The Social Epistemology of Consensus and Dissent

Boaz Miller
Department of Community Information Systems
Zefat Academic College
boaz.m@zefat.ac.il

Forthcoming in *The Routledge Handbook of Social Epistemology*, edited by David Henderson, Peter Graham, Miranda Fricker, and Nikolaj J.L.L. Pedersen.

Abstract This paper reviews current debates in social epistemology about the relations between knowledge and consensus. These relations are philosophically interesting on their own, but also have practical consequences, as consensus takes an increasingly significant role in informing public decision making. The paper addresses the following questions. When is a consensus attributable to an epistemic community? Under what conditions may we legitimately infer that a consensual view is knowledge-based or otherwise epistemically justified? Should consensus be the aim of scientific inquiry, and if so, what kind of consensus? How should dissent be handled? It is argued that a legitimate inference that a theory is correct from the fact that there is a scientific consensus on it requires taking into consideration both cognitive properties of the theory as well as social properties of the consensus. The last section of the paper reviews computational models of consensus formation.

Consensus plays an increasingly growing role in public life. The National Institute of Health (NIH) and the Intergovernmental Panel on Climate Change (IPCC) formulate expert consensus statements to provide authoritative answers to disputed questions. The "Danish Citizen Consensus Model" is used for assessing societal dimensions of science and technology (Medlock et al, 2007; Horst 2008; Einsiedel et al. 2001). Wikipedia promotes consensus for achieving neutrality and verifiability (Wikipedia 2017). Scientific consensus is deferred to when arbitrating between rival experts in legal trials.

But when we think about the relations between knowledge and consensus, two conflicting intuitions arise. On the one hand, since when is truth determined by numbers? Are we back to the notorious Indiana House of Representatives "Pi Bill" of 1897, which redefined Pi by majority voting? (Hallerberg 1977) Wasn't Galileo right although his views contradicted common wisdom? On the other hand, for every Galileo, there are hundreds of misguided contrarians, forgotten by history. Is everybody wrong and only you are right?

Current social epistemology of consensus and dissent sorts out these conflicting intuitions.[1] Social epistemologists distinguish between knowledge-based consensus and mere agreement, explore the relations between consensus and the aims of inquiry, and distinguish legitimate from illegitimate dissent. Section 1 concerns the attribution of consensus to an epistemic community. Section 2 reviews accounts of knowledge-based or epistemically justified consensus. Section 3 addresses consensus as an aim of inquiry and the normative status of dissent. Section 4 reviews computational models of consensus.

1. When Does a Consensus Exist?

Deciding whether a consensus exists is not always trivial. To do so, we need to assess the scope of and depth of the agreement. Start with scope. Tucker (2003, 509-510) argues that a

[1] In philosophy of science until the early 1980s, scientific consensus was addressed mostly in the context of Kuhn's (1962/1970) *The Structure of the Scientific Revolutions* (See Wray in this volume). For a classic account of the interplay between data, theory, method, and rational consensus, see Laudan (1984).

scientific agreement amounts to consensus only when there is no dissent whatsoever. While a universal consensus is not reasonable to expect, because "[t]here are still people who deny that the earth is round" (2003, 511), it is reasonable to expect complete agreement in a community of experts.

Miller (2013, 1297) argues, however, that in a large enough scientific community, chances are that some members will dissent. For example, there are few scientists who deny that HIV causes AIDS, but we cannot deny the existence of a scientific consensus that HIV causes AIDS. Miller argues that a consensus may be attributed to an epistemic community if the consensus view is well institutionalized, and the dissent view is marginalized.

Oreskes (2007, 66-79) analyzes abstracts extracted from the Web of Science using the search-keyword "global climate change." She codes them into groups, and shows that none falls into the climate skeptical camp. She concludes that a consensus on anthropogenic climate change exists. This method can easily be applied to other cases.

Qualitative methods for identifying consensus require extensive manual labour, and their outcome depends on researchers' subjective judgment in classifying papers. Shwed and Bearman (2010) present an algorithm for identifying a consensus that is not susceptible to these vulnerabilities. The algorithm identifies clusters within a network of citations in a large corpus of scientific papers. Shwed and Bearman assume that a citation of a paper is likelier to show agreement with it than disagreement. Bruggeman et al. (2012) argue, however, that when positive and negative and negative citations are distinguished, other clusters emerge. This criticism may take the wind out of Shwed and Bearman's method, since it reintroduces human labour and susceptibility to subjective judgement (see Shwed and Bearman (2012) for a rebuttal).

