Submitted to the CODYM Workshop, ECCS 2011.

Advances in Complex Systems

OPINION DYNAMICS WITH SOCIAL CONSTRAINTS AND EXOGENOUS DRIVERS

H. VAN DYKE PARUNAK

SVEN A. BRUECKNER

ELIZABETH DOWNS

ANDREW YINGER

Vector Research Center, Jacobs Engineering Group, 3520 Green Court, Suite 250, Ann Arbor, MI 48105 USA

{*van.parunak, sven.brueckner, liz.downs, andrew.yinger*}*@jacobs.com*

Received (received date)

Revised (revised date)

A number of studies have explored the dynamics of opinion change among interacting knowledge workers. We are particularly interested in the transition from cognitive convergence (a positive group phenomenon) to collapse (which can lead to overlooking critical information). This paper extends previous agent-based studies of this subject in two directions. First, we allow agents to belong to distinct social groups and explore the effect of varying degrees of within-group affinity. Second, we provide exogenous drivers of agent opinion in the form of a dynamic set of documents that they may query. We exhibit a metastable configuration of this system with three distinct phases, and develop an operational metric for distinguishing convergence from collapse in the final phase. Then we use this metric to explore the system's dynamics, over the space defined by social affinity and precision of queries against documents, and under a range of different functions for the influence that an interaction partner has on an agent.

Keywords: opinion dynamics; cognitive convergence; cognitive collapse; agent-based modeling; emergent behavior.

1. Introduction

Humans form opinions and interests by interactions with their peers (collaboration) and with information sources (search). The pattern of these interactions emerges from the structure of a person's environment (social network, access to information sources) and the person's current opinions. In turn, the evolution of a person's interest may shape the environment, resulting in a complex feedback process. As a result of such dynamics, collective cognitive effects may emerge at the system level (across groups of people) that can dominate the individuals' opinion evolution without the person's being aware of them. One common phenomenon is alignment of opinions, a process that is sometimes called "consensus formation" [8] or "collective cognitive convergence" [15, 14]. This phenomenon contributes to the power of collaborative groups, but it poses a threat if a group's convergence turns into collapse, blinding it to new ideas or data contrary to its current opinion.

Coordination of autonomous actors by feedback through a shared environment is called "stigmergy," and often yields emergence of global features [17]. While stigmergy as a "design pattern" is best known in systems such as social insect colonies, there are many examples of humans coordinating through simple interactions in a shared environment to accomplish a common goal [13]. One individual may select an interaction partner (person or information source) based on her current opinions. The interaction may change not only her opinions, but also the social network relations, affecting subsequent interaction decisions by others.

We construct models to learn how to measure and modulate cognitive convergence in a group of knowledge workers. We and others have studied such dynamics using agent-based models. The models to date are driven solely by the initial opinions of the agents, and social connections, if represented at all, form a connected graph among all agents. In many government and business settings, a population of analysts formulates recommendations for policy makers. Analysts not only talk to one another, but also consult exogenous information, in the form of a dynamic collection of documents. Furthermore, the analyst population is divided into separate communities, within which analysts interact preferentially. Each community starts with a tasking, a description of the subject that they are to explore. Exploring the dynamics of such a system requires two extensions that go beyond previous work by ourselves and others: interaction of disjoint social groups, and the influence of exogenous information.

Section 2 surveys previous work on opinion dynamics, and highlights the new contributions of our research. Section 3 outlines the structure of our model and formal measures we use to observe its behavior. Section 4 reports experiments with the model over the space defined by varying levels of group affinity and varying precision in retrieval of exogenous information sources, leading to two suggestions for modulating convergence among knowledge workers. Section 5 concludes.

2. Previous Research on Opinion Dynamics

One recent review of computational studies of consensus formation [7] traces relevant studies back more than 50 years [6], including both analysis and simulation. These studies differ in the belief model and the topology, arity, and preference of agent interactions.

