

Peer disagreement under multiple epistemic systems

Rogier De Langhe

Received: 12 May 2012 / Accepted: 18 June 2012 / Published online: 31 July 2012
© The Author(s) 2012. This article is published with open access at Springerlink.com

Abstract In a situation of peer disagreement, peers are usually assumed to share the same evidence. However they might not share the same evidence for the epistemic system used to process the evidence. This synchronic complication of the peer disagreement debate suggested by Goldman (In Feldman R, Warfield T (eds) (2010) *Disagreement*. Oxford University Press, Oxford, pp 187–215) is elaborated diachronically by use of a simulation. The Hegselmann–Krause model is extended to multiple epistemic systems and used to investigate the role of consensus and difference splitting in peer disagreement. I find that the very possibility of multiple epistemic systems downgrades the epistemic value of consensus and makes difference splitting a suboptimal strategy.

Keywords Peer disagreement · Epistemic systems · Alvin Goldman · Hegselmann–Krause model · Netlogo

1 Introduction

In the early days of economics, Adolphe Quetelet aspired to treat the social world like geometry. He assumed that there are general principles to be found in the social world that are concealed randomly distributed errors. Quetelet looked for general principles with a (potentially very large) number of allowances for those factors that only exert a minor influence. His aim was to statistically characterize the regularities of the social world like astronomers statistically reconstruct orbits from scattered observations. John Stuart Mill, on the other hand, believed that even small factors could have a large

R. De Langhe (✉)
Tilburg Center for Logic and Philosophy of Science (TiLPS),
Tilburg University, Tilburg, The Netherlands
e-mail: rogiерdelanghe@gmail.com

influence on the entire system. Hence he was suspicious of allowances and instead looked for the exceptionless tendencies behind the phenomena.¹ This is an example of peer disagreement caused by different epistemic systems. Now imagine Quetelet and Mill were given the exact same dataset and were asked to estimate the probability of a double-dip recession in the US. Would they reach a consensus in a finite period of time? Should they try to do so by splitting the difference? And what would the epistemic significance of that consensus be?

Peer disagreement is an important topic in social epistemology. A common answer to the question of how to deal with a disagreeing peer is to 'split the difference' (Christensen 2007; Elga 2007; Feldman and Warfield 2010), viz. adjust our belief to make it a mixture of the other's belief and one's own. On such a view, sticking to one's opinion is irrational. Among others such as Kelly (2010), Douven (2009, 2010) has challenged this view. Douven claims that longstanding disagreements, viz. a refusal to split the difference, is endemic in science. Rather than concluding that scientists are irrational, Douven uses computer simulations to demonstrate that it can indeed be rational to stick to one's opinion. He shows that a community of scientists sticking to their opinion will come near the truth faster, although less accurately so. In situations where speed is more important than accuracy, for example when time and resources are scarce, sticking to one's opinion is therefore rational. Yet, Douven does not make true on his promise entirely. Sticking to one's opinion might well be rational when time and resources are scarce, but these are exactly factors that abound in longstanding disagreements. Although he adduces a good reason for rational disagreement, he therefore fails to explain longstanding disagreement. A more promising candidate for the explanation of longstanding disagreement in science is Alvin Goldman's notion of "epistemic system". The literature on peer disagreement, including Douven (2009, 2010), usually does not consider peer disagreement under multiple epistemic systems. (Goldman 2010, p. 208) This can readily be explained by the fact that most scholars are objectivists. Objectivism holds that there is objective rightness in epistemic standards, principles and norms. It is often understood as denying the possibility of multiple correct standards or norms for belief formation. By contrast, Goldman (2010) proposes a new form of relativism, *objectivity-based relativism*, which saves objectivism but allows for two people with the same evidence for a proposition to reasonably disagree about that proposition. This is possible because peers sharing the same evidence for a belief might not share the same evidence for the epistemic system used to form this belief. Hence rational scientists can disagree indefinitely.