Another factor that needs to be addressed is the depth of agreement. Fuller (2002, 208-210) distinguishes essential from accidental consensus. In an essential consensus, a group forms a collective decision for the same thing using shared standards of evidence and sense of relevance. In an accidental consensus, individuals form the same belief on and for their own reasons. A paradigmatic accidental consensus is found among a random group of surveyed people. The subjects may not understand the question in the same way, and by paraphrasing it the extent of consensus may be manipulated. For Fuller, only an essential consensus is a genuine consensus. Against this, Miller (2013, 1302-1303) argues that an agreement on shared formalisms, ontological schemes, and evidential standards is sufficient for attributing a consensus to the community, because it provides a sufficient basis for pursing research together, even if members differ in their metaphysical beliefs.

2. When is Consensus Knowledge-Based or Epistemically Justified?

A common intuition is that a majority or consensus of epistemically-competent agents is likelier to be right than wrong. Several philosophers try to justify it. Jones (2002) argues that ceteris paribus, a scientific consensus is likelier to be knowledge-based than not because science strives at the truth. Anderson (2011) and Collins (2010), while not naive about the influence of biases and interests on the outcomes of scientific controversies, argue that science is still the most reliable way to the truth, thus laypeople have no better alternative than trusting a scientific consensus.

Another rationale for deferring to the majority or consensus comes from Condorcet's jury theorems, which state, roughly, that sufficiently large groups in which there is a sufficiently large subgroup of individuals who have a higher than 0.5 probability to form a correct belief on a given matter will reach the correct decision by majority voting (Odenbaugh 2012). It is unclear, however, in which concrete cases the conditions of statistical

independence and higher than 0.5 probability should obtain, or how to judge whether they obtain (Vermeule 2009, 28-33).

Indiscriminate lay deference to a consensus of experts has its risks. "If the history of science teaches anything, it's humility. There are numerous historical examples where expert opinion turned out to be wrong" (Oreskes 2007, 66). Moreover, a community of experts may mask internal disagreements and present a façade of consensus for reasons like keeping their professional reputation intact and maintaining public trust in science (Beatty 2006). Beatty (2017) further argues that the existence of consensus on a theory does not imply that the theory has been rigorously tested, but that researchers trust each other about their results. The validity of the consensus depends on whether these mutual trust relations are justified (Miller 2015a).

Thus, other scholars recommend deference to a consensus of experts only when there are reasons to think it is *knowledge-based* or epistemically justified. We may distinguish *cognitive* from *social* accounts of knowledge-based or epistemically justified consensus. Cognitive accounts specify empirical and theoretical virtues of a credible or justified consensual view. Social accounts specify societal properties of such a consensus. There are also hybrid accounts, both cognitive and social.

Start with the cognitive approach. Oreskes identifies five general hallmarks of a good scientific theory: inductive support, predictive success, resistance to falsification attempts, consilience of evidence (convergence of evidence of different types), and explanatory success. Oreskes argues that because the theory of anthropogenic climate change bears these hallmarks, the scientific consensus over it is knowledge-based (2007, 79-92).

It is disputed, though, that this is sufficient for rendering a consensus knowledge-based. Miller (2013, 1012) argues that even if a scientific community unanimously accepts a theory that seemingly enjoys strong evidential support, we should still examine the societal properties of the agreement for deciding whether to trust it. Suppose there is a scientific consensus that passive smoking does not raise the chances of lung cancer. Suppose this conclusion is supported by seemingly robust studies of different types, e.g., epidemiological, in vivo, and in vitro studies. If all these studies were supported by tobacco companies, there is still good chance their conclusion is false. The upshot, so Miller argues, is that for inferring truth from social agreement, in addition to being evidentially supported, we would like it to socially diverse in relevant ways, in this case, shared by publicly-funded researchers, smokers and non-smokers, etc.

By contrast, Stegenga (2016) argues that a purely cognitive epistemic evaluation of a consensus is possible. Stegenga identifies three conditions the procedures leading to a knowledge-based consensus should satisfy: (1) *Inclusivity*: the consensus-forming processes should include all available evidence; (2): *Constraint*: these processes should constrain intersubjective assessments of the hypothesis of interest; (3) *Evidential Complexity*: these processes should assess evidence on multiple relevant evidential criteria. Stegenga acknowledges that social diversity is positively correlated with meeting Inclusivity and Evidential Complexity, but insists that it is subordinate to them and not principally necessary. The dispute between Stegenga and Miller is whether social diversity is required for achieving Constraint by limiting subjects' excessive interpretive flexibility in assessing the evidence. While Miller thinks that it is necessary, Stegenga (2016, 46) holds that Constraint can be achieved by more rigorous evidence assessment methods, and that too much diversity actually militates against Constraint.

