Incorporating Representation When Modeling Cultural Dynamics: Analysis of the Bounded Influence Conjecture

Shane T. Mueller (smueller@ara.com), Ben Simpkins (bsimpkins@ara.com), and
Louise Rasmussen (lrasmussen@ara.com)
Klein Associates Division, ARA Inc
Fairborn, OH, 45434, USA

Abstract

A number of social modeling and simulation approaches have attempted to understand the exchange and convergence of ideas based on simple representations and rules for information exchange. However, this focus on finding emergent properties of simple agents and rules has led this research area to avoid handling many genuinely cognitive phenomenona, and ignore many important questions regarding the transmission of cultural ideas. We propose that one deficiency is in the treatment of knowledge in very simple structures, in contrast to the richness of true cultural knowledge. This simplification has lead to the bounded influence conjecture, which we show is not necessary to produce distinct groups of disagreement within an interactive group. Instead, we hypothesize that aspects of the space of valid knowledge states can create a similar outcome. We provide a demonstration of this hypothesis through simulations in a multi-agent interactive model.

Cultural as Shared Knowledge

One common perspective on culture is that it consists of shared beliefs, attitudes, mental models, and customs developed by a group of individuals who interact frequently. This view (cf. Atran, Medin, & Ross, 2005) is somewhat at odds with perspectives of culture that focus on demographic or geographic variables, because it asserts that culture is primarily a phenomenon of cognition, rather than behavior.

A number of methods have been developed to measure this notion of culture-as-shared-knowledge. Although initial perspectives focused on culture as consensus (e.g., Cultural Consensus Theory; Romney, Weller, & Batchelder, 1986; Romney et al., 1996; Romney, 1999), more recent developments have allowed culture to be characterized by a variety of distinct opinion groups (e.g., Cultural Mixture Modeling; Mueller & Veinott, 2008). These approaches focus on formal statistical methods to represent culture as shared knowledge, and enable detailed cultural models to be identified for interview and survey data.

Simulating Cultural Knowledge and Consensus

The culture-as-shared-knowledge perspective is also present among researchers working with large-scale simulations. For example, Axelrod (1986) famously define a framework for simulating the exchange of ideas in a multi-agent system that enabled formation of local consensus but global diversity. Here, a local group of agreement could be considered a single culture, but the simple act of idea exchange did not create a global monoculture. One important driver of this was the assumption of homophily or bounded influence: an agent is only willing to listen and be influenced by viewpoints that are sufficiently similar to its own. In those models, similarity is both a function of geographic closeness or closeness in a social network (agents only talk to adjacent or nearby agents) and conceptual similarity (agents only talk to agents who are identical on a minimal proportion of features). Thus, after distinct groups emerge, it becomes impossible for members to jump from one group to another, because they simply refuse to be influenced by the extreme beliefs of the other groups.

Opinion Dynamics and Representation

In recent years, a community of practice has emerged in the field of opinion dynamics (see Lorenz, 2007, for a survey of the field), which has explored many of the same ideas as Axelrod, perhaps without arguing the models are cultural. Nevertheless, if one views culture as a consensus of ideas, researchers in this community models the process by which the consensus may develop. In addition, if one takes the view that culture can be a diversity of ideas, this community also models the conditions under which consensus fails to emerge.

Although earlier work by Axelrod used somewhat complex knowledge representations (a multi-dimensional vector containing numbered tokens), the opinion dynamics community has focused primarily on simpler representations that take their inspiration from physical models (see Ball, 2003 for a broad review). So, some opinion dynamics models represent a single 'opinion' as taking on either 'agree' or 'don't agree' (or possible 'disagree'; e.g., Latane & Nowak, 1997; Kacperski & Holyst, 2000, Sznajd-Weron & Sznajd, 2000). These models map closely onto the so-called "Ising spin model", which has been used for years to model physical ferromagnetism. As a consequence, it has many well-understood properties. However, it prevents incorporating any interesting notion of bounded influence, because other members of the population are either identical to an agent, or completely different. Any bound on influence in these models has no effect, because those who are similar to an agent are identical to the agent, and so cannot change its beliefs.

