More Contentious Issues Lead to More Factions:
Bounded Confidence Opinion Dynamics of Bayesian Decision Makers

Kush R. Varshney
Business Analytics and Mathematical Sciences Department
IBM Thomas J. Watson Research Center
1101 Kitchawan Road, Yorktown Heights, New York, USA
krvarshn@us.ibm.com

People are often asked for their opinion. Will *Son of Sardaar* be a blockbuster or a flop? Who will win the presidential election, Obama or Romney? Will Oracle Database or IBM DB2 have more sales next year? Which HMO health plan will enroll more patients, AARP MedicareComplete Plan 2 or Anthem Senior Advantage Value? Are burnt sienna sweaters going to be popular this winter?

In all of these questions, one of the hypotheses is true and the other one is false after the fact (assuming consistent definitions of blockbuster, flop, win, and popular). The job of the opinion-giver is to try to report the true hypothesis as his or her opinion. Thus, such opinion questions can be framed as Bayesian binary hypothesis testing or signal detection problems because human decision makers are tasked with predicting the true hypothesis based on noisy observations of the world and their prior belief with minimal probability of error or Bayes risk [1]. Under such a model, an opinion is composed of both an observation and a prior.

The observation component of opinion can have varying levels of noise. For some opinions, observable features are very predictive of the hypothesis whereas in others, the observable features are not very predictive at all. In society, the prior beliefs of decision makers are mutable and influenced by their network of acquaintances. People sway others one way or the other. However, most people can only be swayed a little bit at a time; they are not influenced by people with a very different prior. This non-influence of people with beliefs too far away is known as bounded confidence [2].

One may ask why signal detection questions rather than more direct consumer choice questions such as whether the decision maker will personally buy a burnt sienna sweater or who the decision maker will vote for are of interest. One reason is that it has been empirically shown that aggregating answers from the former type of question (expectation question) yields more accurate forecasts than the latter type of question (intention question) [3]. Also, the former questions are more in line with the popular definition of opinion.

Opinion dynamics—how opinions evolve among the people in a social network over time—has been studied in a variety of literatures recently [4]. Two discrete-time models of note that incorporate the idea of bounded confidence are the Krause model and the Deffuant-Weisbuch model [2],[5],[6]. The Krause model in particular begins with an initial assignment of opinions for each person; for all people in the society, it works by finding all others whose opinion is within a certain absolute value of one's own and updating the own opinion to be the mean of those opinions. The model has been analyzed to show that opinions quickly converge to a few values [7], resulting in clusters of people with the same final opinion.

The existing literature on bounded confidence opinion dynamics only considers opinions as abstract real-valued numbers without connection to a decision-making formulation. In this work, it is proposed that opinion-giving be treated as a Bayesian signal detection task and that bounded confidence dynamics be applied to the decision makers' prior probabilities. Such a view of opinion dynamics through a decision-theoretic lens seems not to exist in the literature.

When one is to quantify dissimilarity in prior probabilities of Bayesian hypothesis testing, absolute error is not appropriate because it treats all prior probabilities the same. The difference between a prior probability of 0.1 and 0.2 leads to a Bayes risk difference that may be profoundly different than the Bayes risk difference between a prior probability of 0.4 and 0.5, which is also different when there are different observation models and signal-to-noise ratios. Bayes risk error divergence, a member of the family of Bregman divergences, is a criterion by which the dissimilarity of two input prior probabilities is examined appropriately according to detection performance [8],[9]. Thus, in this work, it is proposed that Bayes risk error be used to define bounded confidence. Due to a property of Bregman divergences [10], the Bayes risk error centroid of a set of prior probabilities is simply the mean; thus, the opinion update remains unchanged from the standard Krause model when Bayes risk error is considered.

As seen in Fig. 1 empirically, the convergence behavior under Bayes risk error bounded confidence and absolute error bounded confidence is qualitatively similar. Both converge quickly to a few clusters while maintaining the initial ordering. Under Bayes risk error, the convergence is faster and the middle cluster contains fewer people. Most people end up on one side of the fence or the other rather than straddling it.

