

Editors' introduction: social dynamics and collective rationality

Frank Zenker · Carlo Proietti

Received: 15 January 2014 / Accepted: 15 January 2014 / Published online: 22 March 2014
© Springer Science+Business Media Dordrecht 2014

Abstract We provide a brief introduction to this special issue on social dynamics and collective rationality, and summarize the gist of the papers collected therein.

Keywords Belief polarization · Belief merging · Debiasing · Doxastic disagreement · Dynamic epistemic logic · Echo chambers · Group topology · Informational cascades · Pluralistic ignorance · Truth approximation

1 Social dynamics and collective rationality

This issue collects ten papers selected from five international *Copenhagen Lund Workshops in Social Epistemology*, held between December 2010 and December 2012, that were made possible through a research grant to Vincent F. Hendricks and Erik J. Olsson by the Einar Hansen Research Fond and RQ8-funds from Lund University. Besides providing a platform for basic-level research, workshops focused on applied research related to social dynamics that lead group members to influence each other's beliefs and choices in ways that can produce not only suboptimal but genuinely catastrophic, and thus pathological, collective outcomes.

The specific phenomena treated here under applied research are known as *pluralistic ignorance*, *belief polarization*, *informational cascades*, and *echo chambers*. Constituting genuine puzzles of collective rationality, they prompt social epistemologists to ask, among others, whether some degree of irrationality, broadly conceived, is nec-

F. Zenker (✉) · C. Proietti
Department of Philosophy and Cognitive Science, Lund University,
Kungshuset, Lundagård, 22222 Lund, Sweden
e-mail: frank.zenker@fil.lu.se

C. Proietti
e-mail: carlo.proietti@fil.lu.se

essarily presupposed or whether they may arise also among perfectly rational agents. Originally studied primarily by social psychologists and economists, until recently such phenomena by and large fell outside the scope of social epistemology—let alone traditional epistemology with its focus on individual agents' knowledge and belief, rather than that of collectives.

The transition from individual to collective phenomena generally goes along with added complexities which appear to be inherent to group interaction. Some of these complexities have now been successfully addressed. At the same time, however, it is safe to say that such phenomena continue to constitute challenges both for conceptual analysis and computational modeling.

In what follows we summarize the gist of the ten papers, which comprise informal analyses as well as formal probabilistic and logical approaches. We consider the first four as instances of basic research, the latter six provide more applied work. As editors, moreover, we remain indebted to a number of anonymous reviewers who assisted in selecting workshop contributors and provided selected participants with critical feedback on their submissions to this issue.

2 Overview

Peter Brössel and Anna-Maria Eder (*How to Resolve Doxastic Disagreement*) address how individuals in a group should rationally revise their beliefs in cases of disagreement. An adequate answer would obviously also provide an important norm to guide collective decision-making. Analyzing the standard framework of *Monistic Bayesianism*, the main challenge is found to be the formulation of an aggregation-rule for individual credences that fulfills the following normative requirements: irrelevance of alternatives, convexity, unanimity, commutativity with learning, no-zero preservation, and preservation of independence. Previous work had already shown that these requirements cannot be jointly satisfied unless “dictatorship” is embraced. Additionally, the authors show that certain disagreements cannot even be expressed within the standard framework, and therefore opt instead for *Pluralistic Bayesianism*—so called because two probability functions are essential for describing an individual's epistemic state. More specifically, an individual's epistemic state is now represented not just by her rational credences, but by an ordered pair consisting of her confirmation commitments and her total evidence. Moreover, the individual's rational credences are obtained by conditionalizing her confirmation commitments on the total evidence available to her. The revised framework permits formulating an aggregation rule that fulfills most, but not all, normative requirements. Eder and Brössel go on to argue that the non-fulfilled normative requirements are in fact not essential.

Gustavo Cevolani (*Truth Approximation, Belief Merging and Peer Disagreement*) investigates a social epistemological issue that also touches upon philosophy of science and political philosophy. It is often assumed that rational disagreement among several opinions, or among several opposing scientific explanations, has a (positive) truth-conducive character. But attempts at justifying this contention raise at least two questions: (i) How should a group of recognized epistemic peers combine their beliefs to obtain a theory that is both consistent and represents their individual opinions in the most faithful way? (ii) How should one assess whether the new collective theory is

closer to the truth? (i) and (ii) translate into assessing the conditions under which the merging of different theories, T_1, \dots, T_n , tracks approximation to a given truth. On the basis of extant limiting results for standard theories of belief change, Cevolani shows that belief merging only retains approximate truth under trivial preliminary conditions. To obtain a more explicit characterization of the conditions under which belief merging approximates truth, he elaborates on the verisimilitude of a theory within the so-called basic feature approach (developed jointly by Roberto Festa, Vincenzo Crupi, and himself). As a new result, this modified perspective yields the outcome of merging theories T_1 and T_2 as *more* verisimilar (than the initial T_1 and T_2) when, insofar as their “basic” claims are concerned, weak disagreement outweighs strong disagreement.

