalize our analysis in two directions: we consider the general case of merging three or more theories; and we take into account more specific models of belief merging, including "integrity constrained" belief merging and "majority" vs "arbitration" merging operators (Konieczny 2011).

Finally, we show how our results may shed new light on some hot topics in formal social epistemology and related areas, like the problem of truth tracking in belief merging (Everaere et al. 2010) and judgment aggregation (Pigozzi 2006, Hartmann 2010), and the role of peer disagreement in truth approximation (Douven 2011).

## **Peter Brössel (University of Bochum): Degrees of belief, confirmation commitments, and peer disagreement**

Peer disagreement is in many respects one of the biggest problems of the Bayesian approach to social epistemology. On the one hand we have to answer the question under which conditions an epistemic agent should revise her belief state in the light of peer disagreement, respectively whether a given case of epistemic disagreement is a case of peer disagreement. On the other hand we have to answer the question how the agent should revise her belief state in such a case. Since almost everyone agrees that there are cases of peer disagreement in which an agent should change her belief state, the latter problem can arguably be labeled the scandal of Bayesian social epistemology. However, the problems associated with the second task (inconsistencies between various very plausible rationality requirements on peer update) arise under the assumption that an agent's epistemic belief state is best represented by her degree of belief function. In this paper I first argue that for modeling cases of peer disagreement Bayesian agents should better be represented by their confirmation commitments and the evidence available to them. Second, I will discuss in how far known inconsistencies between rationality constraints on peer update can be resolved if we model the epistemic belief states of Bayesian agents as suggested.

## **Fabrizio Cariani (Northwestern University): Why deliberate?**

This paper studies the epistemic value of collective deliberation. Much recent work in social epistemology has aimed to explain why groups should aggregate their opinions via democratic procedures (see, e.g. the epistemic analysis in List and Pettit 2011). It is notable that the Condorcet-style analysis that sustains these arguments does not explain why we design groups with the ability to deliberate.

The lack of an explanation for the value of deliberation may turn into full-blooded skepticism once we consider evidence from experimental psychology. If polarization (see e.g. Myers and Lamm 1975) and irrational belief-forming practices are as pervasive as some psychological research suggests, collective deliberation may well turn out to have epistemic perils rather than benefits.

Yet this skeptical position should be surprising: collective deliberation is a cornerstone of ordinary and scientific inquiry. Many important groups (parliaments, families, scientific committees) are organized so as to have the ability to deliberate. Before caving in to skepticism about the epistemic value of deliberation, we should explore alternatives.

I develop a theory of deliberation that can acknowledge the facts behind these challenges but also clarify the epistemic role of deliberation. The facts that the skeptic's argument relies upon show that we cannot expect deliberation to *guarantee* any epistemic values. However, the emphasis on "guaranteeing" may be misplaced. It is enough, I argue, if we can show that deliberation *promotes* certain values. "Promotion" here simply means that, under suitably advantageous conditions, those values are more likely to be instantiated by designing a deliberative group rather than by designing a merely aggregative group.

In this paper, I explore the hypothesis that one of the central epistemic values that deliberation promotes is *understanding*. *Understanding* is an important epistemological category, but it has been vastly neglected in contemporary epistemology (except

for virtue epistemology, as in Kvanvig (2003), which, on the other hand, hasn't been in much contact with the sort of formal model I aim to construct).

In order to sustain the argument that deliberation promotes understanding, one could try and operationalize the concept of understanding and then prove that deliberation, under suitable conditions, can promote the operationalized concept. It is doubtful, however, that understanding itself could be operationalized.

For this reason, I propose a less direct argument, articulated in three steps:

- (i) operationalize specific dimensions of understanding
- (ii) prove that deliberation can promote these operationalized dimensions; and
- (iii) leave it up to a philosophical inference to the best explanation to sustain claims about deliberation promoting understanding.

The two dimensions of understanding for which I provide a formal analysis are: (1) having better or worse reasons in favor of one's beliefs and (2) carving logical space in better or worse ways. I construct formal models for these dimensions and prove that, under appropriate conditions, deliberation can promote them.

List and Pettit (2001), *Group Agency*, O.U.P.

Kvanvig (2003), *The Value of Knowledge and the Pursuit of Understanding*. C.U.P.

Myers and Lamm (1975), "The Polarizing Effect of Group Discussion", *The American Scientist*, 63(3), pp. 297-303.