But cognitive accounts of knowledge-based consensus, such as Oreskes' and Stegenga's, have a more fundamental problem. For Oreskes, the existence of consensus does not actually constitute a reason for (or against) trusting it. For suppose that the theory of

anthropogenic climate change that currently enjoys a consensus were supported by a minority of the scientific community, and a rival theory, which does *not* bear Oreskes' five hallmarks of good science, were supported by the majority. By Oreskes' reasoning, the minority theory would still be more credible, because it would still be better evidentially supported.

Social evidence such as the number of people who support a view, their credentials, and their social background *is* relevant to assessing the credibility of a consensus. But can the credibility of a consensus be assessed by looking *only* at social factor? *Social* accounts of knowledge-based consensus try to do that. Their advantage is that they do not require a layperson to evaluate expert claims, but only to examine social features of the consensus, which are more epistemically accessible to her.

Longino (2002a, Ch. 6) provides an influential social account of knowledge-based consensus, called Critical Contextual Empiricism, which specifies necessary and sufficient conditions for an agreement in an epistemic community to count as collective knowledge. When researchers theorize from empirical data, they may make biased background assumptions that may taint their resulting theories. Critical Contextual Empiricism prescribes four social norms of critical deliberation to expose and eliminate such biases: (1) there are public venues of criticism; (2) there is uptake of criticism; (3) there are publicly recognized standards of evaluation of theories; (4) there is tempered equality of intellectual authority – intellectual capacity and relevant expertise are the only criteria by which people are given the right to participate in the collective critical discussion. Critical Contextual Empiricism echoes Habermas' (1984) "ideal speech situation," in which fully rational agents reach a consensus converging on the truth through free discourse.

Critics argue that Longino's four norms are either too permissive or too restrictive, and neither sufficient nor necessary for knowledge. Goldman (2002) argues that interpreted too permissively, a group of creationists with their own peer-reviewed journals may be said to satisfy Longino's norms. Interpreted too restrictively, evolutionary biologists who do not engage with creationists may be said to fail to meet them. Solomon and Richardson (2005) argue that we have successful knowledge achieved without following Longino's norms, while critical norms do not guarantee actual criticism, thus unwarranted views may survive long. They claim that conditions for knowledge cannot be formulated solely procedurally, they must say something substantive about the conditions the end product that the community reaches must meet. This criticism questions the viability of purely social accounts of knowledge-based consensus.

Empirical evidence about the epistemic performance of deliberative models of consensus is mixed. Habermas (2006) cites psychological studies that show that critical public discussion converges on the truth in real-world, and not just ideal conditions. Sunstein (2006) and Solomon (2015, 92-102) cite other studies according to which deliberation leads to group polarization and extremism. Sunstein and Solomon argue that knowledge from a group is best extracted by aggregating members' independently formed judgements, rather than deliberation within the group (see Tollefsen (2006) for a middle-ground position). Sunstein's and Solomons' views accord with Surowiecki's (2004) theory of the Wisdom of Crowds, according to which individuals' aggregated judgments may outperform experts' judgment (cf. Thicke 2017).

Tucker (2003) gives another social account of knowledge-based consensus. Tucker identifies three conditions under which shared knowledge is *the best explanation* of a consensus. First, the consensus is uncoerced. Second, the consensus is uniquely heterogeneous, namely no subgroup shares an extraneous property that may otherwise explain the agreement within it. Just like in a controlled experiment, if members of a test

group do not share any property other than the one being tested, an observed effect may be attributed to the tested property, when there is a consensus in a group of people who do not share a property such a mutual power relationship, joint interest, shared ideology or bias, the consensus may be attributed to knowledge (2003, 506). This is why a third condition, that the group is sufficiently large, is required. While an accidental consensus in a small uniquely heterogeneous group is likely, it is unlikely in a large heterogeneous group.

Against Tucker, Solomon (2001) and Miller (2013) argue that an accidental aggregation of views toward a non-knowledge-based consensus is plausible. It therefore cannot be eliminated as the best explanation of a consensus even when Tucker's conditions hold. Solomon argues that the history of science shows that consensuses often emerge out of an accidental aggregation of interests and influences on individual members of a scientific community, when there are rival theories that do not get the attention they deserve (2001, 121-35). Miller (2013, 1306-1308) adds that accidental aggregation of influences that leads to a consensus is likely from a probabilistic perspective, because people have many influences on a variety of issues, thus they are likely to happen to agree on some of them. For example, while leaders of the world's orthodox religions disagree on the most fundamental doctrinal issue, they all reject homosexuality. This accidental agreement, however, does not give their view any epistemic credit (for a rebuttal, see Tucker 2014).