However, Goldman's framework is not suited to evaluate peer disagreement, as Goldman himself admits.² Peer disagreement is a diachronic problem while Goldman offers only a synchronic treatment. To remedy this I will make use of a simulation. Simulations are commonly used for diachronic analysis. Scholars in fields such as physics, computer science, mathematics and philosophy have developed simulations of opinion dynamics, e.g. Lehrer and Wagner (1981), Hegselmann and Krause (2002,

¹ "The body does not only move in that manner unless counteracted, it tends to move in that manner even when counteracted." (Mill 1872, p. 444)

² "The framework I employ here lacks sufficient tools for a detailed analysis of [this] problem." (Goldman 2010, p. 210)

2005, 2006). These descriptions usually focus on the conditions under which opinions converge or polarize. Although an agent-based model will lie at the heart of my paper, my aim is normative rather than descriptive: to investigate how an epistemic agent *should* react when disagreeing with an epistemic peer.

Models of opinion dynamics have already proven to be ideal testing grounds for exploring peer disagreement. On the Lehrer–Wagner model, agents in a group update their beliefs based on the weighted averages of the beliefs of all others in the group. The Hegselmann–Krause model assumes that agents only update their beliefs when other agents are within a certain (symmetrical or asymmetrical) “confidence interval”. This reflects the idea that agents will not take into account views that radically diverge from their own. In both the Lehrer–Wagner model and the Hegselmann–Krause model agents must split the difference with each other in order to arrive at consensus. Hence not splitting the difference and agreeing to disagree is considered to be irrational. Douven (2010) uses the model of opinion dynamics developed by Rainer Hegselmann and Ulrich Krause because finds this model particularly suited for the purpose of simulating peer disagreement. In order to develop a diachronic perspective on multiple epistemic systems I will therefore use the Hegselmann–Krause model (HK-model), as well.

As such this paper is relevant to the literature in four different ways. (1) It will evaluate peer disagreement under multiple epistemic systems. (2) This will be done by extending the Hegselmann and Krause (2002, 2005, 2006) to multiple epistemic systems, thus (3) providing a diachronic treatment of Goldman (2010) and (4) meet Douven (2010)’s demand for an account of rational longstanding disagreement.

2 Epistemic systems

According to Goldman (2010), *epistemic systems* are “sets of norms, standards, or principles for forming beliefs and other doxastic states.” (Goldman 2010, p. 187) A crucial difference between relativists and objectivists is that the latter believe there is a uniquely correct epistemic system, while the former do not. Objectivism, the view that there is a uniquely correct epistemic system, seems to imply the *Uniqueness Thesis*: there are different attitudes one can have toward a particular proposition (namely believing, disbelieving or suspending judgment), but given a body of evidence there is only one correct attitude that is the rationally justified one. As such, the Uniqueness Thesis renders impossible any rational disagreement. According to Feldman, epistemic peers who have shared their evidence cannot have rational disagreement.

Just as Douven (2010), Goldman is unsatisfied with frameworks failing to account for rational disagreement. In his words, he wants to save “the intuition that differences in intellectual procedure found in diverse cultures, communities, and historical periods do not reflect wholesale irrationality or epistemic depravity.” (Goldman 2010, p. 202) To argue against the necessary connection between the impossibility of rational disagreement and objectivism he develops a new form of relativism which he calls “*objectivity-based relativism*”. On this view, people that share the same evidence for a proposition might still disagree because they have different evidence for their respective epistemic systems. Because they occupy different evidential positions, they

can be *objectively justified* in believing different epistemic systems and hence have different beliefs.

