Understanding 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

When a set of people interact frequently with one another, they often grow to think more and more along the same lines, a phenomenon we call "collective cognitive convergence" (C^3). We discuss instances of C^3 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.

1 Introduction

When the same people interact frequently with one another, the dynamics of the *collective* can lead to a *convergence* in *cognitive* orientation, thus "collective cognitive convergence" (C^3). C^3 is seen in many different contexts, including research subdisciplines, political and religious associations, and persistent adversarial configurations such as the cold war. Tools that support collaboration, such as blogging, wikis, and 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 may accelerate the growth of knowledge.

But *convergence* can lead to *collapse*. It reduces the diversity of concepts to which the group is exposed and thus leaves the group vulnerable to unexpected changes. For example, 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]

---

[1] This paper was motivated by frustration in the industry track at AAMAS07 that some subdisciplines of agent research were becoming so ingrown, focusing only on problems

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, self-reinforcing collapse may lead a group to focus on a cognitive construct meaningful only to the group's members. Highly specialized academic disciplines can become increasingly irrelevant to people outside of their own circle.

We became interested in this phenomenon by observing increasing balkanization in multi-agent research. Since we work in multi-agent simulation, it occurred to us that some light might be shed on the phenomenon with a multi-agent model. Many groups do not collapse, and a simulation can also help us understand how they maintain their diversity. 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 and extends previous work in sociology (both empirical and computational) and evolutionary biology.

2.1 Sociological Antecedents

There is abundant **empirical** evidence that groups of people who interact regularly with one another tend to exhibit C^3. Sunstein [27] 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.

One recent review of **computational** studies of consensus formation [14] traces relevant studies back more than 50 years [11], including both analysis and simulation. They differ in the belief model and the topology, arity, and preference of agent interactions. Rather than attempting an exhaustive review, we situate our work in these dimensions.

Belief.—An agent's belief can be either a single variable or a vector, with real, binary, or nominal values. Vector models usually represent a collection of beliefs, but

---

defined by other members of the subdiscipline, that it was difficult to apply them to real problems.

Table 1: Representative Studies in Consensus Formation

Study	Belief	Topology	Arity	Preference?
Krause [17]	Real variable	Random	Many	Yes
Sznajd-Weron [28]	Binary variable	Lattice	Two	No
Deffuant [7]	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

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.

Topology.—Some models constrain interactions by agent location in an incomplete graph, usually a lattice (though one study [18] considers scale-free networks). In others any agents can interact (often called the "random choice" model).

Arity.—Agents may interact only two at a time, or as larger groups.

Preference.—The likelihood of agent interaction may be modulated by their similarity.

Table 1 characterizes several previous 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 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

[2] [10] 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.

convergence under preferences have been obtained [17], but only for six or fewer agents [14]. 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 [15, 24]).

In a forthcoming paper [21], we have derived formal convergence results for a very simple case of our model (all agents interacting as a single group), but even this model is unable to capture the full richness of behavior that we observe. Given this research context, in this paper we focus our attention on simulation results, to develop intuitions that may reward future analytical exploration.

2.2 Biological Antecedents

The subgroups that form and cease to interact when convergence turns to collapse resemble biological species, which do not interbreed. So we look for insight to research in biological speciation (see [5, 12] for reviews). Compare Dawkins' notion of the meme in cultural evolution [6]. The most commonly proposed speciation mechanisms are allopatric, sympatric, and parapatric.

- In allopatric speciation, genetic barriers gradually evolve between two or more geographically isolated species (for instance, organisms living on separate islands). One configuration of our model can be interpreted as allopatric speciation.
- Parapatric speciation has no discrete barrier between populations; individuals are distributed along a geographic continuum and are separated by distance. Finally, in sympatric speciation a single population with no physical or geographic gene flow barriers divides into separate species.
- Sympatric speciation requires two interacting forces: 1) a force that drives sympatric speciation (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, 9] refers to the differential mating success of individuals in a population, and can be based on either 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, Fisher's runaway process, 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, 12, 16]. 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 [13]. 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 [25]. 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. Each participant's interests are 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. Similarity between two vectors is measured by the normalized Hamming similarity, which is the number of positions at which the two vectors agree, divided by the overall length of the vectors. At each step, each participant

- identifies a neighborhood of other participants based on some criteria (here, usually similarity between their interest vectors greater than a similarity threshold θ, but alternatively geographical proximity, or proximity in a social network),
- 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. Because we are working with vectors of beliefs rather than a single belief, simple summaries such as the mean and variance of a scalar (commonly used in studies of the system $\boldsymbol{x}(t+1) = \boldsymbol{A}\boldsymbol{x}(t)$) are not available to us. We have examined a number of possible measures over sets of belief vectors.