Solomon and Miller both offer hybrid accounts of the social epistemology of consensus, which attend to cognitive as well as social factors. In *Social Empiricism* (2001), Solomon addresses the question of when a distribution of research efforts in a scientific community is epistemically justified. Solomon conceptualizes cognitive diversity in terms of factors ("decision vectors") that influence individuals' and communities' theory choice. She distinguishes between empirical and non-empirical decision vectors. Empirical decision vectors are factors that make scientists prefer theories with empirical success. Non-empirical decision vectors are other reasons or causes for theory choice, such as ideology, pride, or preference for simpler theories (2001, 51-63).

Solomon argues that empirical decision vectors should be equitably distributed among theories in proportion to the theories' empirical success. For instance, if some theory has some technological success, a proportional number of scientists should be drawn to it *because of this success*. It follows that a consensus is epistemically justified only in rare cases in which one theory has all the empirical success (2001, 117-20).

Miller (2013) offers also a hybrid theory of knowledge-based consensus. Miller identifies four types of consensuses, three of which are not knowledge based and one is. The first is a non-cognitive consensus, which aims at promoting non-epistemic aims. The second is a "vertically-lucky" consensus. This is an agreement that happens to be correct, but could have easily been wrong. The third is an epistemically unfortunate consensus, in which parties to the consensus have the bad luck of being systematically or deliberately mislead or biased. When a consensus belongs to none of these three types, it is likely to be knowledge-based.

Miller thus argues that when we can eliminate non-epistemic factors, veritic epistemic luck, and epistemic misfortune as the best explanations of a consensus, knowledge remains its best explanation. He identifies three conditions for knowledge being the best explanation of a consensus: (1) *social calibration*: researchers give the same meaning to the same terms and share the same fundamental background assumptions; (2) *apparent consilience of evidence*: the consensus seems to be built on an array of evidence that is drawn from a variety of techniques and methods; (3) *social diversity*: the consensus is shared by men and women, researchers from the private and public sectors, liberals and conservatives, etc. As mentioned, Miller argues that the third condition is not reducible to the second.

So far we distinguished cognitive, social, and hybrid accounts of knowledge-based or epistemically justified consensus. Following Gasparri (2013), we may also distinguish knowledge-*indicative* from knowledge-*conducive* accounts of consensus. Knowledge-indicative accounts identify signs that indicate that a consensus is credible or epistemically justified. Knowledge-conducive accounts specify conditions which should hold if a community of researchers is to converge on the truth (see Table 1).

Gasparri draws this distinction to scrutinize Tucker's and Miller's knowledge-indicative accounts of consensus. He compares their accounts to Lehrer and Wagner's (1981) knowledge-conducive account of consensus, which consists of a mathematical model of a group of fully-informed rational Bayesian agents who iteratively update their degrees of belief on a given proposition and the degree of credibility they assign to other members regarding this proposition until all agents converge on the same degree of belief, and a rational, knowledge-based consensus is formed. While Tucker and Miller regard Lehrer and Wagner's idealized model as irrelevant to real-world consensuses, Gasparri argues that Lehrer and Wagner's account addresses exactly the missing elements in Tucker's and Miller's accounts; namely, the dynamics of rational aggregation of beliefs and trust-building in an epistemic community.

	Cognitive	**Social**	**Hybrid (Cognitive and Social)**
Indicative	Oreskes (2007)	Tucker (2003, 2014)	Solomon (2001); Miller (2013, 2016)
Conducive	Lehrer and Wagner (1981); Habermas (1984); Stegenga (2016)	Longino (1992, 2002a)	

Table 1: A taxonomy of accounts of knowledge-based or epistemically justified consensus.

Solomon's Social Empiricism is similarly criticized for overlooking the dynamics leading to consensus. Wylie (2006) argues that the same social factors that Solomon identifies as obstructing epistemically effective deliberation also obstruct a justified distribution of decision vectors. Thus, a question arises how such an obstruction is to be avoided. Longino (2008) argues that an epistemic appraisal of consensus should answer who should be trusted in scientific debates. Since Social Empiricism does not analyze how decision vectors affect social inquiry, so Wylie and Longino argue, it has no sufficient resources to address these questions.

3. Aiming for Consensus, Dealing with Dissent

Knowledge-conducive accounts of consensus assume that consensus is an aim of inquiry. We may question this assumption. Longino's Critical Contextual Empiricism strives at rational consensus, yet Longino (2002b) acknowledges that permanent rational dissent due to incompatible theoretical background assumptions between dissenters and the majority is possible. Kitcher (2002a; 2002b) denies this possibility, while Andler (2013) argue that it is both possible and normal.

For Solomon, because both consensus and dissent emerge due to accidental factors, dissent is not a temporary glitch to be overcome, and consensus is not the end of inquiry (2001, 117-20). Miller (2016) warns that through the interaction of scientific and non-scientific factors, a scientific consensus can emerge prematurely, and be unjustly deferred to in policy and legal contexts.