Goldman offers an objectivist account of how rational persons can rationally disagree (synchronic), but on his account the question of how to react to such disagreement through time (diachronic) is left unexplained. (Christensen 2007, p. 188) describes this diachronic problem as: “How should I react when I discover that my friend and I have very different beliefs on some topic? [...] Should my discovery of her differing degree of belief in P lead me to revise my own confidence in P?” In particular, Goldman wonders “Must their beliefs converge? Must the mode of convergence involve splitting the difference?” (Goldman 2010, p. 210) As Goldman notes himself, his framework lacks sufficient tools for a detailed analysis of these questions. In the next section I will extend to HK-model to include multiple epistemic systems and as such allow for a diachronic analysis of peer disagreement under multiple epistemic systems. In Sect. 3 I will demonstrate how the model can provide an answer to these two questions. Finally Sect. 4 discusses to what extent the results in this paper hinge on one’s definition of “peer”.

3 A model of peer disagreement under multiple epistemic systems

Peer disagreement is a diachronic problem because it involves the change of beliefs through time. A common way to investigate the dynamics of a process is the use of simulations. Simulations have already been used to investigate peer disagreement, for example by Douven who notes that the Hegselmann–Krause model “*is particularly suited for this purpose*” (Douven 2010, p. 149). Yet the Hegselmann–Krause model presupposes that all agents adopt the same epistemic system. As a consequence, to investigate peer disagreement under multiple epistemic systems I will need to extend the model to multiple epistemic systems.

The Hegselmann–Krause model is a model of the opinion dynamics of a group of agents that are attempting to find the true value of a certain parameter (for example the likelihood of a double-dip recession, as in the example above). The value of the parameter lies within the interval $]0,1[$. Initially agents’ beliefs are distributed uniformly across the range. At each turn agents update their beliefs based on their own research, which brings them closer to the truth, and the research of others (testimony). Characteristic for the HK-model is that an agent will only update his belief with the beliefs of peers that do not differ from its belief too substantially, viz. when the belief is within the agent’s (symmetrical or asymmetrical) “confidence interval” ε . In an attempt to stay as close as possible to the original HK-model and Douven’s use of it, I will first replicate Douven’s description of the HK-model exactly and consequently demonstrate that my model is indeed able to reproduce Douven’s results. Subsequently I will extend the model beyond existing scholarship toward a model of opinion dynamics under multiple epistemic systems.

Let $\tau \in]0,1[$ be the true value of the parameter, $x_i(u)$ the opinion of agent x_i after the u -th update, and $\alpha \in]0,1[$ the weighting factor used. Further define $X_i(u) := \{j : |x_i(u) - x_j(u)| \leq \varepsilon\}$, and let $|X_i(u)|$ be the cardinality of $X_i(u)$. Then the opinion of agent x_i after the $(u + 1)$ -st turn is represented by the equation:




Fig. 1 A typical run of the HK-model

$$x_i(u + 1) = \alpha \frac{1}{|X_i(u)|} \sum_{j \in X_i(u)} X_j(u) + (1 - \alpha) \tau \tag{1}$$

The normative question of how to react to peer disagreement can then be translated into the model by manipulating the value of ε . With $\varepsilon = 0$ the agent sticks to its opinion. At the next turn, its belief is simply the average of its previous belief and the true value of τ . A higher value of ε means that the agent is willing to split the difference (although only with those agents within its confidence interval). I have reproduced this model in Netlogo.³ A typical run of the model with $\tau = 0.75$, $\alpha = 0.5$ and $\varepsilon = 0.1$ results in a quick convergence on the true value of τ as shown in Fig. 1.

Douven (2010) extends this model by adding some noise to the data. At each turn the agents do not receive information about the exact value of τ through their own research, but about a value that fluctuates around it.

$$x_i(u + 1) = \alpha \frac{1}{|X_i(u)|} \sum_{j \in X_i(u)} X_j(u) + (1 - \alpha) (\tau + \text{rnd}(\zeta)) \tag{2}$$

where $\text{rnd}(\zeta)$ is a uniformly distributed random real number chosen by each agent at each turn from the interval $[-\zeta, +\zeta]$, with $\zeta \in [0,1]$. In such cases the model demonstrates that agents sticking to their opinion ($\tau = 0.75$, $\alpha = 0.5$, $\varepsilon = 0$ and $\zeta = 0.2$) will actually converge faster to values near the truth than agents that split the difference

³ This simulation was programmed in Netlogo v4.1.3. Code available on request.