- Bednar et al. [3] deal with belief vectors, but where the elements reflect (possibly inconsistent) manifestations of the same belief under different circumstances. The attractor of interest for a single agent is thus a vector all of whose entries are the same, and for a community, concurrence across these single values.

- Some speciation models that model genomes as Boolean vectors [15] use the number of shared 1's in the vectors $\boldsymbol{x_j}, \boldsymbol{x_k}$ of two agents a_j, a_k to measure their similarity, and average this quantity over a population to estimate the closeness of the population. This measure has the property that two vectors of all zero's are maximally separate; we prefer a measure (such as the Hamming distance) that recognizes identical vectors as maximally similar, whether the agreement is in 1's or 0's. Simply averaging the similarity loses the important distinction between highly clustered agents and uniformly distributed agents.
- Sophisticated statistical techniques exist for estimating the "true" social beliefs of a population empirically, based on their (noisy) responses to questionnaires [26]. In our case, the vectors are accurate representations of the agents' beliefs.
- One interesting class of measure that we have not pursued, and that might usefully supplement our measures, is the persistence of agent associations over time [20].

In the original version of this research [22], we derived a convergence measure from the distances at which agents cluster in a hierarchical clustering. This measure is costly to compute and does not lend itself to analytical treatment. A much more satisfactory measure is the mutual information between topics and agents. The mutual information between two features (*a, b*) over a set of data is

$$MI = \sum_{a,b} p(a,b)\log\left(\frac{p(a,b)}{p(a)\,p(b)}\right)$$

In our case, *b* indexes over topics and *a* indexes over agents. $p(a,b)$ is the probability that the *a*th agent is interested in the *b*th topic, while $p(b)$ is the sum of this probability over all agents, and $p(a)$ is the sum for agent *a* over all topics. Consider two cases.

First, if all agents have the same interests, $p(a)$ is independent of $p(b)$. Then $p(a,b) = p(a)p(b)$, and the logarithm (and thus *MI*) vanishes.

Second, if each agent has a distinctive set of interests, $p(a,b)$ will differ from $p(a)p(b)$. Sometimes it will be less, sometimes more. When it is less, the logarithm will be negative, but will be weighted by a relatively small value of $p(a,b)$. When it is more, the logarithm will be positive, and will be weighted by the larger $p(a,b)$. The resulting *MI* will be greater than zero, with an upper bound equal to the lesser of $log(m)$ and $log(n)$.

Thus *MI* is an easily computed measure of the degree of diversity in an agent population. 0 indicates that all agents

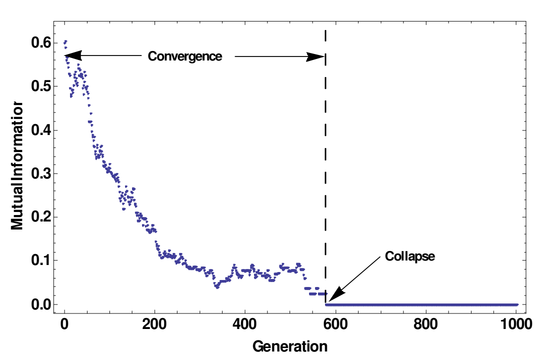

Fig. 1. Evolution of 20 agents, 10 interests, neighborhoods defined by similarity > 0.5

have identical interests, while a higher value indicates divergence.

Fig. 1 shows the behavior of the *MI* over a sample run 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 is greater than θ=0.5. All agents collapse to the same interest vector at about generation 600. 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 $p_{interest}$.).