Solomon (2015, Ch. 2-5) is suspicious of consensus building initiatives. While such initiatives help the public legitimation and dissemination of information, they neither expedite the discovery of truths nor warrant their final conclusions. Focusing on the NIH Consensus Development Program, Solomon argues that when the available evidence is insufficient for reaching a conclusion, a consensus statement is premature, and when the evidence is clear, it is redundant. A consensus conference is useful only when the available evidence is sufficient but not yet clear. But the NIH usually misses this window of opportunities. Moreover, the main vehicle of scientific progress is empirical research, rather than social deliberation, hence the epistemic contribution of consensus conferences is minimal.

Against this, Kosolosky and Van Bouwel (2014) argue that the NIH Consensus Development Program and similar initiatives can make genuine epistemic contributions. They distinguish academic consensus, which is about issues internal to science, from interface consensus, which is about the interface of science and policy. Consensus initiatives may give reasoned answers to interface question that do not arise in standard scientific practice. Consensus initiatives may help explicate the reasoning behind their conclusions, which cannot be done by mere aggregation of evidence.

Melo-Martín and Intemann (2013) identify epistemic and political difficulties with demands by the public or decision makers for a unified scientific front as a condition for action. Such demands convey a distorted image of science as a body that speaks unanimously. They may silence legitimate dissent and set too high a bar for science to meet (cf. Miller 2015b; 2016, 29-31).

These worries raise the problem of dealing with dissent. On the one hand, dissent may have epistemic benefits; "controversies are indispensable for the formation, evolution and evaluation of theories, because it is through them that the essential role of criticism in engendering, improving, and controlling the 'wellformedness' and the 'empirical content' of scientific theories is performed" (Dascal 1998, 147). Beatty and Moore (2010, 203) argue that effective critical deliberation requires researchers genuinely committed to a dissent view. But when they exist, the deliberation may not produce a consensus.

On the other hand, dissent may be epistemically detrimental, especially dissent stemming from *manufactured uncertainty* or *doubt mongering*. Affluent bodies opposed to a particular piece of knowledge may inhibit the formation of consensus or create the perception that it does not exist. They may insist on more and ever more critical scrutiny, no matter how strong the evidence is. Oreskes and Conway (2010) argue that the climate skeptical claims have not originated from within the scientific community but rather from politically motivated external actors who, consciously and cynically, have been manufacturing controversy on the subject.

Intemann and Melo-Martín (2014) argue that Longino's Critical Contextual Empiricism faces special difficulties dealing with manufactured uncertainty, because of its prescribed norms of uptake of criticism. Who is to judge whether dissenters fail to follow this norm or whether their concerns are genuine? When may a community stop engaging with them and move on? Defending Critical contextual Empiricism against this charge, Borgerson (2011, 445) argues that if we distinguish the level of certainty required for action from that required for knowledge, interested parties will be less motivated to manufacture uncertainty. In response, Miller (2015b, 118-119) argues that Critical Contextual Empiricism should still be able to determine when closure in an epistemic community is warranted despite incessant criticism.

Unlike Critical Contextual Empiricism, Tucker's and Miller's accounts of knowledge-based consensus arguably do not face difficulties with manufactured dissent. Tucker and

Miller hold that consensus is knowledge-based when knowledge is the best explanation thereof. If knowledge remains the best explanation of a consensus excluding a dissent, and the dissent is best explained by non-epistemic factors, such as financial interests, then we may still legitimately infer that the consensus excluding the dissent is knowledge-based.

Biddle and Leuschner (2015) provide an inductive-risk account for distinguishing epistemically beneficial from epistemically detrimental dissent. Inductive risk is the risk of making a wrong epistemic judgment, such as accepting a false hypothesis and rejecting a true hypothesis (Douglas 2009, Ch. 5). Drawing on Wilholt (2009), who characterizes conventional scientific epistemic standards, e.g., using a critical *p* value of five percent, as reflecting conventional trade-offs between inductive risks, Biddle and Leuschner identify four conditions jointly sufficient for a dissent over a hypothesis H to be considered epistemically detrimental: (1) the non-epistemic consequences of wrongly rejecting H are likely to be severe; (2) the dissenting research that constitutes the objection violates established conventional standards; (3) The dissenting research involves intolerance for producer risks at the expense of public risks; (4) producer risks and public risks fall largely upon different parties.

Beatty and Moore (2010, 209) propose that a scientific community allow dissent and aim at a meta consensus shared by the entire community that the deliberation process on the theory was adequate, and in this respect the theory can stand for the view of the group. The realistic prospects of this proposal remain to be seen.