- An agent's **belief** can be either a single variable or a vector, with real, binary, or nominal values. Vector models can be either independent, in which an agent can hold any combination of beliefs concurrently, or correlated, in which there is pressure for consistency among an agent's beliefs.
- Different **topologies** can constrain interactions. Some models constrain interactions by agent location in an incomplete graph, usually a lattice (though one study [10] considers scale-free networks). In others any agents

Table 1. Representative studies in opinion dynamics

Study	Beliefs	Topology	Arity	Bias?
Krause [8]	Real variable	Choice	Many	Yes
Sznajd-Weron [18]	Binary variable	Lattice	Two	No
Malinchik [11]	Real variable	Lattice, Random, or Hierarchy	Two	No
Deffuant [5]	Real variable	Choice	Two	Yes
	Binary vector, independent	Choice	Two	Yes
Axelrod [1]	Nominal vector, independent	Lattice	Two	Yes
Bednar [2]	Nominal vector, correlated	Choice	Many	No
Lakkaraju [9]	Real vector, correlated	Complete, Lattice, Regular, Small-world	Two	No
Parunak [14]	Binary vector	Choice	Many	Yes
This paper	Real vector, independent	Arbitrary, Unconnected	Two	Yes

can interact (the “choice” model).

- Interaction **arity** determines whether agents interact only two at a time, or as larger groups.
- The likelihood of agent interaction may be modulated by their **bias** for similar agents.

Table 1 characterizes several studies in this area in terms of these dimensions. Our work extends this field in two ways. First, it supports multiple disjoint social networks. Second, it provides exogenous influences, in the form of a collection of documents that agents can query. These extension allows us to model a situation in which groups of agents are collectively analyzing information from a changing collection of information sources.

3. A Agent-Based Model

This section describes our model and the metrics we use to monitor its dynamics. A wide range of configuration parameters are available to configure the initial set-up of a scenario (discussed under “model components”) and govern the execution cycle (discussed under “model execution”). For each model component and execution step, we identify the main parameters that our model exposes.

3.1. *Model components*

Our model has five main components.

Topic Space.-Analysts and documents live in an abstract Euclidean space defined by a set of k topics. In the real world, a topic is a probability distribution over lexicographic terms (e.g., domain-relevant key words), constructed from a large collection of relevant documents using techniques such as Latent Semantic Analysis (LSA) [4] or Latent Dirichlet Allocation (LDA) [3]. In our model, topics have no semantics. The topic space is a hypercube of dimensionality k, with a range of [0, 1] on each dimension. A location in this space is a Topic Model Vector (TMV). A theme is a region in topic space. We generate analysts or documents associated with a theme by sampling a Gaussian with configurable mean and variance, discarding locations outside of [0, 1] and resampling. Relevant parameters are:

- Number of Topics: Dimensionality of topic space
- Theme Mean
- Theme Variance

Social Network.-We organize analysts into (static) groups. This structure models constraints that externally influence the likelihood that two analysts interact. Additional internal interaction preferences within these constraints arise from the preferential selection by analyst interest. Parameters are:

- Number of Analysts
- Number of Groups
- Group Themes

Document.-A document is a Topic Model Vector (TMV). In the real world, a document's TMV is discovered using topic modeling. Real-world document repositories typically contain documents from different sources and with different concerns. We model this clumping of documents with themes, and generate a population of documents by sampling from several themes with specified means and variances. Real-world document repositories are not static, but continually grow as new documents are discovered. We model the arrival of new information during runtime as the delayed introduction of documents sampled from a new theme. Parameters are:

- Number of Documents
- Number of Themes
- Document Themes (means and variances)

Analysts.-An analyst's current interest is also a Topic Model Vector (TMV). A community's tasking is defined as a theme, and we generate the community working on a tasking by sampled that theme. The central object of our study is the movement of the analyst's TMV through topic space, relative to the TMVs representing documents and other analysts. Parameters are:

- Tasking Themes
- Number of Analysts

Document Search.-Real-world analysts use search engines to select documents for review. Depending on the search engine, queries may have different representations (e.g., key words, forms, relation graphs), but they always define topics of interest. In our model, an analyst poses a query as a subset of topics. The search weights documents by the strength of their entries on those topics, and probabilistically selects and returns a single document. Noise in this process models a real-world analyst's willingness to review documents that were not ranked first in their search results. Parameters are:

- Query Temperature
- Document Selection Temperature

3.2. *Model execution*

First we configure a scenario. Then analysts repeatedly execute four steps: choose interaction type, assemble interaction options, select interaction target, and execute interaction. Analysts execute in random order with replacement.

Configuration.-We instantiate a topic space with a specified number of dimensions (10 in the experiments reported here), then a specified number of analysts in a specified number of groups, each with a tasking theme with specified mean and variance, and finally a specified number of documents from a specified number of themes, each with specified mean and variance.

Choose Interaction Type.-The analyst chooses probabilistically whether to interact with another analyst or with a document. On a given step, an analyst interact either with a document or with another analyst. The parameter is:

- Document Query Probability: The probability p_D that an analyst queries a document. With probability $1 - p_D$, it interacts with another analyst.

Assemble Interaction Options.-If an analyst is interacting with a document, all its possible targets for interaction in this cycle are documents currently in the document space. If the agent is interacting with another analyst, all its interaction options are analysts from one of the disjoint groups in the social network. The agent picks from its own group with probability defined by its Affinity parameter, and otherwise picks from another randomly selected group. The parameter is:

- Affinity: the probability that an analyst will choose to interact with a member of its own group rather than an analyst in another group.

Select Interaction Target.-The analyst selects one target from the options assembled in the previous step. Since both document content and analyst interest are represented as TMVs, this step is identical for documents and analysts. This step by the agent models both the analyst's decision what query to construct based on its current interest and what search result to select. Socially, it reflects the analyst's choice of issues to explore with other analysts and then the interaction partner.

A model parameter specifies the number of topics j in a query. The more topics in a query, the more specific it is. The analyst selects this number of topics, favoring topics in which its interest is high, but adding temperature noise to that selection. For zero temperature, the top-j topics are chosen. For high temperature, this choice is practically random. Parameters are:

- Search Topics: Number of topics on which to query
- Topic Selection Temperature: Amount of randomness in topic selection (via Boltzman normalization)

Then an interaction target is selected based on the query. In this "Search Execution", the TMVs of the interaction options are sorted by their values in the topics of the query: TMVs with higher values rank higher. Thus, we create a relevance ranking of the TMVs relative to the given query. Based on that ranking, but again with a temperature noise parameter, we select one TMV. Here the temperature noise models an analyst's result-selection behavior. The parameter is:

- Search Temperature: Amount of randomness in document selection (via Boltzman normalization)

Execute Interaction.-The analyst now updates its interest model by interacting with the target (document or analyst) from the previous step,. If the analyst interacts with another analyst, it samples the Learning Style parameter to determine the personality it should assume in this interaction, standard or curmudgeon. The selected personality sets the update rule for updating each topics interest level as a function of the difference in interest on that topic between the agent and the selected interaction target. Section 4.4 explores the form of this rule. In the standard personality, the agent shifts its interest level in updated topics to be closer to the interest level in the interaction target. In the curmudgeon personality, it shifts away from the other interest level. Most of our experiments are performed with the standard personality. If the interaction target is a document, then the agent always uses the standard personality. Thus, document content always draws the analyst closer to the document. Parameters are:

- Learning style: probability of being a Curmudgeon
- Analyst or Document Learning or Forgetting Rate: the amount an analyst increases a topic value (learning) or reduces it (forgetting) on interaction with an analyst or document, respectively.

3.3. *Performance metrics*

We define a set of component metrics, and a single aggregate metric.