4 Some Experiments

With this model and metric, we can explore C^3 under a variety of circumstances. As we might expect, forming groups based on similar interests leads to rapid cognitive convergence. But other sorts of neighborhoods also lead to convergence.

4.1 Things that Don't Work

We have explored several different kinds of neighborhood formation policies that also lead to convergence.

4.1.1 A Universal Group

Perhaps highly tolerant agents, who consider all agents their neighbors, might avoid convergence. Fig. 2 shows the evolution 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 model a conference with only plenary sessions. The population still collapses.

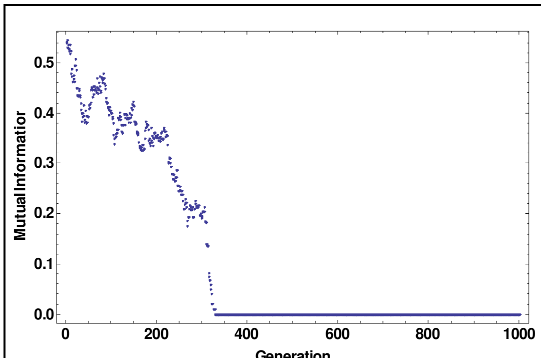

Fig. 2. θ = 0

4.1.2 Fixed Neighborhood Size

Perhaps the problem is that as agents converge, their neighborhoods in-crease in size. Fig. 3 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 on the basis of the common interests of their members. It corresponds to the biological model of 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 *MI* still goes to zero, as agents form subgroups within which interests collapse.

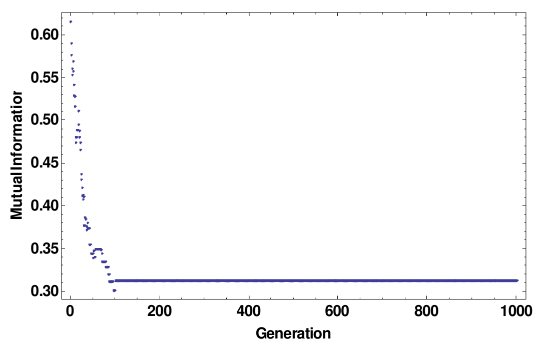

Fig. 3. Fixed-size agent neighborhoods (four closest agents)

4.1.3 Random Neighborhoods

Even more radically, let an agent's neighbors at each step be four randomly chosen agents (like a conference where papers are assigned to tracks, not by topic, but randomly). In spite of the resulting mixing, the population again collapses.

The time to convergence is highly variable, not only among these examples but within a single configuration.[4] 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 fixed interacting subpopulations.

4.2.1 Random Mutation

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. Fig. 4 shows an extended run with parameters the same as Fig. 1 (neighborhoods defined by $\theta = 0.5$), but with $p_{mutate} = 0.06$. Mutation is certainly able to reintroduce variation, but the level is critical. If mutation is too low (say, 1%), it is unable to keep

[4] This variation also shows the limits of applying a simple analytical derivation of convergence time such as that of Bednar et al. [3] to our more complex system. For the simplest case of a universal group, a forthcoming study [21] shows that convergence is proportional to $e^{-t/mn}$, where *m* is the number of topics and *n* the number of agents, but this analysis does not take into account more complex groupings.

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, but then discovers new ideas that reinvigorate it.

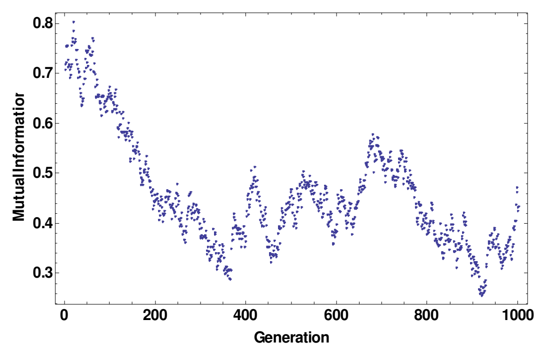

Fig. 4. Adding 6% mutation

Mutation as an abstract mechanism corresponds to several possible effects in real-world group dynamics.