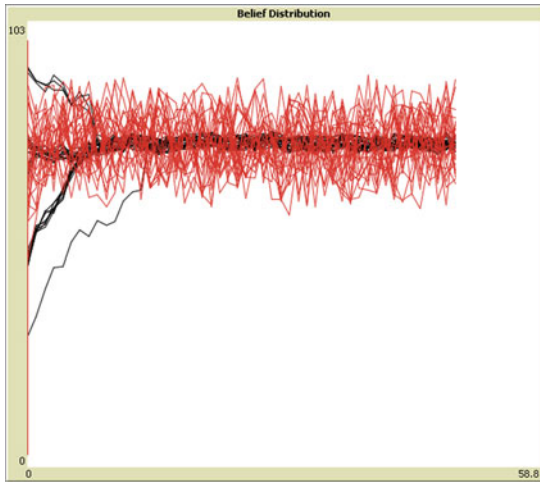


Fig. 2 Noisy data with and without difference splitting




Fig. 4. Noisy data, no difference-splitting.




Fig. 5. Noisy data, repeated difference-splitting.

($\tau = 0.75, \alpha = 0.5, \varepsilon = 0.1$ and $\zeta = 0.2$). My reproduction of both cases respectively is shown on the left side of Fig. 2, with red for difference splitting and black for sticking to one’s opinion. On the right are the results of Douven (2010, p. 151) for the same settings after 50 ticks.

Now that both the HK-model and Douven’s extension have been successfully reproduced, I will extend this model to include multiple epistemic systems. I therefore introduce the complication that τ is relative to the agent’s epistemic system. As a consequence, as stated by Goldman’s synchronic analysis of multiple epistemic systems, agents can now rationally disagree about the true value of τ . Moreover, to allow agents to discriminate between the beliefs of peers within and outside their own epistemic system, a parameter σ is introduced, which is a fraction of ε . As in the HK-model, the agents share the same confidence interval ε , but in this extension they apply that confidence interval only to peers within their own epistemic system. As such, ε denotes the extent to which agents in the model split the difference *within* their system. The confidence interval they apply to peers outside their epistemic system is a fraction of that confidence interval. If $\sigma = 1$ agents update their beliefs on all peers within their confidence interval, regardless of those peers’ epistemic system. If for example $\sigma = 0.5$ then the confidence interval with regard to peers from other epistemic systems is only half of the confidence interval used for peers from within one’s epistemic system. σ then denotes the extent to which agents split the difference *between* epistemic systems. The model takes the following general form

$$x_i(u + 1) = \alpha \frac{1}{|X_i(u)|} \sum_{j \in X_i(u)} X_j(u) + (1 - \alpha)\tau x_i \tag{3}$$

with $X_i(u) := \{j : |x_i(u) - x_j(u)| \leq \sigma \varepsilon\}$.

4 Belief convergence and difference splitting

The reason for extending the HK-model with multiple epistemic systems was to enable a diachronic analysis of peer disagreement under multiple epistemic systems. Goldman's first question was whether it was still desirable for degrees of belief to converge. His second question concerned the desirability of difference splitting.

The model was run with a population of 100 agents divided evenly across 4 different epistemic systems with $\tau_1 = 0.14$, $\tau_2 = 0.37$, $\tau_3 = 0.62$ and $\tau_4 = 0.86$. In the plots the adopters of each epistemic system receive a different color (red, blue, green, black respectively). Figure 3 plots for each epistemic system how far its agents are on average away from its respective truth, accompanied by a plot of the parameter value chosen at each tick by each agent. Three typical runs of the model are shown with each time $\alpha = 0.5$ and various settings of σ and ε .