Fenrong Liu, Jeremy Seligman and Patrick Girard (*Logical Dynamics of Belief Change in the Community*) take a step towards combining the study of abstract social networks and dynamic epistemic logic (DEL) within the theoretical framework that the same authors had inaugurated with “Logic in the community” (2009), then carried through in a series of papers on *Epistemic Friendship Logic* (EFL). Their present contribution introduces “an account of how standard models of belief revision ... can be extended to models of social influence on one’s belief.” Exploring various options for modelling social influence, the authors lay out research lines in what is “an open-ended survey of some of the possibilities for modeling the way in which people’s beliefs are influenced by their social relations.” Specifically, they discuss the “stability” of individual belief in a community of agents and how it depends both on the network structure and the different norms for belief revision that agents are endowed with. The paper also provides an extensive appendix that sees the authors compare their modelling with more established research on social networks such as models of diffusion, learning, and games on networks.

George Masterton (*Topological Variability of Collectives and its Import for Social Epistemology*) critically engages with the presuppositions of studying the “effect of social interaction on the epistemic properties of a group;” particularly the computational study of group topology against the background of evaluating the veritistic value—which, roughly, is the value placed on leading to true rather than false beliefs—of social practices. Because social practices are variably implementable, it stands to reason that a practice’s veritistic value could depend on the group topology, e.g., who in the group can communicate, in a direct or a mediated way, to other group members. Hence, certain topologies may have a comparatively greater veritistic value than others, *ceteris paribus*. As Masterton demonstrates, however, the sobering truth is that “[n]o matter how one approaches the task of sampling social networks, the task of evaluating such topologies with respect to their veritistic value is computationally challenging. Indeed, depending on specifics the task may be beyond what is practically feasible given present computing technology.” Consequently, though such computational studies bear useful results for networks with “low single digit numbers of participants,” they are prone to generate misleading epistemic evaluations when applied to larger groups. This, as the author argues, severely limits social epistemology as a tool for evaluating the epistemic implications of social practices.

Jens Christian Bjerring, Jens Ulrik Hansen, and Nikolaj Jang Lee Linding Pedersen (*On the Rationality of Pluralistic Ignorance*) investigate the socio-psychological phe-

nomenon of pluralistic ignorance. This occurs “when a group of individuals all have the same attitude towards some proposition or norm, all act contrary to this attitude and all wrongly believe that everyone else in the group has a conflicting attitude to the proposition or norm.” Here, two questions are fundamental: (i) what exactly is the phenomenon of pluralistic ignorance? (ii) Can the phenomenon arise among perfectly rational agents? To answer the first, the authors scrutinize instances of the phenomenon from the literature. They demonstrate the (complete or partial) inadequacy of current definitions by providing counterexamples, i.e., fictive situations that cannot be cases of pluralistic ignorance although they fit extant definitions. According to their own characterization, pluralistic ignorance refers to a situation where the individual members of a group (i) all privately believe some proposition p ; (ii) all believe that everyone else believes *not* p ; (iii) all act contrary to their private belief that p (i.e., act as if they believed *not* p); and (iv) all take the actions of the others as strong evidence for others’ private beliefs about p . As the authors recognize, rationality may come in different shapes such as *epistemic* rationality or *pragmatic* rationality. Bypassing this concern, however, they argue that pluralistic ignorance may arise among perfectly rational agents.

Rasmus K. Rendsvig (*Pluralistic Ignorance in the Bystander Effect*) offers a step-by-step reconstruction, followed by a logical analysis, of the process that (often) leads larger groups of bystanders to fail to intervene, say, on behalf of a victim of an incident—an unfortunate social mistake that appears to be frequent especially in urban contexts. Contrary to the popular tendency of attributing such events to bystanders’ apathy, he pursues an alternative explanation (owed to Latané and Darley) according to which those in doubt about the gravity of an incident first observe what others do. Should everybody do so, then collective inaction must result in a shared misperception of the situation. This condition, in turn, may constitute a reason for everyone to evade the situation. Rendsvig then proposes a formal modelling based on “augmented” models of dynamic epistemic logics (DEL). Such models describe changes of factual states such as the incident as well as the relative changes in an agent’s beliefs. He then shows that DEL can encode—through specific model transition rules—different types of agents (e.g., first responders, city dwellers, and hesitators). Should every agent follow the hesitator’s rule, then a state of *pluralistic ignorance* is generated. In such circumstances, agents decide how to act (intervene or evade) upon observation of others’ behavior, and by relying on so-called *social proof*—likewise encoded as an action rule. Results obtained for abstract ideal agents are then confronted with empirical data to show the adequacy of the model. His formalization also clears ground for computer simulations that may support empirical research in this area.

