A short version of this paper is forthcoming in the Proceedings of AAMAS 2008, Estoril, Portugal.

Modeling and Managing Collective Cognitive Convergence

H.V. Parunak, T.C. Belding, R. Hilscher, S. Brueckner
NewVectors division of TTGSI
3520 Green Court, Suite 250
Ann Arbor, MI 48105 USA
+1 734 302 4684

{van.parunak, ted.belding, sven.brueckner, rainer.hilscher}@newvectors.net

ABSTRACT

When the same set of people interact frequently with one another, they grow to think more and more along the same lines, a phenomenon we call "collective cognitive convergence" (C^3). In this paper, we discuss instances of this phenomenon and why it is advantageous or disadvantageous; review previous work in sociology, computational social science, and evolutionary biology that sheds light on C^3; define a computational model for the convergence process and quantitative metrics that can be used to study it; report on experiments with this model and metric; and suggest how the insights from this model can inspire techniques for managing C^3.

Categories and Subject Descriptors

J.4 [**Computer Applications**]: Social and Behavioral Sciences – *Sociology*; I.6.3 [**Simulation and Modeling**]: Applications

General Terms

Measurement, Experimentation, Human Factors, Theory.

Keywords

Groupthink, cognitive convergence, modeling, social simulation.

1. INTRODUCTION

When the same set of people interact frequently, they grow to think more and more along the same lines. We call this phenomenon "collective cognitive convergence" (C^3), since the dynamics of the *collective* lead to a *convergence* in *cognitive* orientation.

C^3 arises in many contexts, including research subdisciplines, political and religious associations, and even persistent adversarial configurations such as the cold war. Tools that support collaboration, such as blogging, wikis, and communal tagging, make it easier for people to find and interact with others who share their views, and thus may accelerate C^3. This efficiency is sometimes desirable, since it enables a group to reach consensus more quickly. For instance, in the academy, it enables coordinated research efforts that accelerate the growth of knowledge.

But convergence can go too far, and lead to collapse. It reduces the diversity of concepts to which the group is exposed and thus leaves the group vulnerable to unexpected changes in the environment. Here are two examples.

In academia, specialized tracks at conferences sometimes become unintelligible to those who are not specialists in the subject of a particular track, and papers that do not fit neatly into one or another subdiscipline face difficulty being accepted. The subdiscipline is increasingly sustained more by its own interests than by the contributions it can make to the broader research community or to society at large.[1]

In military operations, the force-on-force orientation developed during the Cold War left both the former Soviet Union and NATO ill-prepared to deal with insurgencies and asymmetric warfare.

Groups that have undergone cognitive collapse will only produce output conforming to their converged set of ideas, and will be unable to conceive or explore new ideas. In the worst case, collapse may lead a group to focus its attention on a cognitive construct with little or no relation to the real world. For example, highly specialized academic disciplines become increasingly irrelevant to people outside of their own circle.

We became interested in this phenomenon by observing increasing balkanization in the research field of multi-agent systems. Since we work in the area of multi-agent simulation, it occurred to us that some light might be shed on the phenomenon, and on how it can be managed, with a multi-agent model. This paper presents some preliminary results.

Section 2 discusses previous work related to our effort. Section 3 describes our model, and a metric that we use to quantify C^3. Section 4 outlines a series of experiments that exhibit the phenomenon and explore possible techniques for managing it. Section 5 suggests directions for further research, and Section 6 concludes.

2. PREVIOUS WORK

Our research on C^3 builds on previous work in sociology (both empirical and theoretical) and evolutionary biology.

Cite as: Modeling and Managing Collective Cognitive Convergence (Short Paper), H.V. Parunak, T.C. Belding, R. Hilscher, and S. Brueckner, *Proc. of 7th Int. Conf. on Autonomous Agents and Multiagent Systems (AAMAS 2008)*, Padgham, Parkes, Müller and Parsons (eds.), May, 12-16., 2008, Estoril, Portugal, pp. XXX-XXX.

[1] This paper was motivated by frustration voiced in the industry track at AAMAS07 about how some subdisciplines of agent research were becoming so intellectually ingrown, focusing only on problems defined by other members of the subdiscipline, that it was difficult or impossible to apply them to real problems.

There is abundant empirical evidence that groups of people who interact regularly with one another tend to exhibit C^3. Sunstein [24] draws attention to one version of this phenomenon, "group polarization": a group with a slight tendency toward one position will become more extreme through interaction. This dynamic suggests that confidence in group deliberation as a way of reaching a moderating position may be misplaced. He summarizes many earlier studies, and attributes the phenomenon to two main drivers: social pressure to conform, and the limited knowledge in a delimited group. Our model captures the second of these drivers, but not the first. Sunstein suggests some ways of ameliorating the problem that we explore with our model.

Table 1: Representative Studies in Consensus Formation

Study	Belief	Topology	Arity	Preference?
Krause [16]	Real variable	Random	Many	Yes
Sznajd-Weron [25]	Binary variable	Lattice	Two	No
Deffuant [6]	Real variable	Random	Two	Yes
	Binary vector	Random	Two	Yes
Axelrod [2]	Nominal[2] vector	Lattice	Two	Yes
Bednar [3]	Nominal vector[3]	Random	Many	No
This paper	Binary vector	Random	Many	Yes

Computational social science has long been preoccupied with the dynamics of consensus formation. One recent review [13] traces relevant work back more than 50 years [10]. These studies include analysis, simulation, or sometimes both. Their models differ along several important dimensions, including include the belief model and three characteristics of agent interaction (topology, arity, and preference). Rather than attempting an exhaustive review, we situate our work in these dimensions.