4. Computer Simulations of Consensus and Dissent Formation

Some scholars develop computational models of a group of epistemic agents who exchange information, and revise their beliefs accordingly.[2] The models inter alia try to identify the salient factors that lead to knowledge-based consensus or permanent dissent.

Zollman presents two such models. One model (Zollman 2010) finds that in a sufficiently dense community, there is an epistemic trade-off between diversity of strong views and fast communication channels. When only one factor is present, the community is likely to converge on the truth, but not when both are. The rationale is that when both strong views and fast communication channels are present, scientists with weaker beliefs are more easily persuaded to abandon their beliefs and adopt the beliefs of scientists with stronger beliefs, such that if the abandoned beliefs are, in fact, true, the scientific community will not discover that. Another model (Zollman 2008) finds a similar trade-off between the speed of communication and its reliability. Holman and Bruner (2015; forthcoming) relax Zollman's assumptions and add factors such as intransigent biases. They demonstrate that convergence on a falsehood or permanent dissent may also occur under some circumstances.

Thagard presents two models that associate consensus with coherence. The first model (Thagard 2000, Ch. 10) is of "cold" (intellectual) cognitive factors. In it, a consensus in a scientific community emerges if the different pieces of information scattered in the community cohere, and the communication between the agents is sufficiently reliable. The second model (Thagard and Kroon 2006) studies the effect of "hot" (emotional) cognitive factors on consensus formation. In it, subjects interact and emotionally affect each other. Consensus on a judgement is reached when it coheres with the subjects' emotional stances toward it. The outcome is path dependent. If Alice and Bob may both emotionally affect Charlie's judgment, it may matter whether Charlie interacted first with Alice or Bob.

Such models do not make traditional philosophical argument redundant. Their external validity and applicability to concrete circumstances hinge on the validity of their

[2] See also Reijula and Kuorikoski in this volume.

background assumptions, which need to be justified by argument. They can thus still benefit from engaging with the work reviewed in the previous sections.

References

Anderson, Elizabeth. 2011. Democracy, Public Policy, and Lay Assessments of Scientific Testimony. *Episteme* 8(2): 144-164.

Andler, Daniel. 2103. Dissensus in Science as a Fact and as a Norm. In *New Challenges to Philosophy of Science*, eds. Hanne Andersen, Dennis Dieks, Wenceslao J. Gonzalez, Thomas Uebel, and Gregory Wheeler, 493-506. Dordrecht: Springer.

Beatty, John and Alfred Moore. 2010. Should we Aim for Consensus? *Episteme* 7(3): 198–214.

Beatty, John. 2006. Masking Disagreement among Experts. *Episteme* 3(1): 52-67.

Beatty, John. 2017. Consensus: Sometimes It Doesn't Add Up. Forthcoming in *Landscapes of Collectivity in the Life Sciences*, edited by Snait Gissis, Ehud Lamm and Ayelet Shavit. Cambridge, MA: MIT Press.

Biddle, Justin B. and Anna Leuschner. 2015. Climate Skepticism and the Manufacture of Doubt: Can Dissent in Science Be Epistemically Detrimental? *European Journal for Philosophy of Science* 5(3): 261-278.

Borgerson, Kirstin. 2011. Amending and Defending Critical Contextual Empiricism. *European Journal for Philosophy of Science* 1(3): 435-449.

Bruggeman, Jeroen, V. A. Traag, and Justus Uitermark. 2012. Detecting Communities through Network Data. *American Sociological Review* 77: 1050–1063.

Bruner, Justin and Bennett Holman. Forthcoming. Complicating Consensus. In *Expert Disagreement and Measurement*, ed. Luciana Garbayo. Dordrecht: Springer.

Collins, Harry M., 2010. *Elective Modernism*. Cardiff: Cardiff School of Social Sciences. www.cardiff.ac.uk/socsi/contactsandpeople/harrycollins/expertise-project/elective%20modernism%204.doc

Dascal, Marcelo. 1998. The Study of Controversies and the Theory and History of Science. *Science in Context* 11(2): 147-154.

de Melo-Martín, Inmaculada and Kristen Intemann. 2013. Scientific Dissent and Public Policy: Is Targeting Dissent a Reasonable Way to Protect Sound Policy Decisions? *EMBO Reports* 14(3): 231-235.

Douglas, Heather. 2009. *Science, Policy, and the Value-Free Ideal*. Pittsburgh, PA: University of Pittsburgh Press.

Einsiedel, Edna F., Erling Jelsøe, and Thomas Breck. 2001. Publics at the Technology Table: The Consensus Conference in Denmark, Canada, and Australia. *Public Understanding of Science* 10(1): 83-98.

Fuller, Steve. 2002. *Social Epistemology*, 2nd ed. Bloomington, IN: Indiana University Press.