More recent advances in the field have begun to incorporate notions of bounded influence, but to do so, they required moving to a knowledge representation that is on a multi-value scale. The most common versions assume support of an opinion takes on a rational value between 0.0 and 1.0 (Dittmer, 2001; Hegselmann & Krause, 2002). When the influence threshold is large enough, a consensus typically emerges; when the influence threshold is smaller, multiple distinct islands of opinions may emerge. Extensions of these models have dealt with rules of bounded influence that are even more complex (Amblard & Deffuant, 2004; Deffuant, 2006; Franks, et al., 2008), and have only recently begun to take seriously psychological functions within their models (cf. Salzarulo 2006; Kopecky, Bos, & Greenberg, 2010).

The bounded influence conjecture is interesting, because it does not appear to be inspired by any documented cognitive or social phenomenon, and it is used to create an effect (distinct groups of disagreement) that is not an empirical phenomonon, but simply a reasonable possible state of the world. To be fair, Mueller & Veinott's (2008) cultural mixture modeling advocates the possibility of multiple distinct communities of belief within a culture, but it does not assume that this happens for a single dimension, but rather as an emergent cluster over a set of beliefs.

One explanation for the bounded influence conjecture is that it occurs because the impoverished representations of knowledge present in opinion dynamics models prevent a richer understanding of knowledge, disagreement, and the relationships between knowledge states. In a standard opinion dynamics model, the simplistic methods of knowledge representation coupled with simple interaction/influence means that in a typical standard situation, all agents will eventually align to form a consensus. These simple models are valuable because they are amenable to formal proofs of convergence, but their simplicity may necessitate assumptions (such as the bounded influence conjecture) that are not necessary and not warranted by the psychology of influence. Without bounded influence, given enough interactions, a system will typically converge to a consensus.

Yet cultural knowledge is much richer than is supposed by these models. If we allow just one level of increase in the complexity of data, we find cultural surveys like the Afrobarometer (e.g., Lewis, 2007), and the General Social Survey (e.g., Davis et al., 1998) that at least require one to view culture as a set of values in some type of space. But even more complexity can exist, which presents additional interesting possibilities.

For example, perhaps some sets of ideas depend on one another, or imply one another. Some sets of ideas may be internally consistent, while others are not. Perhaps social customs prevent discussion and influence about some types of concepts (politics and religion) but not other (television shows). These possibilities suggest an intriguing alternative: perhaps the bounds of influence are not parameters of an individual's willingness to listen to another's opinion, but rather are parameters of the viable states of knowledge.

Knowledge Spaces as a means to represent complex knowledge

Researchers in areas of education and measurement have developed formal algebras for representing complex knowledge systems that may provide a means to understand beliefs across a group (cf. Albert, 1994; Doignon, 1999; Falmagne et al., 2006). According to this research, one can represent a set of concepts as a partially-ordered set of knowledge states, such that any higher state requires the attainment of a lower state, or in other words, attainment of one state implies attainment of all lower states. A common example includes a student's attainment of the mathematical skill of complex division, which requires master of subtraction and simple division. If one knows the student understands complex division, one can assert that she also understands subtraction. These states form a lattice: a set of nodes in a network that are connected if one node implies understanding of another.

Wiley and Martin (1999), Martin and Wiley (2000), and Butts and Hilgeman (2003) have expanded and extended these ideas to characterize a knowledge lattice across a social group. Now, knowledge in this case is not just abstract, but the set of knowledge states found among a population is used to form the knowledge lattice. Each individual in the group lives within one of the nodes, and the network describes the variety of beliefs within the group. For example, Figure 1 shows a potential knowledge lattice for a set of world religions. Here, we describe religions in the Judeo-Christian-Islamic traditions based on the prophets they believe in. A vertical connection implies belief in everything connected lower than the node. So, Christians believe in Moses and Jesus, Muslims believe in Moses, Jesus, and Mohammed, and so on. Each node is closely affiliated with a world religion, and the movement between nodes involves adding or removing a single belief from ones belief system. The set algebra of course could represent a religious group that, for example, Moses was not a prophet but Jesus was, or that both Mohammed and Mormon were prophets but the Bab was not. This example highlights how, for a set of beliefs, there may be a more likely path from one belief to another (Conversion from Mormonism to Islam may be unlikely), or a set of beliefs happens to be inconsistent (e.g., Mormon was a prophet but Jesus was not).