Component Measures.-The topics in a model span a space with valid locations (TMVs) limited to [0, 1] for each topic. As analyst agents update their TMV through interactions with other analysts or documents, they move through this space. We seek metrics of analyst movement that might signal cognitive collapse. Component measures apply to an analyst's single step.

The most basic measure of analyst movement is the magnitude of a single TMV update, which is the length of the vector between the agent's prior and new location in each cycle. The step is a function of the agent's choice of the interaction mode (document or analyst), the agent's limitation of the set of possible interaction partners (In-Group/Out-of-Group for analyst interactions), the emulation of query construction (based on current interest = current location) and relevance selection, and the application of a personality in the calculation of the TMV update as a function of the agent's location and the location of the interaction target.

The length of a step in conveys the absolute magnitude of the impact a particular interaction had on the analyst's interest. It does not show the nature of the step relative to other analysts. A second measure is the distance of the analyst's location to the Center of Gravity of all analysts in the model, that is, the mean over the TMVs of all analysts regardless of group affiliation. The mean TMV may not be near any analyst. Movement of the analysts shifts the location of the mean TMV, thus successive "distance to mean TMV" measurements, unlike "step-length," are not statistically independent.

We explored other measures on the step-by-step movement of analysts, such as the length of the step vector projected onto the vector from the agent's prior location to the mean TMV, and the distance between the agent's initial and current locations. We found that the model dynamics of interest are sufficiently observable in the first two metrics defined above.

An Aggregate Measure.-The measures in the previous section concern a single step of an analyst. Initial explorations show that we also need to discover a directed walk, in which successive steps are correlated. Previously [16], we applied information-theoretic (entropy) measures to detect a directed walk, but encountered idiosyncrasies from the specific definition of the system states whose probabilities are measured in the entropy calculation. For the current research, we developed an aggregate metric that measures the "directedness" in an agent's movement through topic space without the complications of entropy calculations. The delayed step length metric adds the step vectors (delta TMV) for a single agent over the most recent n cycles (configurable, 50 in the results reported here). The vector sum of steps of a random walk is on the order of $\sqrt{n}$, while the vector sum of steps that generally point in the same direction (directed walk) tends to be on the order of n.

4. Experiments with the Model

This section walks through an example scenario, exhibits the system's metastability, derives an objective way to measure the cognitive collapse of a knowledge community, and uses this measure to explore the space defined by community affiliation and precision of interaction with exogenous information. Both of these dimensions are new to the simulation study of opinion dynamics, and our experiments explore only a small portion of the space that they define. Nevertheless, our results suggest two practical principles for managing convergence and preventing collapse.

4.1. *A representative scenario*

We consider two distinct document themes and two groups of analysts. We sample 25 documents for each theme. One theme is the tasking for all 6 analysts of the

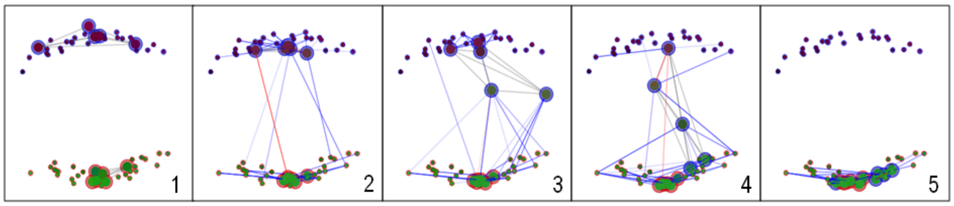

Fig. 1. Stages in the evolution of the model in the example scenario.

first group, and the other theme initializes the 5 analysts of the second group. The documents and analysts are embedded in a 10-dimensional topic space.

The screenshots in Fig. 1 illustrate the evolution of the model. As the simulation runs, we visualize the recent interactions of agents with other agents (red lines) and documents (blue lines) fading away into history (line transparency).