- The model of an agent's belief can be either a single variable or a vector, with values that can be real, binary, or nominal. Vector models usually represent a collection of beliefs, but in one study [3] the different entries in the vector represent the value of the same belief that underlies different behaviors, to explore of internal consistency.
- In some models agent interactions are constrained by agent location in an incomplete graph, usually a lattice (though one study [17] considers scale-free networks). In others any agents can interact (often called the "random choice" model).
- Agents to interact only two at a time, or as larger groups.
- The likelihood of agent interaction may be modulated by their similarity.

Table 1 characterizes several papers in this area in terms of these dimensions. Our work represents a unique combination of these characteristics. In particular,

- We consider a vector of m beliefs, rather than a single belief. This model allows us to look at how an individual may participate in different interest groups based on different interests, but also makes describing the dynamics much more difficult than with a single real-valued variable. In the latter case, individuals move along a linear continuum, and measures such as the mean and variance of their position are suitable metrics of the system's state. In our case, they live on the Boolean lattice $\{0,1\}^m$ of interests, and our measures must reflect the structure of this lattice.
- We allow many individuals to interact at the same time. This convention captures the dynamics of group interaction more accurately than does pairwise interaction, but also means that our agents interact with a probability distribution over the belief vector rather than a single selection from such a distribution.
- We allow our agents to modulate the likelihood of interaction based on how similar they are to their interaction partners. This kind of interest-based selection is critical to the dynamics of interest to us, but makes the system much more complex.

One consequence of selecting the more complicated options along these dimensions is that analytic results, accessible with some (but by no means all) simpler models, become elusive. Almost all analytical results in this discipline are achieved by modeling the belief of agent i as a single real number x_i and studying the evolution of the vector $\boldsymbol{x}$ over time as a function of the row-stochastic matrix $\boldsymbol{A}$ whose elements a_{ij} indicate the weight assigned by agent i to agent j's belief, $\boldsymbol{x}(t+1) = \boldsymbol{A}\boldsymbol{x}(t)$. This model captures interaction arity greater than two, but not vector beliefs or agent preferences. Conditions for convergence under preferences have been obtained [16], but only for six or fewer agents [13]. Bednar et al. [3] have derived convergence times for a form of vector belief, but only for binary interactions and with no preferences. Even for binary interactions, the combination of vector-based beliefs and preferences has resisted analytical treatment (in studies of an isomorphic system, bisexual preferential mating [14, 21]).

Given this research context, in this paper we focus our attention on simulation results, to develop intuitions that may reward future analytical exploration.

The subgroups that form and cease to interact when convergence turns to collapse are reminiscent of biological species, which do not interbreed. So we look for insight to research in the field of biological speciation (see [5, 11] for reviews). The most commonly proposed speciation mechanisms are allopatric speciation, sympatric speciation, and parapatric speciation. In allopatric speciation, genetic barriers gradually evolve between two or more geographically isolated species. This might happen for instance between organisms living on separate islands. These barriers could evolve either through natural selection or through other means such as the founder effect (i.e., differences in genes between populations due to the small sample sizes of the founding populations). One configuration of our model can be interpreted as exhibiting allopatric speciation.

In parapatric speciation, there is no discrete barrier between populations; individuals are distributed along a geographic continuum and are separated by distance. Finally, sympatric speciation refers to instances where a single population with no physical or geographic gene flow barriers divides into separate species. Two interacting forces are required for sympatric speciation to occur: 1) a force that drives sympatric speciation

[2] [9] finds faster convergence when some elements in the vector function as interval variables.

[3] All entries reflect the same belief in different behavioral settings, and pressure toward internal consistency is part of the model dynamics.

(e.g. resource competition or sexual selection) and 2) assortative mating that generates phenotypic variability and maintains evolving phenotypic clusters that eventually become species. Assortative mating refers to a mating system where different individuals express preferences for different phenotypes (e.g. some female birds prefer males with red feathers and other females prefer males with blue feathers). Some configurations of our model correspond to sympatric speciation.

Sexual selection [1, 8] roughly refers to the differential mating success of individuals in a population, Sexual selection can either be based on an asymmetric mating system (males compete and females choose) or a symmetric mating system (mutual mate choice where both sexes compete and choose). One sexual selection mechanism is Fisher's runaway process, which leads to extravagant traits in males that are detrimental to their survival.

While the relative importance and frequency of these speciation mechanisms in nature are still heavily debated, the mathematical prerequisites for each mechanism have been extensively studied [5, 11, 15]. This work could be adapted to predict when and how C^3 will develop, and how it can be managed.

Our C^3 model can be considered an instance of a runaway sexual selection speciation model with mutual mate choice. We assume a homogenous environment, no physical barriers for the exchange of ideas and a symmetric "mating system" where individuals express their "mating preferences" (i.e. their preference for an atomic interest; see Section 3 below) mutually. In our model, a preference for extreme traits is modeled as the probability of adopting an interest based on the prevalence of this interest in a given neighborhood. A successful runaway process in our model can be viewed as the development of academic specializations with little practical relevance.

There has also been much theoretical work done to study the amount of gene flow or migration that is necessary to prevent isolated populations of organisms from diverging or losing diversity due to genetic drift, or sampling error [12]. Sewall Wright argued in his Shifting Balance Theory that a subdivided population with intermittent migration could exhibit more rapid evolutionary change than a single cohesive breeding population [22]. The mathematical frameworks for studying migration could be applied to modeling the exchange of ideas or individuals between groups in C^3, and the amount of exchange that is necessary to prevent intellectual isolation.