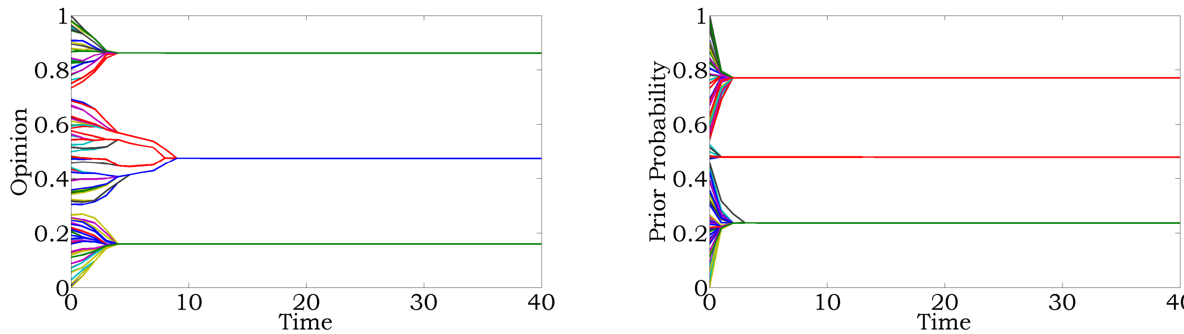

(a) (b)

Figure 1. Dynamics of 100 people with bounded confidence threshold 0.1 for (a) standard Krause model and (b) Bayes risk error divergence Krause model with univariate Gaussian observations having means 0 and 1, and standard deviation 4.

Since (unlike absolute error) the Bayes risk error divergence depends on the likelihood functions of the observation model, bounded confidence opinion dynamics can be investigated under different signal-to-noise ratios. For smaller and greater noise levels, as seen in Fig. 2, the convergence pattern as well as the number of final clusters is different than an intermediate amount of noise.

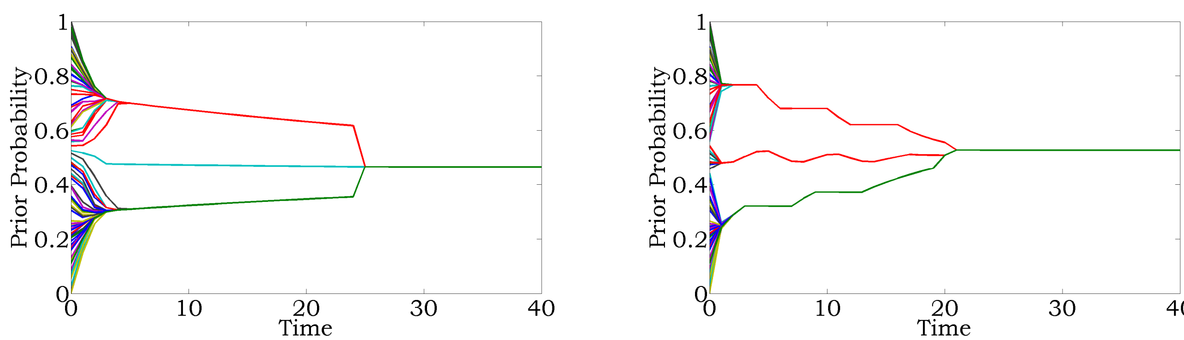

(a) (b)

Figure 2. Dynamics of 100 people with bounded confidence threshold 0.1 for Bayes risk error divergence Krause model with univariate Gaussian observations having means 0 and 1, and (a) standard deviation 1 and (b) standard deviation 16.

The number of converged clusters (averaged across several trials of initial decision maker prior probabilities) as a function of noise level for a fixed bounded confidence is shown in Fig. 3. The number of clusters first increases and then decreases with the noise level. When observations are either very certain or very uncertain, there is convergence to a single cluster; however, in the intermediate case there are more clusters. The interpretation is that at very low noise, the observation readily indicates the true hypothesis so priors do not matter, and thus everyone converges to an average, maximally ignorant prior. The intermediate noise regime corresponds to the most contentious issues and opinions because the observation is