A knowledge space may have a similar effect on preventing consensus that the bounded influence conjecture does, but without making the bounded influence assumption. Instead of simply saying that one will not be influenced by individuals who have an extremely different set of beliefs, one might hypothesize that movement between beliefs can only happen when the influence process lands a person in a valid belief state. So, in our example, perhaps a Jew and a Universalist will indeed listen to one another, and although they are at opposite ends of the belief spectrum, it may be possible that

Figure 1: Knowledge space characterizing several world religions. A connection implies that the higher-level concept in some sense requires or implies the beliefs at lower levels.

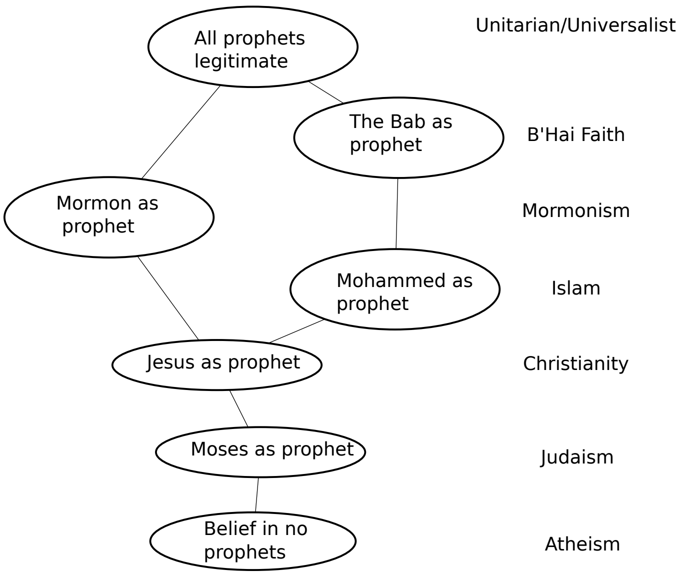

Figure 2: Knowledge lattice used in present simulations. Each of the fifteen states corresponds to a binary vector of twenty features. An arrow implies that a higher node contains all the positive (1) values that a lower node contains.