(1) The agents (large circles) and documents (small circles) are separated into two groups, reflecting the two document themes and analyst taskings. Lines among agents reflect their social groups.
(2) The agents in each group, initially spread out in their tasking, converge and thus form tighter clusters. Most interactions occur within the tightly clustered groups and their surrounding theme.
(3) Interactions of low-affinity analysts with the other group eventually lead to the defection of two analysts from the smaller (upper) group and their movement toward the larger group.
(4) Once the first analyst defectsp, others follow rapidly. Eventually, all analysts abandon the upper document theme.
(5) Both groups of analysts have converged on the same set of interests exemplified by the lower document set. Interactions among analysts within a group are no different from Out-of-Group interactions. Interactions with documents are (mostly) confined to the lower theme.

4.2. *A metastable transition*

Our metrics reveal three distinct phases of interest evolution dynamics (Fig. 2, metrics applied to a single analyst from the upper, smaller group). The agents from the smaller (upper) group with lower affinity first remain in their separate interest area (Phase 1), eventually defect one-by-one to the interest of the other group (Phase 2), and then explore the other interest area jointly with the agents from the larger group (Phase 3). The three distinct phases in the model dynamics are reflected in the step length and distance to mean TMV metrics.

In Phase 1. all analysts are far from the center of gravity and long steps are fairly frequent. The large distance to the mean TMV reflects the initial separation of analysts' interests into two groups, so the center of gravity of all analysts is

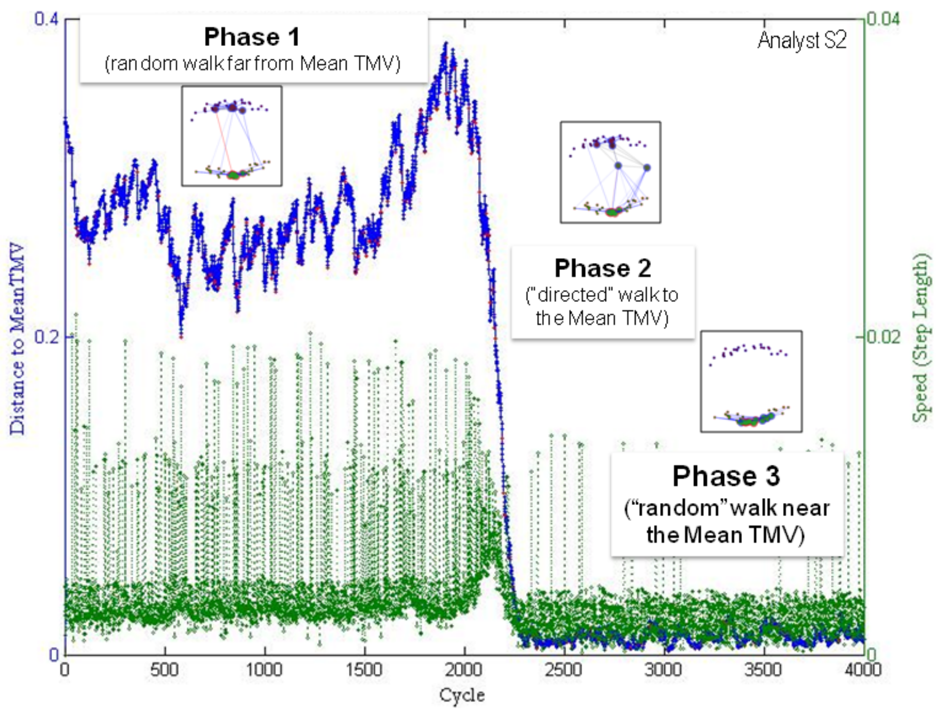

Fig. 2. The dynamics of a single analyst from the top group. Both step length (green) and distance to mean TMV (blue) display three distinct phases.

far from any individual analyst. Individual steps are of three types. Short steps (most frequent) are interactions with other analysts from the same group (and thus similar interest) or documents near the analyst's initial tasking. Medium length steps are interactions with documents from the other theme. We set the Analyst or Document Learning or Forgetting Rates so that document interactions have less impact on analyst interest than analyst-to-analyst interactions. The longest steps, and the least frequent, are interactions with analysts from the other group. In Phase 1, the roughly constant distance to the mean TMV shows that the agent's successive steps with respect to the center of gravity are random (not correlated).