3. A MODEL AND METRICS

We have constructed a simple multi-agent model of C^3 to study this phenomenon. Our model represents each participant's interests as a binary vector. Each position in the vector corresponds to an atomic interest. A '1' at a position means that the participant is interested in that topic, while a '0' indicates a lack of interest. At each step, each participant

- identifies a neighborhood of other participants based on some criteria (which may include proximity between their interest vectors, geographical proximity, or proximity in a social network, criteria that correspond to differences among various forms of biological speciation),
- learns from this neighborhood (by changing an interest *j* currently at 0 to 1 with probability $p_{interest}$ = proportion of neighbors having interest *j* set to 1), and
- forgets (by turning off an interest *j* currently at 1 to 0 with probability $1 - p_{interest}$).

One boundary condition requires attention. If an agent has no neighbors, what should $p_{interest}$ be? We take the view that interests are fundamentally social constructs, persisting only when maintained. Thus an isolated agent will eventually lose interest in everything, and in our model, a null community leads to $p_{interest}$ = 0 for all interests. Alternative assumptions are certainly possible, and would lead to a different model.

We need quantitative measures of agent convergence to study C^3 systematically. (Sophisticated statistical techniques exist for estimating the consensus of a group of people empirically, based on their responses to questionnaires [23]. For our purposes, the abstract measures here are more suitable.) To derive our measures, we cluster the population hierarchically based on cognitive distance between agents (in our case, the Jaccard distance between their interest vectors). Each node of the resulting cladogram forms at a specific distance (the "diameter" of the cluster represented by that node). The root has the highest diameter. In a random population of agents, the distances at which

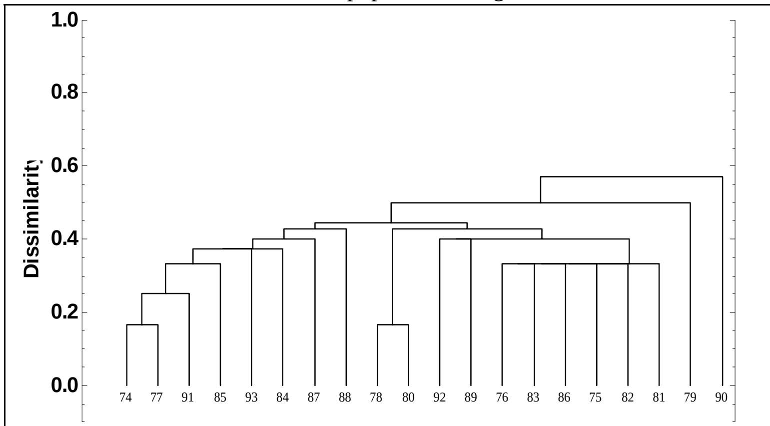

Figure 1 Random interest vectors, median min-max ratio = 0.583.

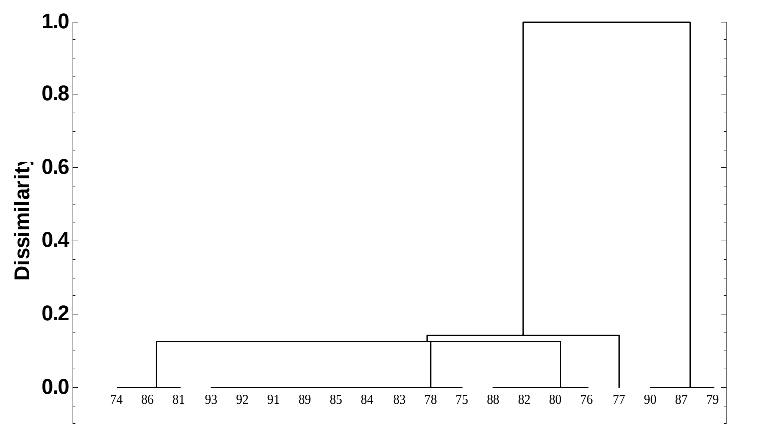

Figure 2 A highly converged population, median min-max ratio = 0

lower-level nodes join the tree is not much less than the diameter of the root (Figure 1), while in highly converged populations, the diameters of lower-level nodes are much less than the diameter at the root (Figure 2, where agents grouped at diameter 0 have identical interest vectors). Thus we compute the ratio of node diameter to root diameter (the "min-max ratio") for each node, and use the median of this ratio as a measure of overall system convergence. A ratio of 0 (as in Figure 1) means that more than half of the agents belong to groups within which all interest vectors are identical. We also record the maximum diameter of the clustered population at each generation. Convergence can lead to a diameter of 1 (when the population fragments into groups with orthogonal interest vectors that collectively span the interest space), a low value, asymptotically 0 (when all agents collapse toward a single group that is restricted to a single point in the interest space), or intermediate values (when groups have overlapping interests but no way to communicate about them to drive further convergence).

Figure 3 shows the behavior of the min-max ratio over a sample run of the system with 20 agents and interest vectors of length 10, where the probability of learning and forgetting is equal, and where agents are considered to be in the same group if the similarity between their interest vectors (the similarity threshold) is greater than 0.5. It takes only about 80 generations for the median min-max ratio to reach 0. (A generation consists of selecting one agent, choosing its neighbors, choosing with equal probability whether it shall attempt to learn or forget, selecting a bit in its interest string at random, then if it is learning and the bit is 0, flipping the bit with probability p_{learn} * $p_{interest}$, or if it is forgetting and the bit is on, flipping the bit with probability p_{forget} * (1 $p_{interest}$.)) Figure 2 shows the state of this system at generation 300. By generation 370 it has collapsed into two groups of completely homogeneous agents of sizes 3 and 17 respectively. The diameter of the overall population in this configuration is 1, representing orthogonal interest vectors. Among themselves, the agents still cover the entire interest space, but because they choose to interact only with the agents nearest themselves in that space, they form separate islands that cannot interact.