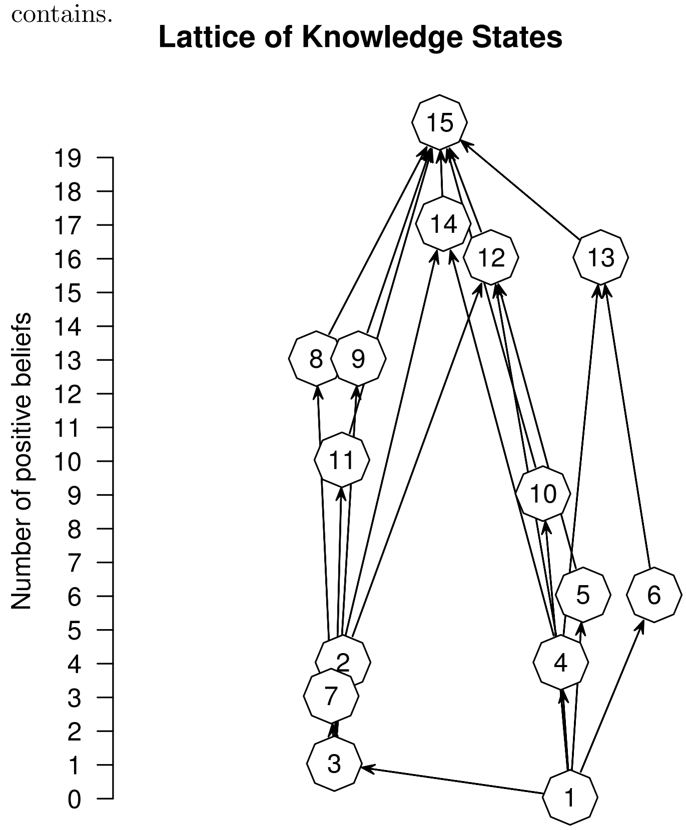

one moves toward the other *as long as they end up in a valid belief state*. So, the Universalist may not convert to Judaism, but may be convinced that aspects of Mormonism are false. However, if there is no consistent set of beliefs that allow one to hold "every religion but Mormonism is acceptable", they may not move away from their initial knowledge state. This provides a mechanism for belief clustering and consensus avoidance that critically depends on the knowledge structure, not simply arbitary sharing rules.

Simulation Model of Influence Dynamics via Belief States

To test the hypothesis that consensus and disagreement behaviors could arise from properties of the knowledge system rather than a parameterized bounded influence, we implemented a simple simulation to test the principles.

We began by creating a knowledge space based on twenty binary features, and selecting fifteen distinct states within the space. Figure 3 shows the space selected, and Table 1 shows the binary feature representations for each state.

Each simulation involved 100 agents, initialized with a specific distribution across knowledge states. We assumed that these fifteen states were the only valid states an agent could hold, but did not make assumptions about why. At each step of the simulation, a pair of agents were chosen randomly, and with probability $\mu = .3$, each feature of each agent was changed to be

Table 1: Feature representation of the fifteen knowledge states in Figure 1.

Knowledge State	Knowledge
1	00000000000000000000
2	00001000001001000010
3	00000000000000000010
4	00001100100000100000
5	00100100010101100000
6	01000001001111000000
7	00011000000000000010
8	00100111111110011110
9	10010111010111100111
10	01000001111000011101
11	01100011011100010011
12	11101110111101111110
13	11011111111111110001
14	11111110111110111110
15	11111111111111111111

identical to the other agent's. At the end of this exchange, each agent's new knowledge state was examined, and if it were identical to one of the predetermined knowledge states, this new state was retained. This last restriction forces all agents to live somewhere in the predetermined knowledge space. We computed the distribution across knowledge states after every ten exchanges to visualize the evolution of belief across time.

Figure 3 shows a fairly typical simulation for a starting configuration in which individuals were uniformly spread across belief states. Because there are many individuals at points along the path from one belief to another, initial conditions such as this tend to converge to a single belief state or a pair of adjacent states. The simulation continued 40000 exchanges, after which 86 agents were in state 1, and 14 were in state 13. Mixed initial configurations typically converge to a single belief state, near the extreme, although two adjacent belief states can be fairly stable. We have not witnessed convergence to a compromising belief state.

Opinion dynamics simulations often start out in random configurations, but this is probably unrealistic–we are all born into a society that has established norms and beliefs. Another starting configuration to consider is one in which individuals are in either one or another of the most extreme states (akin to the organization of the U.S. Senate, as discussed by Mueller & Veinott, 2008). Figure 4 shows the evolution of belief over time for such a starting configuration. The belief distribution remains fairly stable, at the two extremes of the belief spectrum (states 1 and 15). Although other states do emerge in their proportion, these quickly get subsumed into the larger more extreme groups.

Summary and Discussion

In this paper, we demonstrate how one common assumption about the way information is shared among people (the bounded influence conjecture) may not be a property of influence per se, but rather could be an aspect of a more complex network of potential knowledge states. This highlights one of the benefits of taking cognition more seriously when modeling social science. Previous assumptions about bounded influence were made in order to produce a behavior that seemed reasonable, but apparently without reference to psychological study of influence.