- An agent might exhibit spontaneous curiosity about some topic that it has not previously thought important.
- An agent's attention might be drawn to a topic because of exogenous events such as news reports. In effect, the agent is in multiple groups concurrently, the community whose members are modeled by the C^3 agents, and a broader community (e.g., subscribers to the *New York Times*). We model this interplay of information from multiple communities explicitly later in this section.
- Yves Demazeau[5] has suggested the dynamic of a changing set of topics over time. Mutation approximates this dynamic, in the following sense. Mutation replaces a single agent's position on a topic by a randomly chosen 0 or 1. If the population adds a new topic, concurrently deleting an old one so that the total number of topics remains the same, the effect is the same as if each agent were to mutate the same position in its vector, so that agents have random values in that position. Thus topic addition can be viewed as mutation across the whole population rather than just a single agent.

4.2.3 Curmudgeons

A curmudgeon is a non-conformist, someone who regularly questions the group's norms and assumptions. Sunstein [27] 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.

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

---

[5] In very helpful discussion at the MABS workshop where this paper was presented.

Fig. 5 shows the effect when 1% of the agent decisions are reversed, again with the baseline configuration of Fig. 3. The system clearly converges, but does not collapse. Furthermore, p_{cur} can achieve this balancing effect over a much wider range than p_{mutate}. 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, as long as they are not excluded from the community.

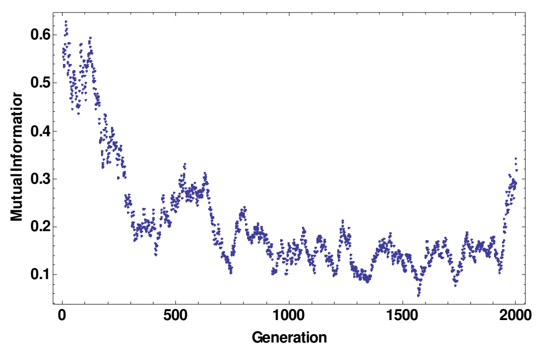

Fig. 5. 1% curmudgeons

4.2.3 Fixed Interacting Subpopulations

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

The behavior depends on the structure of the graph induced by a given threshold. Fig. 6 shows how the number of components depends on the threshold for groups formed in populations of 20 agents with 10 interests each. The shift from many components at 0.9 to a few at 0.6 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 [8], in this case as a result of lowering the threshold. Four cases merit attention.

If θ 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.

With a low θ, all agents will form one group, and converge as in Fig. 2.

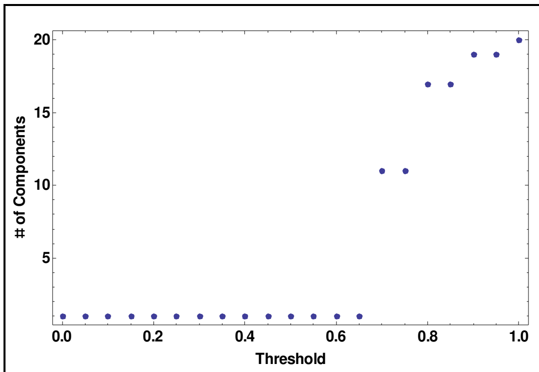

Fig. 6. Number of components as function of θ

At intermediate θ above the phase shift, each agent's set of neighbors the agents

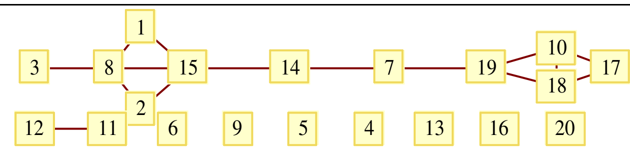

Fig. 7. Neighborhood relations for threshold of 0.7

clump into small disjoint components. Each of these groups evolves independently, yielding high diversity among groups but collapse within groups. This model corresponds to allopatric speciation, in which physical separation allows groups to evolve separately.

For intermediate θ below the phase shift, the agents form a number of neighborhoods, but some agents ("bridging agents") belong to more than one neighborhood. Fig. 7 is a graph of one such case with θ = 0.7, with an edge between two agents if the similarity between those agents is greater than θ. 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 bridging agents, keeps nearby neighborhoods from collapsing. 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.