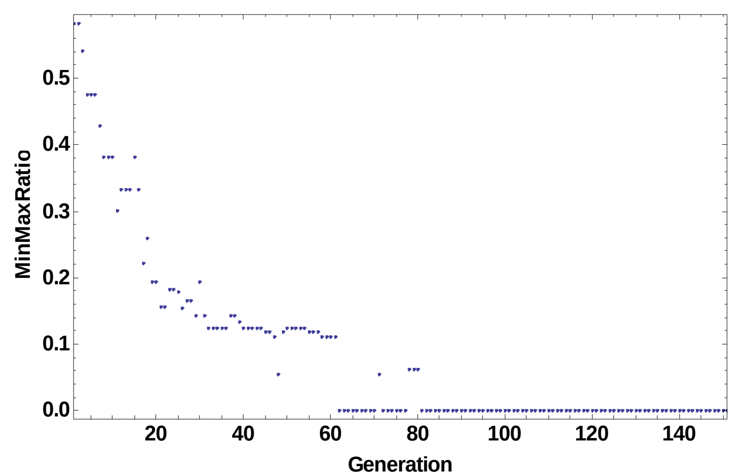

Figure 3 Evolution of 20 agents with length-10 interest vectors, neighborhoods defined by similarity > 0.5

4. SOME EXPERIMENTS

Armed with this model and metric, we can explore the dynamics of C[3] under a variety of circumstances. As we might expect, forming neighborhoods based on similarity of interest leads to rapid cognitive convergence. But surprisingly, other sorts of neighborhoods also lead to convergence.

4.1 Things that Don't Work

We might think that highly tolerant agents, those that consider all agents their neighbors, might be more robust to convergence. Figure 4 shows the evolution of the same population of agents when two agents consider one another neighbors if their similarity is greater than 0 (that is, they have at least one bit position in common). This configuration might be a model for a conference that has only plenary sessions. The population still collapses—this time, toward a maximum diameter of 0, indicating that the entire population tends to a single point in interest space.

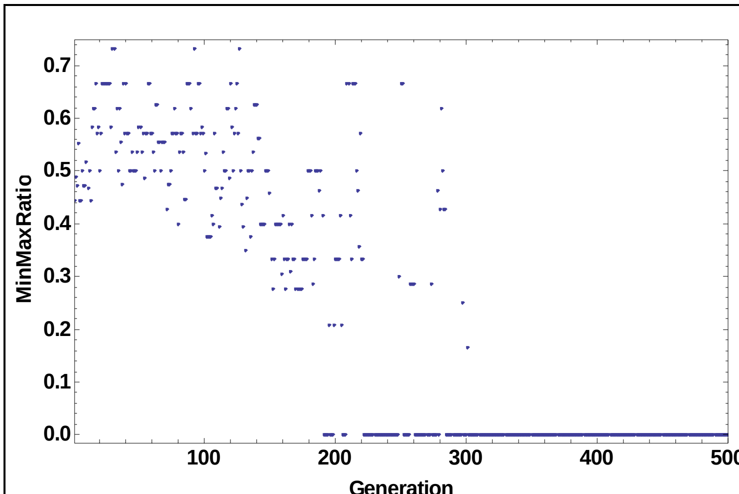

Figure 4: Zero Threshold

Perhaps the problem is that as agents converge, their neighborhoods increase in size. Figure 5 shows the effect of defining an agent's neighborhood at each turn as the group of four other agents that are closest to it. This configuration models a conference with separate tracks, organized by the common interests of their members. It corresponds to sympatric speciation: the assortative component is provided by the preference for partners with similar interests, while the limit on group size provides pressure toward diversity. Though agents base their adaptation at each turn on only 20% of the other agents, the min-max ratio still goes to zero, as agents form subgroups within which interests collapse. The maximum diameter freezes at an intermediate value (in this case, 0.6). The population has lost some but not all of its variation, but as in Figure 3, the selection of partners by interest proximity means that agents never interact with those who differ with themselves.

Figure 6 shows an even more radical approach. Here an agent's neighbors at each step are four randomly chosen agents. Imagine a conference at which papers are assigned to tracks, not by topic, but randomly. In spite of the mixing that this random selection provides, the population again collapses. The population diameter in this case asymptotes to 0, a single point in interest space

These examples differ in how long it takes the system to converge to a min-max ratio of 0. The time to convergence is highly variable, even within a single configuration. Repeated runs show that we should not assume that because (say) Figure 5 converges faster than Figure 4, small groups will always lead to faster convergence than highly tolerant agents. The one constant across all runs is that the system does converge, in fewer than 500 generations (often far fewer).

4.2 Introducing Variation

The collapse of agent interests is abetted by the lack of any mechanism for introducing variation. Once the population loses the variation among agents, it cannot regain it. We have explored three mechanisms for adding variation to the population: random mutation, curmudgeons, and interacting subpopulations.