Rogier De Lange (*To Specialize or to Innovate?*) develops an internalist account, and a NetLogo computer model, of pluralistic ignorance in economics, “viz. one that explains the phenomenon purely in terms of factors internal to science.” In what is perhaps an unconventional move, he treats social pressure as being on a par with factors more standardly conceived of as being internal to science such as experiments and proofs, for instance. “[S]ocial factors,” De Lange holds, “can be internalist because they promote scientific specialization and diversity.” He seeks to explain the “boom-and-bust-like features of pluralistic ignorance by a tension [i.e., a goal-conflict] between

specialization and diversity.” His model exploits the numerical difference between the ‘internal value’ of a contribution to an established paradigm (in the Kuhnian sense) and the ‘overall utility’ of this contribution. Suitably defined, the former value is set privately by each agent, the latter is acquired publically. Under further assumptions, De Lange’s computer model behaves as is characteristic for pluralistic ignorance. In particular, the model predicts vast and sustained adherence to a single paradigm as well as a comparatively sudden switch to a new one. Vis-à-vis the recent economic crisis—in which the ‘efficient markets hypothesis’ fell victim to pluralistic ignorance—De Lange also suggests that a deliberate separation of the relevant research communities may serve as a counter-measure to adverse effects expectably brought about by prolonged adherence to a single paradigm.

Jon Robson (*A Social Epistemology of Aesthetics*) addresses belief polarization and echo-chambers as socio-informational pathologies related to the formation and transmission of human aesthetic judgment, thus as hitherto uncharted applications in the contemporary debate on the epistemology of testimony. He investigates and rejects two claims that he dubs normative and descriptive pessimism. According to these, “we should not form aesthetic judgements on the basis of testimony” and, respectively, “as a matter of fact our aesthetic beliefs are not formed on the basis of testimony.” Defending opposing claims (i.e., normative and descriptive optimism), Robson relies on empirical evidence that “people [do] form aesthetic beliefs on the basis of testimony,” then seeks to support that “in many cases such beliefs are perfectly legitimate.” The empirical side in place, he argues that normative pessimism “commits us to a troubling form of skepticism” because one could at best “know that [aesthetic] beliefs are partially testimonial and partially based on our critical and perceptual judgement.” Normative pessimists then, argues Robson, “will have no reliable way of gauging which of these sources (and to what extent) any given [aesthetic] belief is dependent on.”

Tim Kenyon (*False Polarization: Debiasing as Applied Social Epistemology*) presents false polarization—i.e., the disputants’ overestimation of their difference of opinion that is owed to a biased perception of the other’s actual position—as “a barrier to reliable belief-formation in certain social domains.” Importantly, the false polarization-phenomenon appears to be expectable independently of interlocutors taking a prior stance on an issue, or not. Moreover, it is prone to leave discussants “stuck” in artificial disagreements that forestall the fruitful exchange of views. Based on extant empirical evidence, Kenyon reveals what might amount to a most troubling insight to those engaged in teaching: “intuitively attractive strategies for debiasing are not very effective, while more effective strategies are neither intuitive nor likely to be easily implemented.” Identifying reasons why some de-biasing strategies are prone to fail while others promise to at least mitigate the phenomenon, he concludes that the “practical epistemology of debiasing [...] turns out to be social in a very rich sense: that is, only in the presence of significant socio-institutional scaffolding is there realistic hope of minimizing biases like false polarization in real life situations.” Kenyon thus suggests that—because a phenomenon such as false polarization is sustained by habits of reasoning (aka biases) and the absence of an effective counteracting infrastructure—a societal effort is needed that not only improves extant de-biasing strategies and makes these more readily available than is currently the case, but that also turns their regular use into a literal habit.

Bert Baumgärtner's agent-based model (*Yes, No, Maybe So*), coded in NetLogo, addresses echo-chambers, i.e., socio-technological implementations (such as niche-websites) that let users only hear their own voices, metaphorically speaking. This leads to levels of belief amplification that would be unexpected were agents to communicate regularly with those who disagree with their prior doxastic attitudes. The model-population consists of designated and ordinary agents, where the former either believe or disbelieve a proposition p , while the latter may additionally be indifferent towards p . Moreover, agents believing or disbelieving p do so with varying degrees of entrenchment (or strength). Designated agents, finally, believe or disbelieve p at an entrenchment level one unit higher than ordinary agents. Consonant with an echo chamber, agents are modelled as updating their entrenchment for p by adopting the entrenchment level that dominates their own, provided agents have the same doxastic attitudes, and they do so until a maximum (called "entrenchment ceiling") is reached. Vis-à-vis these assumptions, Baumgärtner argues for maintaining sufficiently low entrenchment levels, and for keeping populations doxastically sufficiently diverse to avoid, or reverse, echo chambers.