## **Chris J. Thompson (Cambridge University): An epistemic case for deliberative democracy**

Epistemic democracy involves two claims: firstly, that at least some political decisions are about matters of fact and therefore can be correct or incorrect; and, secondly, that some features of democratic decision-making procedures make them reliable at making correct decisions. Claims of epistemic democracy require some mechanism or justification for why it is that democratic decision-making can be reliable. Most existing accounts look to judgment aggregation procedures such as majority rule and the classic Condorcet Jury Theorem (CJT) to provide the required social epistemic mechanism. Very informally, majority rule serves as a mechanism to 'pool' the information, dispersed across the electorate, into the social choice.

To the extent that advocates of epistemic democracy have considered deliberation, the assessment has tended to be negative. Sunstein, for example, points to both informational failures of deliberation (such as information cascades); and social failures (such as group think). The idea that deliberation could act as an information pooling mechanism, alongside majority rule, is dismissed in favor of mechanisms such as information markets.

I argue that while deliberative democracy is subject to some (though not all) of the problems highlighted by Sunstein, deliberation has an essential role as an information pooling mechanism in at least some social choice cases.

To make the epistemic case for deliberation we must firstly fill out the details tacitly assumed in much of the literature on judgment aggregation and the CJT. Firstly, agents extract information from the environment about the correct alternative on the agenda. Secondly, these pieces of information act as causal factors on the judgments of agents – the information makes agents more or less likely to vote correctly. Thirdly,

when agents express their judgments via their votes and the votes are aggregated into the social choice, the private information found by individuals is shared with the group.

There are two cases where deliberation may be an essential information pooling mechanism. Firstly, it may be the case that agents extract relevant information from the environment, but this information does not act as a causal factor on that agent's judgment. We may require deliberation to generate epistemic closure on propositions, so that the conclusion deduced from the conjunction of propositions may act as a causal factor on agent's votes.

Secondly, there may only be a limited amount of evidential information. If this information only influences the judgment of a single agent, the probability of a correct social choice will be low. However, if the evidential information is shared with a wider group via deliberation, agents can draw implications from the evidence by applying their private background information. In expressing their judgments agents thereby share relevant background information with the wider group.

Thus, although standard judgment aggregation procedures such as majority rule can provide an epistemic justification for democracy in many cases, there are other cases where group deliberation will be an essential social epistemic mechanism.

### **George Masterton (Lund University): Does Laputa solve Goldman's computational problem?**

In a contribution to the field of social epistemology Goldman (1999) proposed that collective epistemic means should be evaluated according to their *veritistic value*. This proposal faces a serious difficulty known as the computational problem: it is extremely difficult to calculate the veritistic value of any non-trivial inquisitional practice. Olsson (2011) has proposed that this practical problem can be addressed by using computers to simulate such practices. As an example of such an approach, Olsson offers the agent based model Laputa. But does Laputa really solve the computational problem of social epistemology? Herein, Laputa's contribution to solving Goldman's problem is thoroughly reviewed. It is found that, despite much promise, problems surrounding the specification of inquisitional social practices currently prevent Laputa from achieving this goal. Prospects for extending Laputa beyond this limitation are considered and judged to be feasible though challenging. Consequently, though Laputa currently fails to be a generally applicable answer to the computational problem, there are foreseeable developments of Laputa that might realise this promise.

### **Bert Baumgärtner (University of California at Davis): An agent-based model of Millian mingling**

In this work I focus on a social epistemic practice I call "Millian Mingling", which is inspired by a suggestion made by John Stuart Mill (1848): "It's hardly possible to overstate the value, in the present state of human improvement, of placing human beings in contact with other persons dissimilar to themselves, and with modes of thought and action unlike those with which they are familiar." It is often said that making good societal decisions requires communities to have some degree of heterogeneity in opinion, and epistemic enterprises are no different: scientific inquiry is thought to benefit from interactions between various viewpoints.

I give an evaluation of Millian Mingling with respect to "echo chambers", a sociological setting where people's prior beliefs are "echoed back" by other like-minded people, giving the (false) impression that their beliefs are correct. Not surprisingly, echo chambers tend to be regarded negatively. That makes them a good metric for the evaluation of Millian Mingling as an epistemic practice.

I model echo chambers with a kind of "peer"-alignment; agents that interact with one another very frequently will form the same beliefs. This occurs because agents with highly "entrenched" beliefs act as attractors in the model. I then use an agent-based model, which affords an explicit space, to implement Millian Mingling and two other practices. Some agents, which I call Millians, seek out agents that are different than themselves with respect to their beliefs. Other agents, which I call Chums, seek out agents that they already share some beliefs with. In the middle, we have Neutrals, which are agents that seek out other agents impartially.

I then do a two part analysis of Millian Mingling. First, I evaluate how well Millian Mingling does with respect to the other two practices by using average times to convergence. Here, practices that resist echo chambers are taken to be those that take longer to converge (on average). And not surprisingly, populations consisting mainly of Chums align themselves more quickly than populations consisting mainly of Millians and Neutrals (if they align themselves at all).