The simplest approach is mutation. At each generation, with some small probability p_{mutate}, after learning or forgetting, the active agent selects a bit at random and flips it. This mechanism models spontaneous curiosity on the part of agents. Figure 7 shows an extended run with parameters the same as in Figure 3 (neighborhoods defined by a similarity threshold of 0.5), but with p_{mutate} = 0.03. Mutation is certainly able to reintroduce variation, but the level is critical. If mutation is too low (say, 1%), it is unable to keep up with the pressure to convergence, while if it is too high (10%), the community does not exhibit any convergence at all (and in effect ceases to be a community). The nature of its contribution follows a clear pattern. When it is in the critical range, the system occasionally collapses to a min-max ratio of 0, but then discovers new ideas that reinvigorate it. The population diameter under mutation converges to 1, since even when mutation is too low to avoid collapse within groups, it can introduce new interest vectors that are orthogonal to the converged groups.

A curmudgeon is a non-conformist, someone who regularly questions the group's norms and assumptions. Sunstein [24] observes that "group members with extreme positions generally change little as a result of discussion," and serve to restrain the polarization of the group as a whole.

To model curmudgeons, recall that ordinarily agents learn by flipping a 0 bit to 1 with probability $p_{interest}$, the proportion of neighbors that have the bit on, and forget by flipping a 1 bit with probability equal to $1 - p_{interest}$. To model curmudgeons, when an agent decides to learn or forget, with probability p_{cur},, it reverses these probabilities. That is, its probability of forgetting when it

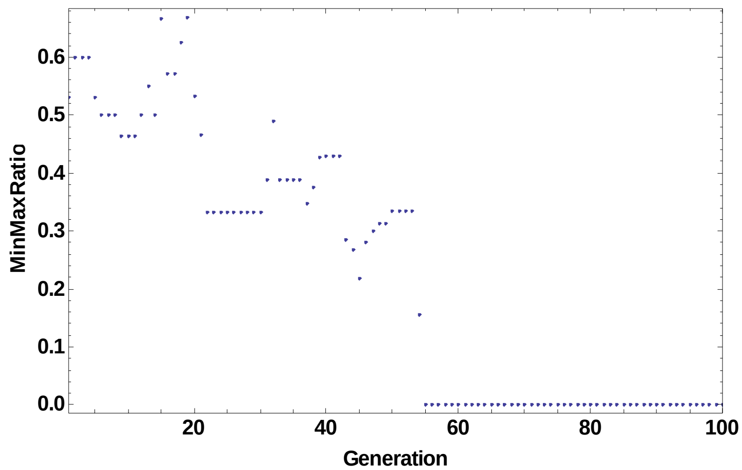

Figure 5 Fixed-size agent neighborhoods (four closest agents)

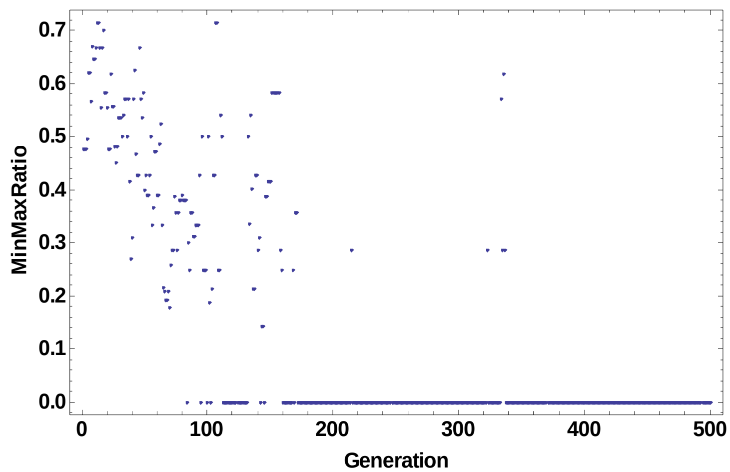

Figure 6 Neighborhoods of four randomly-chosen agents

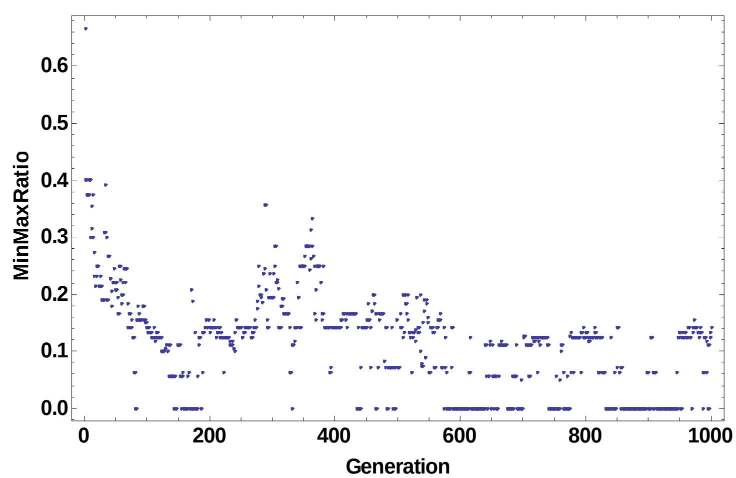

Figure 7 Adding 3% mutation

is curmudgeonly is $p_{interest}$ (instead of 1 - $p_{interest}$ in the non-curmudgeonly state), and its probability of learning is $1 - p_{interest}$.

Figure 8 shows the effect of 10% curmudgeons, again with the baseline configuration of Figure 3. The system clearly converges, but seldom reaches a min-max ratio of 0. Furthermore, p_{cur} can achieve this balancing effect over a much wider range than p_{mutate}. The population diameter tends to 1, reflecting the addition of diversity. As much as researchers may resent reviewers and discussants who "just don't get it," curmudgeons are an effective and robust way of keeping a community from collapsing.