We believe that there are additional benefits that can be gained from considering more complex knowledge structure as the medium of cultural exchange in opinion dynamics models. For example, much cultural knowledge can take the form of cultural schemas and mental models Cultural Schemas/Models (e.g., D'Andrade, 2001), causal models (Rasmussen, Sieck, & Smart, 2009), decision trees (Gladwin, 1989), scripts (Ryan, 1996). By considering more complex representations, one needs to begin to consider important factors such as how much and what type of information is shared (rather than sharing all beliefs at each interaction), whether there are core aspects of a belief system that logically or traditionally cohere, and the ways messages might be created by one who is attempting to influence another, such that it has the greatest chance of changing their mind.

Acknowledgments

The research reported here was performed under Subcontract No. RQ000746 from SRA International, Inc. (Prime contract contract FA8650-09-D-6939) in support of the Air Force Research Labs, Wright-Patterson Air Force Base, OH. We would like to thank Dr. Michael J. Young with AFRL/RHXB and Scott Macbeth with SRA for their encouragement, guidance, and support.

References

Albert, D. (Ed.). (1994). Knowledge structures. New York: Springer-Verlag.

Amblard, F., & Deffuant, G. (2004). The role of network topology on extremism propagation with the relative agreement opinion dynamics. Physica A: Statistical Mechanics and its Applications, 343, 725-738.

Atran, S., Medin, D. L., & Ross, N. O., (2005). The Cultural Mind: Environmental decision making and cultural modeling within and across populations. Psychological Review, 112, 744-776.

Axelrod, R. (1986). An evolutionary approach to norms. The American Political Science Review, 80(4), 1095-1111.

Ball, P. (2003). The Physical Modelling of Human Social Systems. Complexus, 1(4), 190-206. doi:10.1159/000082449

Butts, C. T., & Hilgeman, C. (2003). Inferring Potential Memetic Structure form Cross-Sectional Data:An Application to American Religious Beliefs. Journal of Memetics - Evolutionary Models of Information Transmission, 7(2), 12-34. doi:Article

D'andrade, R. (2001). A cognitivist's view of the units debate in cultural anthropology. Cross-Cultural Research, 35(2), 242.

Davis, J. A. , Smith, W., & Marsden, P. V. (1998). General Social Survey, 1998. National Opinion Research Center, Chicago, 1998.

Deffuant, G. (2006). Comparing Extremism Propagation Patterns in Continuous Opinion Models. Journal of Artificial Societies and Social Simulation, 9(3), 8.

Dittmer, J. C. (2001). Consensus formation under bounded confidence. Nonlinear Analysis-Theory Methods and Applications, 47(7), 4615-4622.

Doignon, J. (1999). Knowledge spaces. Berlin ;New York: Springer.

Falmagne, J. C., Cosyn, E., Doignon, J. P., & Thiery, N. (2006). The assessment of knowledge, in theory and in practice. Lecture notes in computer science, 3874, 61.

Franks, D. W., Noble, J., Kaufmann, P., & Stagl, S. (2008). Extremism propagation in social networks with hubs. Adaptive Behavior, 16(4), 264.