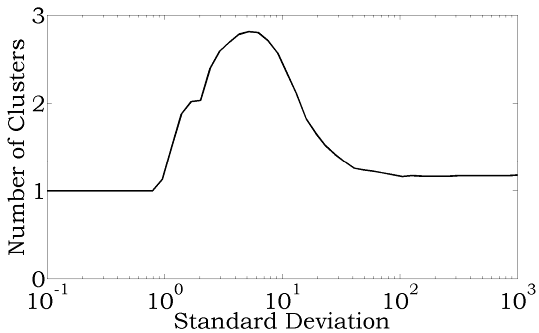

Figure 3. Converged number of clusters of 100 people with bounded confidence threshold 0.1, averaged over 200 trials, for Bayes risk error divergence Krause model with univariate Gaussian observations having means 0 and 1, and standard deviation as plotted.

somewhat informative of the true hypothesis, but not fully. The choice of prior has a large effect on the opinion, and thus these contentious issues lead to factions. With very large amounts of noise, no one has any idea what is going on, and everyone is uninformed; again, convergence is to an average, uninformed consensus of priors. Although such a phenomenon is known to occur in society, to the best of the author's knowledge, there has been no previously proposed opinion dynamics model that generates more factions when the issue is more contentious.

The empirical results shown above look at the basic Krause model extended via Bayes risk error divergence. Further study examines a more realistic version of the Krause model in that only other people with whom a social connection exists are included in the opinion update mean [11]. Other further study incorporates certain decision makers that either do not change their opinion or try to aggressively push their beliefs [12]. A third piece of further study is vector-valued opinions or prior probabilities [13], which are readily handled with Bayes risk error divergence [8].

Acknowledgment

The author thanks Michael Rabbat and Lav R. Varshney for discussions.

References

[1] H. L. Van Trees, *Detection, Estimation, and Modulation Theory*. New York, NY: Wiley, 1968.

[2] R. Hegselmann and U. Krause, "Opinion Dynamics and Bounded Confidence Models, Analysis, and Simulations," *J. Artif. Societies Social Simul.*, vol. 5, no. 3, 2002.

[3] D. Rothschild and J. Wolfers, "Forecasting Elections: Voter Intentions versus Expectations," Univ. Penn., Philadelphia, PA, Working Paper, 2011.

[4] J. Lorenz, "Continuous Opinion Dynamics Under Bounded Confidence: A Survey," *Int. J. Modern Phys. C*, vol. 18, no. 12, pp. 1819–1838, 2007.

[5] U. Krause, "Soziale Dynamiken mit vielen Interakteuren. Eine Problemskizze," in *Model. Simulat. Dyn. vielen interagierenden Akteuren*, Bremen, Germany, pp. 37–51, 1997.

[6] G. Deffuant, D. Neau, F. Amblard, and G. Weisbuch, "Mixing Beliefs Among Interacting Agents," *Adv. Complex Syst.*, vol. 3, no. 1–4, pp. 87–98, 2000.

[7] V. D. Blondel, J. M. Hendrickx, and J. N. Tsitsiklis, "On Krause's Multi-Agent Consensus Model with State-Dependent Connectivity," *IEEE Trans. Automat. Contr.*, vol. 54, no. 11, pp. 2586–2597, 2009.

[8] K. R. Varshney, "Bayes Risk Error is a Bregman Divergence," *IEEE Trans. Signal Process.*, vol. 59, no. 9, pp. 4470–4472, 2011.

[9] K. R. Varshney and L. R. Varshney, "Quantization of Prior Probabilities for Hypothesis Testing," *IEEE Trans. Signal Process.*, vol. 56, no. 10, pp. 4553–4562, 2008.

[10] A. Banerjee, S. Merugu, I. S. Dhillon, and J. Ghosh, "Clustering with Bregman Divergences," *J. Mach. Learn. Res.*, vol. 6, pp. 1705–1749, 2005.

[11] M. Rabbat, "Bounded Confidence Opinion Dynamics with Network Constraints and Localized Distributed Averaging," in *Proc. IEEE Statist. Signal Process. Workshop*, Ann Arbor, MI, pp. 636–639, 2012.

[12] E. Yildiz, D. Acemoglu, A. Ozdaglar, A. Saberi, and A. Scaglione, "Discrete Opinion Dynamics with Stubborn Agents," Lab. Inf. Decision Syst., Mass. Inst. Technol., Cambridge, MA, Tech. Rep. P-2858, 2011.

[13] R. Vicente, A. C. R. Martins, and N. Caticha, "Opinion Dynamics of Learning Agents: Does Seeking Consensus Lead to Disagreement?," *J. Stat. Mech.*, p. P03015, 2009.