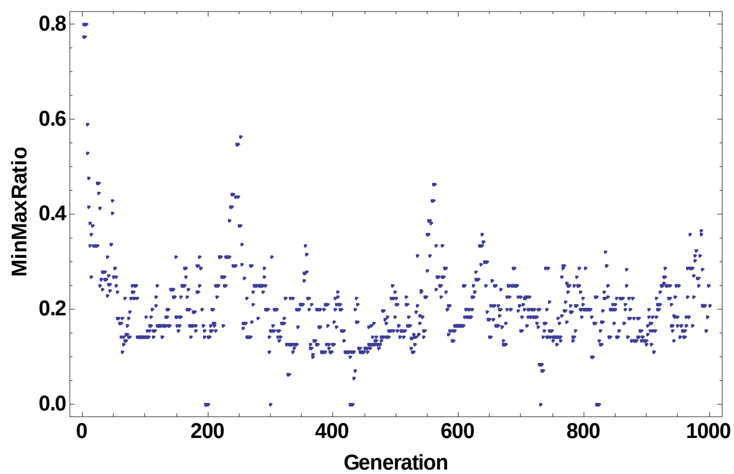

Figure 8 10% curmudgeons

The third source of variation is even more robust, and somewhat surprising, as the source of variation is endogenous rather than exogenous. So far, our agents have chosen a new set of neighbors at every step, based on their current set of interests. What would happen if we assign each agent to a fixed group at the outset, using a fixed similarity threshold that allows groups of various sizes to form?

The behavior depends on the structure of the graph induced by a given threshold. Figure 9 shows how the number of components depends on the threshold for groups formed in populations of 20 agents with 10 interests each. The sudden shift from many components at 0.6 to a few at 0.55 is an instance of the well-known phase transition in random graphs in which a giant connected component emerges as the number of links increases [7], in this case as a result of lowering the threshold. Four cases merit our attention.

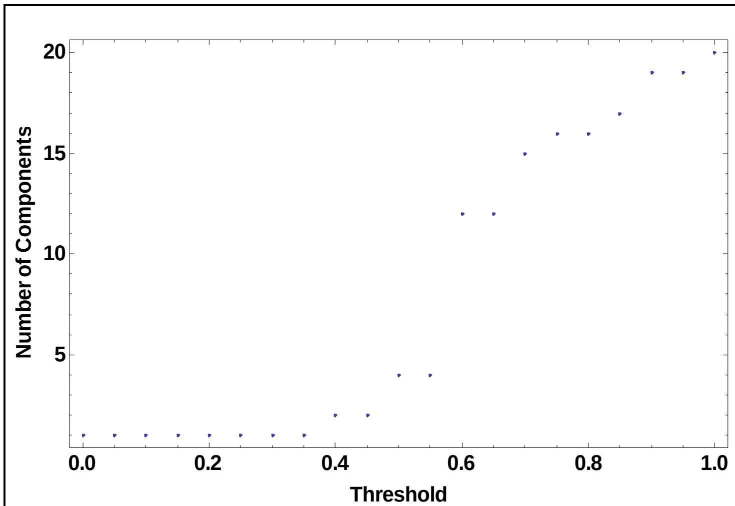

Figure 9: Number of components as function of threshold: 20 agents, 10 interest.

If the threshold is very high, there are 20 components, one for each agent. With no neighbors to reinforce its interests, each agent will begin to forget them, and the agents will independently approach the fixed point of an all-zero interest string.

If the threshold is very low, all agents will form one large group, and converge as in Figure 4.

At intermediate thresholds above the phase shift, the agents clump into small disjoint components. For example, one run at threshold 0.7 yielded two groups of size 3, three of size two, and eight of size one. Each of these groups evolves independently, yielding high diversity among groups (population diameter 1) but collapse within groups (min-max ratio of 0). This model corresponds to the biological concept of allopatric speciation, in which physical separation allows groups to evolve separately.

For intermediate thresholds below the phase shift, the agents form a number of neighborhoods, but some agents ("bridging agents") belong to more than one neighborhood. Figure 10 is a graph of one such case with threshold 0.5, with an edge between two agents if the similarity between those agents is greater than the threshold. Because neighborhoods are fixed over the run, each neighborhood can converge relatively independently of the others, but the bridging agents (in this case, for example, agent 20) repeatedly displace each neighborhood's equilibrium with the emerging equilibrium of another group. Convergence within local neighborhoods provides the source of diversity that, mediated by

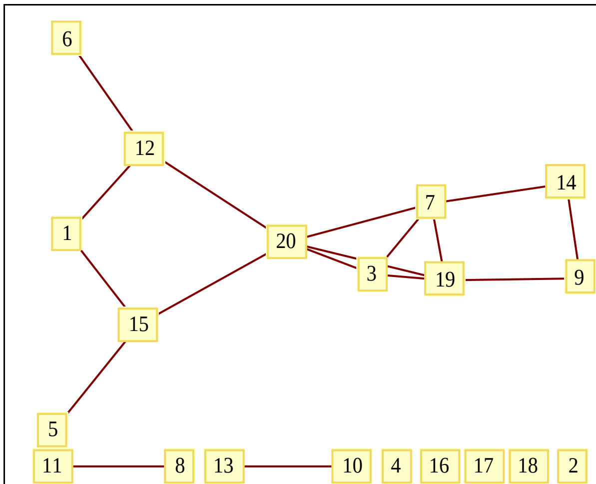

Figure 10 Neighborhood relations for threshold of 0.5

bridging agents, keeps nearby neighborhoods from collapsing. Page [18] discusses the potential for such dynamics, and the model of Bednar et al. [3] can be aligned with this result by drawing on their observation that the pressure to internal consistency for a single agent is formally equivalent to the pressure to conformity among a group of agents.