Figure 3: Proportion of different knowledge states of the population over time

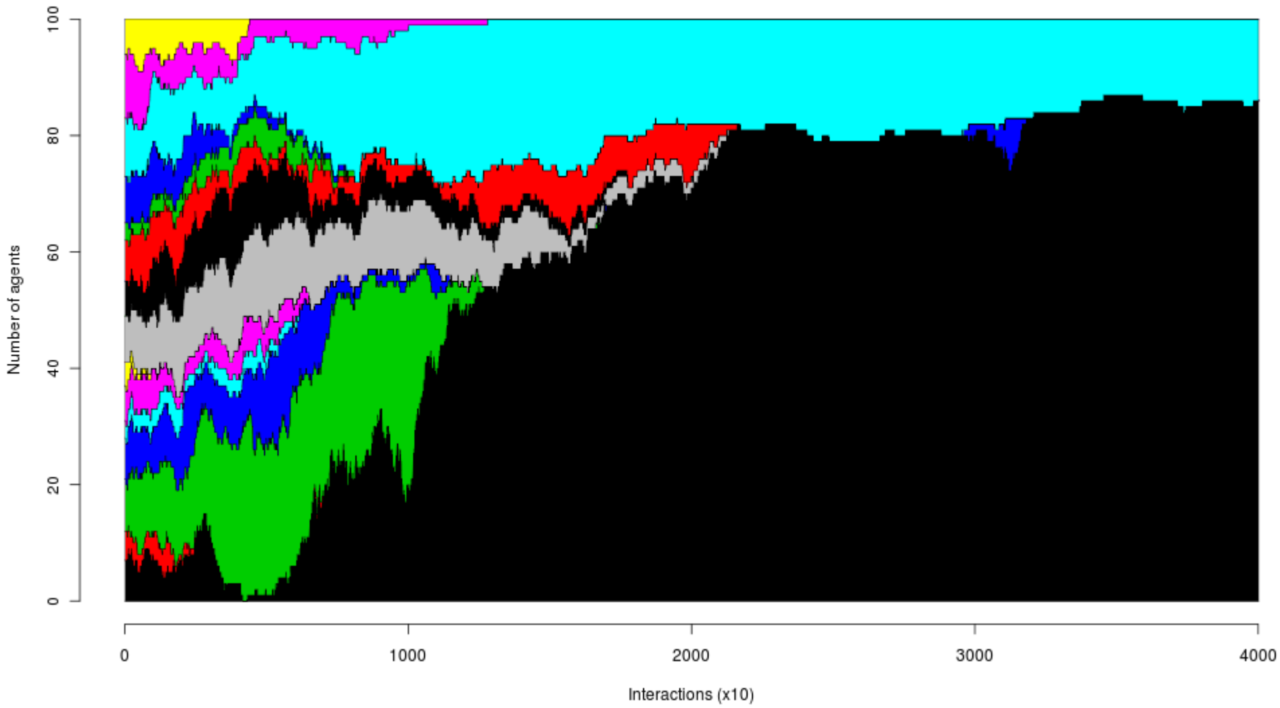

Figure 4: Proportion of different knowledge states of the population over time. Initially, all agents were in only the two most extreme states.