Goldman, Alvin I. 2002. Knowledge and Social Norms. *Science* 296 (June 21): 2148-2149.

Habermas, Jürgen. 1984. *The Theory of Communicative Action*, Vol. 1 & 2. Boston: Beacon Press.

Habermas, Jürgen. 2006. Political Communication in Media Society: Does Democracy Still Enjoy an Epistemic Dimension? The Impact of Normative Theory on Empirical Research. *Communication Theory* 16(4): 411-426

Hallerberg, Arthur E. 1977. Indiana's Squared Circle. *Mathematics Magazine* 50(3): 136-140.

Holman, Bennett and Justin P. Bruner. 2015. The Problem of Intransigently Biased Agents. *Philosophy of Science* 82(5): 956-968.

Horst, Maja. 2008. In Search of Dialogue: Staging Science Communication in Consensus Conferences. In *Communicating Science in Social Contexts*, eds. Donghong Cheng, Michel Claessens, Toss Gascoigne, Jenni Metcalfe, Bernard Schiele, and Shunke Shi, 259-274. Dordrecht: Springer.

Intemann, Kristen and Inmaculada de Melo-Martín. 2014. Are There Limits to Scientists' Obligations to Seek and Engage Dissenters? *Synthese* 91(12): 2751-2765.

Jones, Ward E. 2002. Dissident Versus Loyalist: Which Scientists Should We Trust? *The Journal of Value Inquiry* 36(4): 511-520.

Kitcher, Philip. 2002a. The Third Way: Reflections on Helen Longino's *The Fate of Knowledge*. *Philosophy of Science* 69(4): 549-559.

Kitcher, Philip. 2002b. Reply to Helen Longino. *Philosophy of Science* 69(4): 569-572.

Kosolosky, Laszlo and Jeroen Van Bouwel. 2104. Explicating Ways of Consensus-Making in Science and Society: Distinguishing the Academic, the Interface and the Meta-Consensus. In *Experts and Consensus in Social Science*, eds. Carlo Martini and Marcel Boumans, 71-92. Dordrecht: Springer.

Kuhn, Thomas S. 1961/1970. *The Structure of Scientific Revolutions*, 2nd ed. Chicago: The University of Chicago Press.

Laudan, Larry. 1984. *Science and Values: The Aims of Science and Their Role in Scientific Debate*. Berkeley: University of California Press.

Lehrer, Keith & Carl Wagner. 1981. *Rational Consensus in Science and Society: A Philosophical and Mathematical Study*. Dordrecht: Reidel.

Longino, Helen. 2002a. *The Fate of Knowledge*. Princeton: Princeton University Press.

Longino, Helen. 2002b. Reply to Philip Kitcher. *Philosophy of Science* 69(4): 573-577.

Longino, Helen. 2008. Norms and Naturalism: Comments on Miriam Solomon's Social Empiricism. *Perspectives on Science* 16(3): 241-245.

Medlock, J., R. Downey, and E. Einsiedel. 2007. Governing Controversial Technologies: Consensus Conferences as a Communications Tool. In *The Media, the Public and Agricultural Biotechnology*, ed. Dominique Brossard, James Shanahan, and T. Clint Nesbitt, 308-326.Wallingford, Oxfordshire: CABI.

Miller, Boaz. 2013. When Is Consensus Knowledge Based? Distinguishing Shared Knowledge from Mere Agreement. *Synthese* 190(7): 1293-1316.

Miller, Boaz. 2015a. Why (Some) Knowledge is the Property of a Community and Possibly None of Its Members. T*he Philosophical Quarterly* 65(260): 417-441.

Miller, Boaz. 2015b. 'Trust Me – I'm a Public Intellectual': Margaret Atwood's and David Suzuki's Social Epistemologies of Climate Science. In *Speaking Power to Truth: Digital Discourse and the Public Intellectual*, edited by Michael Keren & Richard Hawkins, 113-128. Athabasca, AB: Athabasca University Press.

Miller, Boaz. 2016. Scientific Consensus and Expert Testimony in Courts: Lessons from the Bendectin Litigation. *Foundations of Science* 21(1): 15-33.

Odenbaugh, Jay. 2012. Climate, Consensus, and Contrarians. In *The Environment: Philosophy, Science, and Ethics*, eds. William P. Kabasenche, Michael O'Rourke, Matthew H. Slater, 137-150. Cambridge, MA: MIT Press.

Oreskes, Naomi and Erik M. Conway. 2010. *Merchants of Doubt: How a Handful of Scientists Obscured the Truth on Issues from Tobacco Smoke to Global Warming.* New York: Bloomsbury.