Strikingly, the members of each system still *converge*, but not always on the right value for τ . As a consequence, whereas models such as the HK-model suggest that consensus has epistemic significance because agents convert on the true value of τ , it is here demonstrated that this property of the model is not robust against an extension of the HK-model to multiple epistemic systems.

To further investigate this finding the model was run 10 times for each combination of σ and ε in an interval $[0,1]$ with an 0,1 increment. For each combination it was calculated how far the agents were on average removed from the true value of their respective epistemic system (Total Average Distance) after 1,000 ticks. The average of this value over 10 runs is represented on the Y-axis of Fig. 4.

Once agents start splitting the difference with agents from other epistemic systems ($\sigma > 0$) a cost in epistemic utility is incurred. This remains invisible (agents converge on the true value of τ_X ; without incurring a cost) if the analysis is restricted to cases where there are no other epistemic systems ($\sigma = 0$). Interestingly, σ and ε cannot be raised simultaneously without a cost in epistemic utility. The simulations reveal a *tradeoff* between splitting the difference within one's system and splitting the difference between epistemic systems. In order to make splitting the difference within one's system epistemically justified one must decline sharing the difference with adopters of other epistemic systems. As a consequence, the imperative to always split the difference is not robust under an extension to multiple epistemic systems.

5 The definition of "peer"

The most important objection against the notion of peer disagreement under multiple epistemic systems is probably the objection that if agents have a different epistemic system, they do not share the same evidence and hence they are not peers in the first place. In this section I offer a phased response to this objection:

- (1) In the literature on peer disagreement, the evidence shared by peers is evidence concerning the belief under consideration, evidence concerning one's own competence and evidence concerning the competence of peers. The evidence used to justify one's choice of epistemic system is usually ignored. (Goldman 2010,