Phase 2 corresponds to the agent's defection from its original tasking to the larger group. The distance to the mean TMV rapidly shrinks. The step-length metric shows an increase in the magnitude in the high-frequency interactions as the agent moves away from the documents and analysts in its own group, due to the TMV update rule, which computes larger changes for larger differences between the analyst's TMV and the TMV of its interaction partner. At the same time, the magnitude of the lower-frequency steps (interactions with the other document set and out-of-group analysts) decreases, as the analyst moves closer to those entities. In

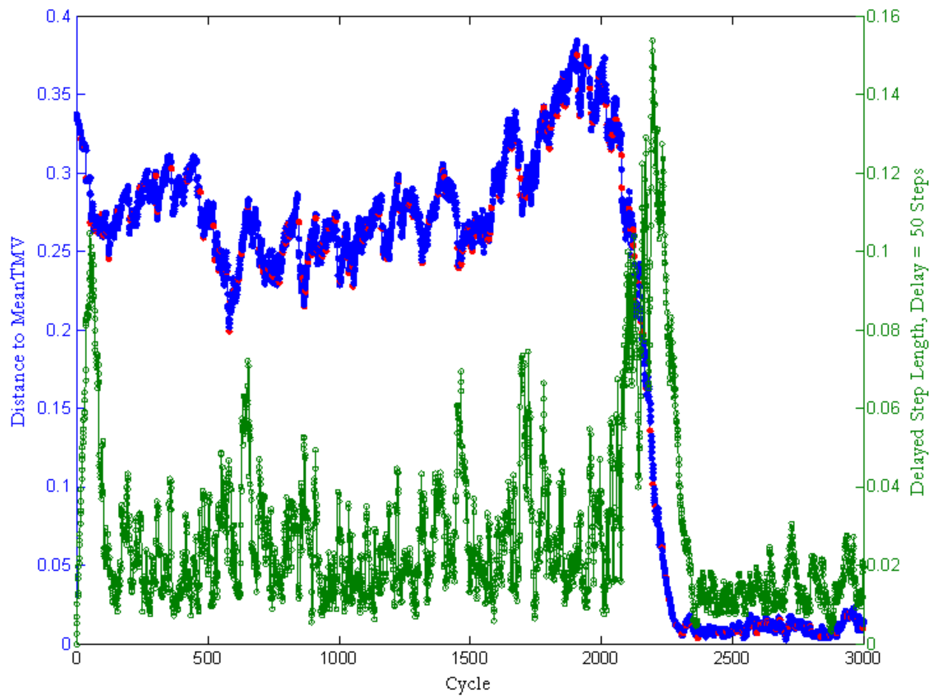

Fig. 3. Re-plotting Fig. 2 with Delayed Step Length reveals a strong peak for phase 2.

this phase, the agent's successive steps are correlated, as the rapidly falling distance to the mean TMV shows.

Phase 3 is similar to Phase 1 in that the agent's successive steps are again uncorrelated. In this phase, all analysts of both groups occupy the same region in topic space and share a common interest in documents of the second document theme. The analysts' distance to the mean TMV is small as they are now all tightly clustered, and most interactions result in only minor changes to an analyst's TMV (small step-length) as they select either nearby documents or analysts that are close by regardless whether they are in the analyst's group. The analysts still interact (infrequently) with documents from the other theme (larger steps), but those interactions have no lasting effect on the analysts' relative location to each other.

Fig. 3 shows again the distance to the mean TMV (blue) for a single analyst as in Fig. 2, but compares it with the delayed step length metric. A strong peak in the delayed step length corresponds to the rapid fall-off in distance to the mean TMV (around cycle 2200). Phase 2 in the agent dynamics is characterized by a directed walk, visible in the delayed step length metric, while Phases 1 and 3 are (generally) less directed.