Oreskes, Naomi. 2007. The Scientific Consensus on Climate Change: How Do We Know We're Not Wrong? In *Climate Change: What It Means for Us, Our Children, and Our Grandchildren*, eds. Joseph F.C. DiMento and Pamela Doughman, 65–99. Cambridge, MA: MIT Press.

Rescher, Nicholas. 1993. *Pluralism: Against the Demand for Consensus*. New York: OUP.

Shwed, Uri and Peter S. Bearman. 2010. The Temporal Structure of Scientific Consensus Formation. *American Sociological Review* 75(6): 817-840.

Shwed, Uri and Peter S. Bearman. 2012. Symmetry is Beautiful. *American Sociological Review* 77(6): 1064-1069.

Solomon, Miriam and Alan Richardson. 2005. A Critical Context for Longino's Critical Contextual Empiricism. *Studies in History and Philosophy of Science* 36: 211-222.

Solomon, Miriam. 2001. *Social Empiricism*. Cambridge, MA: MIT Press.

Solomon, Miriam. 2015. *Making Medical Knowledge.* Oxford: OUP.

Stegenga, Jacob. 2016. Three Criteria for Consensus Conferences. *Foundations of Science* 21(1): 35–49.

Sunstein, Cass R. 2006. Deliberating Groups versus Prediction Markets (or Hayek's Challenge to Habermas). *Episteme* 3(3): 192-213.

Surowiecki, James. 2004. T*he Wisdom of Crowds: Why the Many are Smarter than the Few and How Collective Wisdom Shapes Business, Economies, Societies, and Nations*. New York: Doubleday.

Thagard, Paul and Fred W. Kroon. 2006. Emotional Consensus in Group Decision Making. *Mind & Society* 5(1): 85-104.

Thagard, Paul. 2000. *Coherence in Thought and Action*. Cambridge, MA: MIT Press.

Thicke, Micahel. 2017. Market Epistemology. *Synthes*e, in press. link.springer.com/epdf/10.1007/s11229-017-1464-2

Tollefsen, Deborah P. 2006. Group Deliberation, Social Cohesion, and Scientific Teamwork: Is There Room for Dissent? *Episteme* 3(1-2): 37-51.

Tucker, Aviezer. 2003. The Epistemic Significance of Consensus. *Inquiry* 46(4): 501-21.

Tucker, Aviezer. 2014. Epistemology as a Social Science: Applying the Neyman-Rubin Model to Explain Expert Beliefs. In *Experts and Consensus in Social Science*, eds. Carlo Martini, & Marcel Boumans, 155-170. Dordrecht: Springer.

Vermeule, Adrian. 2009. *Law and the Limits of Reason*. New York: OUP.

Wikipedia. 2017. Wikipedia:Consensus. en.wikipedia.org/wiki/Wikipedia:Consensus

Wilholt, Torsten. 2009. Bias and Values in Scientific Research. *Studies in History and Philosophy of Science* 40(1): 92-101.

Wylie, Alison. 2006. Socially Naturalized Norms of Epistemic Rationality: Aggregation and Deliberation. *The Southern Journal of Philosophy* 44: 43-48.

Zollman, Kevin J. S. 2007. The Communication Structure of Epistemic Communities. *Philosophy of Science* 74(5): 574-587.

Zollman, Kevin J. S. 2010. The Epistemic Benefit of Transient Diversity. *Erkenntnis* 72: 17-35.

Further Reading

Simon (2010) identifies three types of communal epistemic closure in Internet platforms, and discusses their respective virtues and vices. Andler (2012) criticizes the notion of the Wisdom of Crowds as an epistemic panacea. Van Bouwel (2009) draws a comparison between political models of pluralism and social-epistemic models of consensus and dissent. Delborne (2008) and Martin (2008) offer empirically-based taxonomies of dissent and dissenting practices.

Andler, Daniel. 2012. What Has Collective Wisdom to Do with Wisdom? In *Collective Wisdom: Principles and Mechanisms*, eds. Hélène Landemore and Jon Elster, 72-94. Cambridge: Cambridge University Press.

Delborne, Jason A. 2008. Transgenes and Transgressions: Scientific Dissent as Heterogeneous Practice. *Social Studies of Science* 38(4): 509-541.

Martin, Brian. 2008. Varieties of dissent. In *Dissent and the Failure of Leadership*, ed. Stephen P. Banks, 22-36. Cheltenham, UK: Edward Elgar, 2008.

Simon, Judith. 2010. A Socio-Epistemological Framework for Scientific Publishing. *Social Epistemology* 24(3): 201-218.

Van Bouwel, Jeroen. 2009. The Problem With(out) Consensus: The Scientific Consensus, Deliberative Democracy and Agnostic Pluralism. In *The Social Sciences and Democracy*, ed. Jeroen Van Bouwel, 121-142. New York: Palgrave Macmillan.