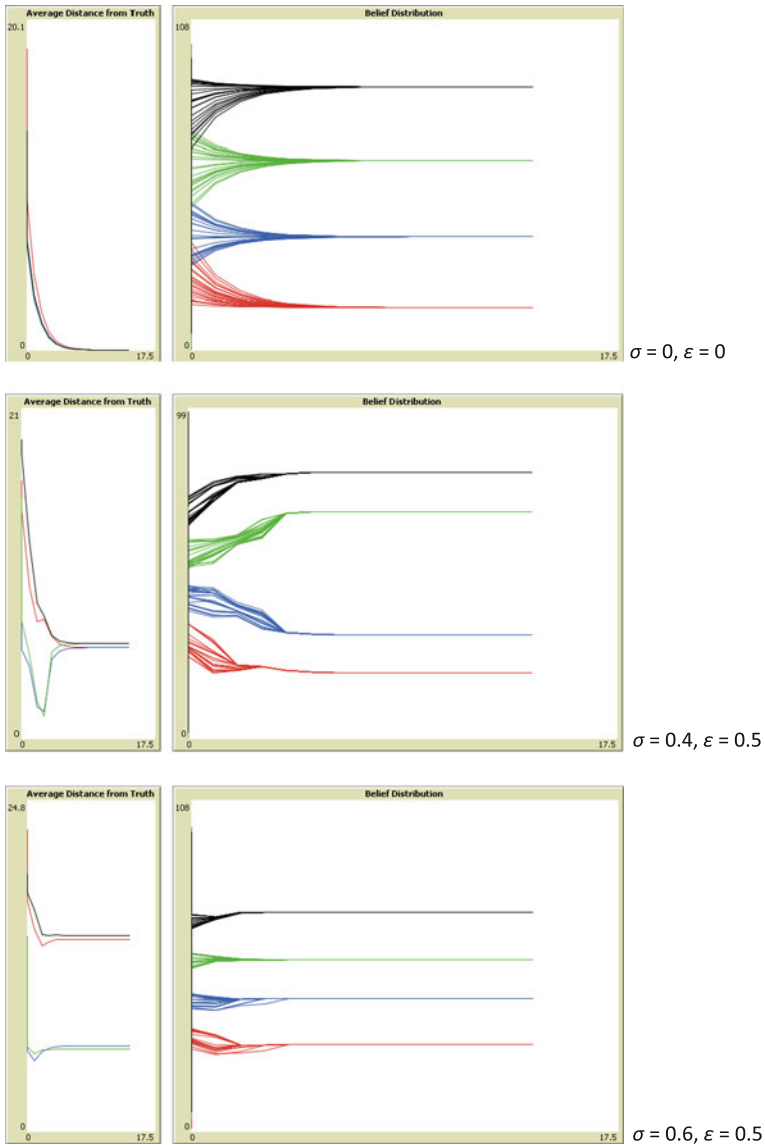


Fig. 3 Typical runs of the HK-model under multiple epistemic systems

- p. 208) As such a minimal result of this paper would be that the notion of “peer” is narrowed down to peers sharing the same epistemic system.
- (2) Narrowing down the definition of “peer” to save the epistemic desirability of difference splitting in the face of peer disagreement is an artificial solution because it requires that agents can discriminate between disagreement within and disagreement between epistemic systems. Whether or not agents share the three usual forms of evidence peers are considered to share can relatively easily be determined because this judgment is relative to the topic under consideration.

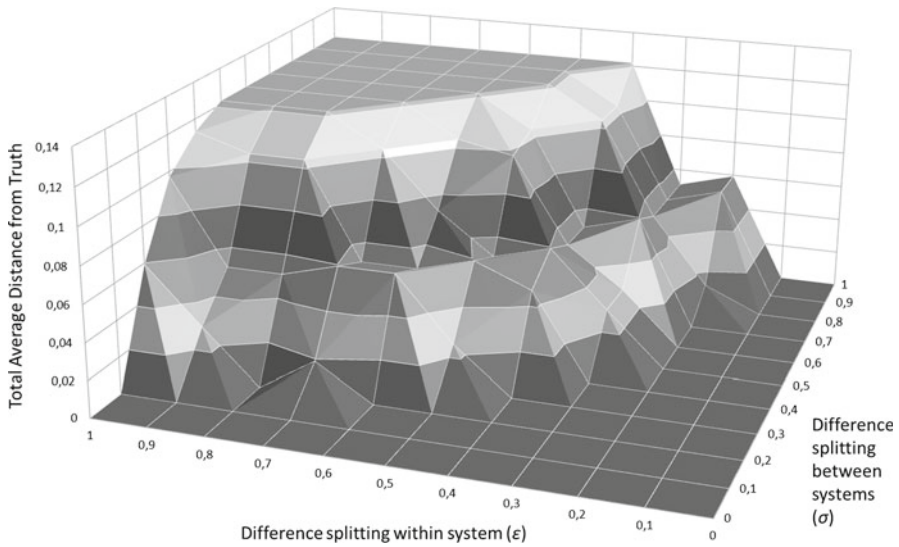


Fig. 4 Behavior space of the extended HK-model

Agents' evidence for preferring one epistemic system over another is however dependent on the entire epistemic history of that agent. As a consequence, the investment needed to establish whether or not an agent shares the same epistemic system is of a different magnitude. It is therefore reasonable to argue that the cost of establishing whether or not another agent is a peer will often outweigh the epistemic benefits of difference splitting.