Fig. 8 shows the behavior of this interplay of separate but linked groups. Diversity actually increases over the first 50 generations, as individual groups explore their local configuration of topics. Then, the interaction between groups leads to convergence, but the population does not collapse because each group exerts a pressure toward its own consensus. This mechanism, like curmudgeons and unlike mutation, is robust against intermittent collapse. It reflects a community whose subdisciplines 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.

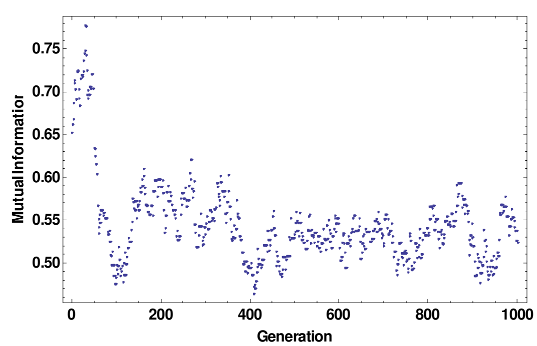

Fig. 8. Behavior of fixed neighborhoods, θ = 0.7.

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 less completely identified with a single 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

A number of directions for further work suggest themselves. For example:

- We allow agents to have only binary interests. In practice, a person's level of interest in a topic is more nuanced, and it would be interesting to explore the behavior of models with real-valued interests.
- Our topics are independent of one another. In practice, topics may have an influence on one another, which might be represented by a weight between any two topics. Increase in an agent's interest in one topic would lead to an increase in interest in related topics.
- An analytical model of C^3 (compare our previous work on multi-agent convergence [23]) might suggest additional mechanisms for monitoring and avoiding collapse. Existing work on the mathematics of speciation offers a promising foundation, though our combination of vector beliefs and preferential group interaction is significantly more complex than the systems that have been analyzed previously. A forthcoming paper [21] presents a formal analysis for the case of a single interacting population ($\theta = 0$).
- How can convergence be monitored in practice? Our metric, while effective for simulation, is impractical for monitoring actual groups of people. Explicit questionnaires [26] are appropriate for experimental settings but not for monitoring groups "in the wild." One might monitor the amount of jargon that a group uses, or lack of innovation. A promising example of initial work in this area is Schemer [4].
- Convergence is a two-edged sword. What is degree of convergence allows the production of specialist knowledge without risking 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

Groups of people naturally converge cognitively. This convergence facilitates mutual understanding and coordination, but if left unchecked can lead to cognitive collapse, blinding the group to other viewpoints.

Experiments with a simple agent-based model show that seemingly obvious mechanisms do not check this tendency. In the domain of academic conferences, these mechanisms represent plenary sessions, special tracks, and 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 help monitor and manage collaboration. For example, consider the problem of academic overspecialization. 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. What if one reviewer for each paper were 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 avoids collapse in our experiments.

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] R. Dawkins. *The Selfish Gene*. Oxford University Press, 1976.

[7] 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.

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

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

[10] 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.

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

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

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

[14] 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.

[15] 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.

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

[17] 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.

[18] 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/.

[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] G. Palla, A.-L. Barabási, and T. Vicsek. Quantifying social group evolution. *Nature* 446(7136):664-667, 2007. http://www.nd.edu/~alb/Publication06/142-SocialGroup_Nature-5Apr07-CN/SocialGroup_Nature-5Apr07.pdf.

[21] H. V. D. Parunak. A Mathematical Analysis of Collective Cognitive Convergence. In *Proceedings of the Eighth International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS09)*, Budapest, Hungary, pages (submitted), 2009. http://www.newvectors.net/staff/parunakv/AAMAS09CCC.pdf.

[22] H. V. D. Parunak, T. C. Belding, R. Hilscher, and S. Brueckner. Modeling and Managing Collective Cognitive Convergence. In *Proceedings of The Seventh International Conference on Autonomous Agents and Multi-Agent Systems*, Estoril, Portugal, pages 1505-1508, International Foundation for Autonomous Agents and Multi-Agent Systems, 2008. http://www.newvectors.net/staff/parunakv/AAMAS08M2C3.pdf.

[23] 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.

[24] 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.

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

[26] 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.

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

[28] 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.