The result of this interplay of separate but linked groups is convergence without collapse (Figure 11). This mechanism, like curmudgeons and unlike mutation, provides robustness against intermittent collapse. It reflects a community with subdisciplines, but subdisciplines that recognize the value of members who bridge with other subdisciplines and exchange ideas between them. Such members are likely to be tolerated better by subgroups than would curmudgeons, because the source of the variation introduced by the bridging individuals is perceived as resulting from their multidisciplinary orientation rather than their orneriness. The population diameter under fixed groups tends to 1.

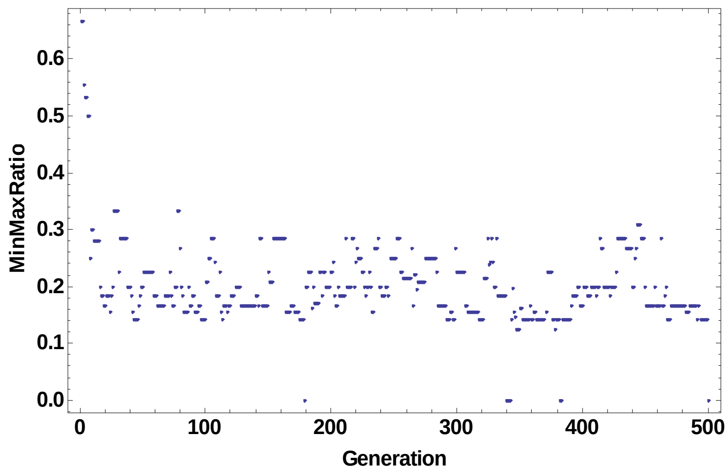

Figure 11 Fixed neighborhoods induced by threshold 0.5

This last mechanism is related to Sunstein's observation that polarization is more likely if people feel strong solidarity with their group. By definition, bridging individuals are part of multiple groups. They are inherently less completely identified with any one group, and thus unlikely to be drawn completely into the group consensus. As a result, they can keep the group leavened with new ideas, protecting against collapse.

5. DIRECTIONS FOR FUTURE WORK

Our simple model has shown a surprisingly rich space of behaviors. A number of directions for further work suggest themselves. For example:

- An analytical model of C[3], along the lines of our previous work on convergence of multi-agent systems [20], would be valuable for suggesting additional mechanisms for monitoring and avoiding collapse. Existing work on the mathematics of biological speciation offers a promising foundation for this analysis.
- How can convergence be monitored in practice? Our metric, while effective for simulation, is impractical for monitoring actual groups of people. Explicit questionnaires [23] are appropriate for experimental setting but cumbersome in monitoring groups "in the wild." One might monitor the amount of jargon that a group uses, or lack of innovation, as indicators of convergence. A promising example of initial work in this area is Schemer [4].
- We have suggested that convergence is a two-edged sword. What is the ideal degree of convergence, to allow the production of specialist knowledge without compromising the ability to escape collapse?
- How does convergence vary with group size? Recent work [19] suggests that convergence in small groups requires specialized knowledge, while convergence in large groups requires a general knowledge base.
- We have assumed homogeneous tendencies to learn, forget, mutate, or behave curmudgeonly over all agents. How does the system respond if agents vary on these parameters? In particular, what is the impact of these parameters for bridging individuals in comparison with non-bridging individuals?

6. CONCLUSION

It is natural for groups of people to converge cognitively. This convergence facilitates mutual understanding and coordination, but if left unchecked can lead the group to collapse cognitively, becoming blind to viewpoints other than their own. Experiments with a simple agent-based model of this phenomenon show that seemingly obvious mechanisms do not check this tendency. In the domain of academic conferences, these well-intended mechanisms include plenary sessions, special tracks, or even random mixing. A source of variation must be introduced to counteract the natural tendency to converge. Mutation is effective if just the right amount is applied, but tends to let the system intermittently collapse. Curmudgeons are more robust, but socially distasteful. Perhaps the most desirable mechanism consists of bridge individuals who provide interaction between individually converging subpopulations. These individuals arise when groups are well-defined, but have thresholds for participation low enough that some individuals can participate in multiple groups.

Insights from this simple model can give guidance in monitoring and managing collaboration. Here are two examples.

Our first example consists of a team of analysts searching for information. In this situation, management may have considerable influence in forming the team, and the actions available for managing C[3] reflect this influence. For example,

- If a group's searches are sparsely distributed in search space, guide more analysts to join this group to cover more areas in this search space.
- If a group's searches are not specific enough, promote the splitting of groups to create smaller, specialist groups (for example, by introducing specialists).
- If a certain convergence threshold is reached (perhaps because the search space has been exhausted), introduce a curmudgeon to guide the group into a new area of the search space.

- If in a group only a few individuals drive convergence, encourage less active individuals to participate more.
- If in a group the majority of people prevent the exploration of novel areas in search space, artificially encourage these people to be more adventurous.

As another example, consider the problem of academic overspecialization that prompted our research in the first place. The association of researchers into subdisciplines is much less amenable to centralized intervention.

Our results suggest that topical conference tracks can contribute to collapse. The narrow focus of such tracks is enhanced by selecting reviewers for each paper who are experts in the domain of the paper. Papers must be well aligned with the subdiscipline to rank high with such experts, and bridging papers are at a disadvantage. One might envision requiring one reviewer for each paper to be a senior researcher (thus capable of discerning high quality in problem formulation and execution) but *not* a member of the paper's main topic (and thus less disposed to exclude papers that cross disciplinary boundaries). Such a scheme might encourage the acceptance of quality papers that would otherwise fall in the cracks between subspecialties, and the presence of these papers in topically-organized conference tracks would then provide the bridging function that proved so successful in avoiding collapse in our experiments.