- (3) Even if one narrows the definition of peers to peers sharing the same epistemic system and one is able to discriminate perfectly between peers that apply the same epistemic system, the problem is not solved but simply displaced. The possibility of multiple epistemic systems creates a new problem, namely the problem of "system disagreement": *how to react to peers with different epistemic systems?* If the adherents of one epistemic system split the difference only with other such adherents, they will indeed converge upon their own truth. But wouldn't it be rational for them to at least take note of the fact that others disagree with them, calling into question their own epistemic system? In other words, the possibility of multiple epistemic systems creates a new dilemma between sticking to one's opinion and splitting the difference, but this time it is situated on a higher level, namely the level of systems. Paradoxically, then, the model in this paper has demonstrated that in order to create the circumstances in which it is always rational to split the difference requires that one sticks to one's opinion with respect to one's own epistemic system. The possibility of multiple epistemic systems thus threatens to render the imperative to always split the difference incoherent. If one decides to narrow the definition of "peer" to agents sharing the same epistemic system, then one is not splitting the difference with regard to one's epistemic system. But if one decides peers can have different epistemic systems (as shown in the simulation) communities of peers will not converge on the truth.

6 Conclusion

In the debate on peer disagreement, difference splitting and consensus are both widely considered to be epistemically desirable. This is supported by models of opinion dynamics such as the HK-model. Both Igor Douven and Alvin Goldman object that these models cannot account for longstanding disagreement. However, Douven only explains short-term disagreement and Goldman only gives a synchronic treatment of epistemic systems unable to evaluate normative claims on peer disagreement. This paper has provided an extension of the HK-model, extending Douven's critique of difference splitting to long-term disagreement and making Goldman's treatment of epistemic systems diachronic. As anticipated, the resulting model was able to analyse peer disagreement under multiple epistemic systems. The model demonstrated that the epistemic desirability of difference splitting and consensus is not robust against an extension to multiple epistemic systems. Whereas the HK-model typically converges to the truth, it does not when multiple epistemic systems are present in the scientific community. This is an important theoretical result because the very *possibility* of multiple epistemic systems already undermines the trust one should put in difference splitting and consensus because of the transaction costs associated with determining the presence of epistemic systems (see Sect. 5). Moreover the possibility of multiple epistemic systems complicates the peer disagreement debate by raising the problem of how to respond to system disagreement.

Open Access This article is distributed under the terms of the Creative Commons Attribution License which permits any use, distribution, and reproduction in any medium, provided the original author(s) and the source are credited.

References

- Christensen, D. (2007). Epistemology of disagreement: The good news. *Philosophical Review*, 116, 187–217.
- Douven, I. (2009). Uniqueness revisited. *American Philosophical Quarterly*, 46, 347–361.
- Douven, I. (2010). *Studies in History and Philosophy of Science*, 41, 148–157.
- Elga, A. (2007). Reflection and disagreement. *Noûs*, 41, 478–502.
- Feldman, R., & Warfield, T. (Eds.). (2010). *Disagreement*. Oxford: Oxford University Press.
- Goldman, A. (2010). Epistemic relativism and reasonable disagreement. In R. Feldman & T. Warfield (Eds.), *Disagreement* (pp. 187–215). Oxford: Oxford University Press
- Hegselmann, R., & Krause, U. (2002). Opinion dynamics and bounded confidence: Models, analysis, and simulations. *Journal of Artificial Societies and Social Simulation*, 5(3). <http://jasss.soc.surrey.ac.uk/5/3/2.html>. Accessed 15 February 2011.
- Hegselmann, R., & Krause, U. (2005). Opinion dynamics driven by various ways of averaging. *Computational Economics*, 25, 381–405.
- Hegselmann, R., & Krause, U. (2006). Truth and cognitive division of labor: First steps towards a computer aided social epistemology. *Journal of Artificial Societies and Social Simulation*, 9(3). <http://jasss.soc.surrey.ac.uk/9/3/10.html>. Accessed 15 February 2011.
- Kelly, T. (2010). Peer disagreement and higher order evidence. In R. Feldman & T. Warfield (Eds.), *Disagreement* (pp. 111–174). Oxford: Oxford University Press.
- Lehrer, K., & Wagner, C. (1981). *Rational consensus in science and society*. Dordrecht: Reidel.
- Mill, J. S. (1872). *A system of logic* (8th ed.). London: Longmans Green, Reader, and Dyer.