The second part of the analysis involves "veridical"-alignment, the idea that the way things are can also shape our beliefs. Epistemic practices ought not to just avoid mere peer-alignment, but should be sensitive to "the truth" (even if that happens to coincide with peer-alignment). I model veridical-alignment by introducing designated agents whose beliefs (the 'true' ones) are maximally entrenched. Initial results suggest that Neutrals do the best when they are the majority of the population. However, Millians tend to do better than Neutrals when the population has a high proportion of Chums.

In sum, Millian Mingling is a social practice that seems to have high epistemic value when a large proportion of the population is easily subject to echo chambers. If it weren't for the possibility of echo chambers, however, it seems that a practice of impartiality is better. This is suggestive of the quote above when Mill said, "..., in the present state of human improvement, ...".

## **Rasmus K. Rendsvig (Copenhagen University): Dynamics of the Bystander Effect**

The bystander effect is a phenomenon well-known from social psychology, where the tendency for less bystanders to come to aid in emergency situations the more bystanders are present is well documented. In this talk, the dynamics of one of the attributed causes, uncertainty about the situation laced with pluralistic ignorance, is modeled using Dynamic Epistemic Logic. The model utilizes epistemic plausibility models updated by event models allowing for ontic fact change, includes a merge operation that allows agents to form beliefs based on social proof, as well as an analysis of pluralistic ignorance seen as a difference between in fact used decision rules and applied action interpretation rules

## **Raphael Künstler (University of Provence): Information cascade and the division of cognitive labor**

If some of its assumptions are relaxed, the knowledge-seeking model of inquiry offered by Hintikka (2007) offers a convenient framework to deal with the threat of an information cascade effect in the division of cognitive labor. According to this model, inquiry is a strategic activity, which consists in deciding, on the ground of given data, whether a rational agent should draw an inference or question an “oracle”. If, instead of assuming that this oracle function is played by the nature, we assume that this function is endorsed by a cognitive agent, this model can be used to raise problems that concern social epistemology. The Hintikka's theory of presupposition can then be used to articulate a 'co-enquiry' relation : the result of one's questioning process can be used by someone else as a presupposition to formulate a new questions or to draw an inference. In such a situation there is an asymmetric epistemic dependence relation between the co-enquirers. Now, if we relax Hintikka's assumption according to which an oracle is always giving the right answer, the situation of epistemic dependence can bring about an information cascade phenomenon. One effect of this phenomenon within the context of enquiry is the formulation of empty questions: accepting a false presupposition leads to raise not answerable questions.

In my paper, I intend to expose and evaluate two kinds of division of the cognitive labor that aim at avoiding or minimizing the risks of formulating empty questions.

The first model is a sequential one. It is defined by a simple rule that prohibit to raise a question that bears on an uncertain presupposition. As a consequence, at a given moment, all the research forces should be devoted to answer only one question, until a certain answer is found. However, this classical way of managing the organization of collective cognitive work suffers a shortcoming: it ignores the fact that theoretical decisions are not only grounded on coherence considerations, but also on research agendas (Olsson & Westlund, 2006). Some questions are too urgent to be avoided until comes their logical turn.

In the second model, a question can be raised before the truth of its presuppositions is beyond any doubt. In order to prevent the occurrence of an information cascade, it uses a “parallel” pattern of work organization: it distributes the work forces among differentiated tasks so that for every question of a group of subordinated questions, there should be a group of individuals that investigates this question. However, this kind of organization implies that some groups will be raising empty questions : this seems unacceptable, from the collective as well as individual good points of view. Bearing on the case of climate change, my paper aims at offering a defense of this second model.

Hintikka (2007) *Socratic Epistemology*. Cambridge University Press.

Olsson & Westlund (2006) “On the Role of the research Agenda in Epistemic Change”. *Erkenntnis* 65: 165-183.

## **Carlo Proietti & Erik J. Olsson (Lund University): Pluralistic ignorance, informational cascades and higher order beliefs**

Recent research in social psychology and behavioral economics witnesses many "Bayesian" explanations/reconstructions of phenomena of "herd behavior". Informational cascades - like buying a best seller book on the basis of a good review - or pluralistic ignorance - e.g. thinking that the stock market is doing all right just because nobody is worried - are explained there by dynamics that combine subjective credence about a given fact and immediate external evidence provided by explicit behavior of other individuals. These explanations mostly disregard individual opinions about others' reasoning patterns as being non important in these dynamics. Is it really so? How could we eventually test relevance? Here we try to outline some possible research questions along these lines.