7. REFERENCES

[1] M. Andersson. *Sexual Selection*. Princeton, NJ, Princeton University Press, 1994.

[2] R. Axelrod. The Dissemination of Culture: A Model with Local Convergence and Global Polarization. *Journal of Conflict Resolution*, 41(2 (April)):203-226, 1997. http://www-personal.umich.edu/~axe/research/Dissemination.pdf.

[3] J. Bednar, S. Page, A. Bramson, and A. Jones-Rooy. Conformity, Consistency, and Cultural Heterogeneity. In *Proceedings of Annual Meeting of the American Political Science Association*, Marriott, Loews Philadelphia, and the Pennsylvania Convention Center, Philadelphia, PA, American Political Science Association, 2006. http://www.allacademic.com/meta/p150817_index.html.

[4] C. Behrens, H.-S. Shim, and D. Bassu. Schemer: Consensus-Based Knowledge Validation and Collaboration Services for Virtual Teams of Intelligence Experts. In R. L. Popp and J. Yen, Editors, *Emergent information technologies and enabling policies for counter-terrorism*, pages 209-230. John Wiley & Sons, Hoboken, NJ, 2006.

[5] J. A. Coyne and H. A. Orr. *Speciation*. Sunderland, MA, Sinauer Associates, 2004.

[6] G. Deffuant, D. Neau, F. Amblard, and G. Weisbuch. Mixing beliefs among interacting agents. *Advances in Complex Systems*, 3(1-4):87-98, 2000. http://www.lps.ens.fr/~weisbuch/mixbel.ps.

[7] P. Erdös and A. Rényi. On the Evolution of Random Graphs. *Magyar Tud. Akad. Mat. Kutató Int. Közl.*, 5:17-61, 1960.

[8] R. A. Fisher. *The Genetical Theory of Natural Selection*. Oxford, UK, Clarendon Press, 1930.

[9] A. Flache and M. W. Macy. What sustains cultural diversity and what undermines it? Axelrod and beyond. University of Groningen, Groningen, the Netherlands, 2006. http://arxiv.org/abs/physics/0604201.

[10] J. R. P. French. A formal theory of social power. *Psychological Review*, 63:181–194, 1956.

[11] D. J. Futuyma. *Evolutionary Biology*. 3 ed. Sundarlund, MA, Sinauer, 1998.

[12] D. L. Hartl and A. G. Clark. *Principles of Population Genetics*. 2 ed. Sunderland, MA, Sinauer, 1989.

[13] R. Hegselmann and U. Krause. Opinion Dynamics and Bounded Confidence Models, Analysis, and Simulation. *Journal of Artifical Societies and Social Simulation*, 5(3), 2002. http://jasss.soc.surrey.ac.uk/5/3/2.html.

[14] P. G. Higgs and B. Derrida. Stochastic models for species formation in evolving populations. *J. Phys.* , A24:L985-L991 1991. http://www.lps.ens.fr/~derrida/PAPIERS/1991/higgs-speciation.pdf.

[15] R. Hilscher. Agent-based models of competitive specitationI: effects of mate search tactics and ecological conditions. *Evolutionary Ecology Research*, 7(7):943-971, 2005.

[16] U. Krause. A Discrete Nonlinear and Non-autonomous Model of Consensus Formation. In S. Elaydi, G. Ladas, J. Popenda, and J. Rakowski, Editors, *Communications in Difference Equations*, Gordon and Breach Science Publ., 2000. http://www.informatik.uni-bremen.de/~krause/download/ConsensusFormation.pdf.

[17] G. Long, C. Yun-feng, and C. Xu. The evolution of opinions on scale-free networks. *Frontiers of Physics in China*, 1(4):506-509, 2006. http://www.springerlink.com/content/74x0w76817014272/.

[18] S. E. Page. *The Difference: How the Power of Diversity Creates Better Groups, Firms, Schools, and Societies*. Princeton, NJ, Princeton University Press, 2007.

[19] G. Palla, A.-L. Barabási, and T. Vicsek. Quantifying social group evolution. *Nature*, 446:664-667, 2006. http://arxiv.org/abs/0704.0744.

[20] H. V. D. Parunak, S. A. Brueckner, J. A. Sauter, and R. Matthews. Global Convergence of Local Agent Behaviors. In *Proceedings of Fourth International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS05)*, Utrecht, The Netherlands, pages 305-312, ACM, 2005. http://www.newvectors.net/staff/parunakv/AAMAS05Converge.pdf.

[21] P. S. Pennings, M. Kopp, G. Meszéna, U. Dieckmann, and J. Hermisson. An Analytically Tractable Model for Competitive Speciation. *The American Naturalist*, 171(1):E44-E71, 2008.

[22] W. B. Provine. *Sewall Wright and Evolutionary Biology*. Chicago, University of Chicago Press, 1986.

[23] A. K. Romney, S. C. Weller, and W. H. Batchelder. Culture as Consensus: A Theory of Culture and Informant Accuracy. *American Anthropologist*, 88(2 (June)):313-338 1986.

[24] C. R. Sunstein. The Law of Group Polarization. 91, University of Chicago Law School, Chicago, IL, 1999. http://ssrn.com/paper=199668

[25] K. Sznajd-Weron and J. O. Sznajd. Opinion Evolution in Closed Community. *International Journal of Modern Physics C*, 11(6):1157-1165, 2000. http://www.ift.uni.wroc.pl/~kweron/articles/Sznajd-WeronSznajd00_IJMPC.pdf.