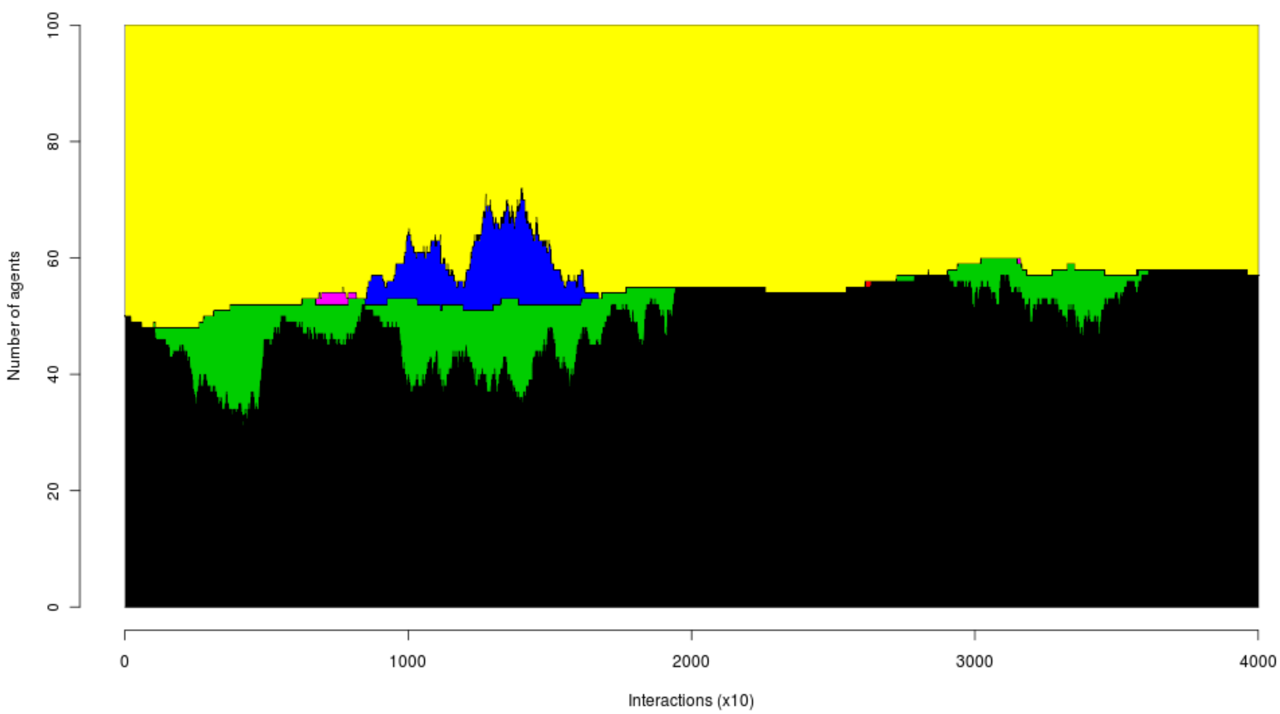

Hegselmann, R. K., & Krause, U. (2002). Opinion dynamics and bounded confidence models, analysis and simulation. Journal of Artificial Societies and Social Simulation, 5(3).

Kacperski K.[1], & Holyst J.A. (2000). Phase transitions as a persistent feature of groups with leaders in models of opinion formation. Physica A, 287, 631-643. doi:10.1016/S0378-4371(00)00398-8

Kopecky, J., Bos, N., & Greenberg, A. (2010). Social identity modeling: Past work and relevant issues for social-cultural modeling. BRIMS 2010, Charleston, VA.

Latan, B., & Nowak, A. (1997). Self-organizing social systems: Necessary and sufficient conditions for the emergence of clustering, consolidation, and continuing diversity. Progress in communication sciences, 43-74.

Lewis, P. M. (2007). Identity, institutions and democracy in Nigeria. Afrobarometer Working Paper #68, 397-424.

Lorenz, J. (2007). Continuous opinion dynamics under bounded confidence: A survey. International Journal of Modern Physics C, 18(12), 1819-1838.

Martin, J. L., & Wiley, J. A. (2000). Algebraic representations of beliefs and attitudes II: Microbelief models for dichotomous belief data. Sociological Methodology, 30, 123-164.

Mueller, S. T., & Veinott, E. S. (2008). Cultural mixture modeling: Identifying cultural consensus (and disagreement) using finite mixture modeling. In Proceedings of the 30th Annual Conference of the Cognitive Science Society (pp. 6470).

Rasmussen, L. J., Sieck, W. R., & Smart, P. (2009). What is a good plan? Cultural variations in expert planners concepts of plan quality. Journal of Cognitive Engineering & Decision Making, 3, 228-249.

Romney, A. K. (1999). Culture consensus as a statistical model. Current Anthropology, 40 (Supplement), S103S115.

Romney, A. K., Boyd, J. P., Moore, C. C., Batchelder, W. H., & Brazill, T. J. (1996). Culture as shared cognitive representations. Proc. Natl. Acad. Sci. USA, 93, 4699-4705.

Romney, A. K., Weller, S. C., & Batchelder, W. H. (1986). Culture as consensus: A theory of culture and informant accuracy. American Anthropologist, 313-338.

Salzarulo, L. (2006). A continuous opinion dynamics model based on the principle of meta-contrast. Jasss-the Journal of Artificial Societies and Social Simulation, 9(1).

Sznajd-Weron, K., & Sznajd, J. (2000). Opinion evolution in closed community. International Journal of Modern Physics C-Physics and Computer, 11(6), 1157-1166.

Wiley, J. A., & Martin, J. L. (1999). Algebraic Representations of Beliefs and Attitudes: Partial Order Models for Item Responses. Sociological Methodology, 29, 113-146.