The fine-structure in the metric leading up to the merger of the analyst with the other group's interest merits further investigation. We hypothesize the following. The first peak in the metric (around cycle 50) corresponds to the initial convergence of interests within the analyst's group, dispatching of the noise in the group-member's probabilistic initialization. Subsequent peaks (e.g., near cycles 600, 1500, and 1700), corresponding to significant drops in the agent's distance to the mean TMV, are failed "attempts" of the analyst to free itself from its group that are thwarted in subsequent interactions with its group members and its current document theme. Eventually, the agent succeeds in defecting to the other group.

4.3. *Defining and measuring collapse*

Thus far, we have analyzed cognitive convergence in the interest evolution of the analyst agents in the two groups in our reference model. The delayed step length metric peaks as an agent distinctly reacts to information exposed by the other group and the other document theme.

We informally defined "cognitive collapse" as the inability of an agent or a group of agents to respond to new information. We can now operationalize this definition:

An agent is in cognitive collapse if it is not in phase-2 dynamics and if it does not return to phase-2 dynamics when qualitatively new information is introduced.

We introduce new information by adding documents from a third theme to the model once when all agents show Phase-3 dynamics. This new document theme (located between the two initial themes in topic space) probes the analyst dynamics. If analysts in Phase 3 are collapsed and not just converged, they should not respond to new information by returning to Phase 2. In other words, the self-reinforcement of interests among analysts should outweigh the pull from new information.

We detect Phase-2 dynamics visually using the delayed step length metric. To facilitate automated exploration of a system's tendency to collapse, we use the nonparametric Mann-Whitney-U test [12] to compare step lengths before and after probe insertion. We explore the region of parameter space that includes query temperatures in [0.12, 0.32] and affinities in [0.4, 0.9], and run 20 replications at each point. We allow each configuration to reach Phase 3, then insert a probe, and compute the percentage of analysts who respond, indicating that they are not collapsed. Fig. 4 shows this metric over the parameter space. The configuration used in the results reported so far is at location *a*. Increasing the query temperature, and thus exposing analysts to unexpected documents, dramatically reduces collapse. Surprisingly, the likelihood of collapse does not vary systematically as we change analysts' affinity for their own group. Practically, if one wishes to modulate the rate of convergence among analysts, adding a variable quantity of noise to their queries against exogenous information appears to be more effective than motivating more or less interaction with other teams of analysts.

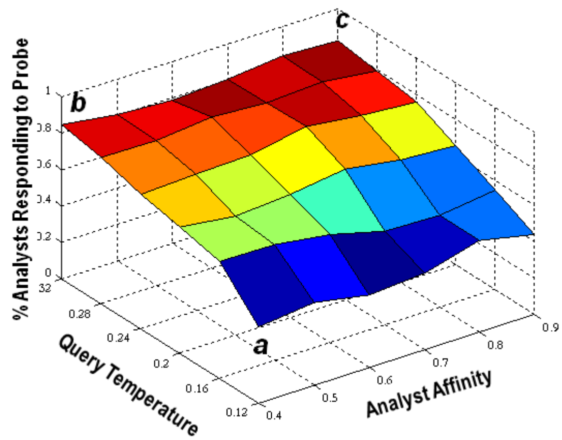

Fig. 4. Collapse as function of search temperature and analyst affinity.

4.4. *Exploring the update rule*

The TMV update rule translates the difference in TMV elements between an analyst and an interaction partner into the analyst's step length (toward the partner for an ordinary analyst, and away for a curmudgeon). In the experiments thus far, the magnitude of the change in interest in a particular topic in the TMV is proportional to the difference in that topic between the agent and the interaction target. This assumption models **curiosity**: an interest very different from mine stimulates my interest and has a larger effect on me than an interest that is similar to my own. (Parents of college freshmen often observe this principle when their children return home at their first vacation.) An alternative model is **homophily**: I am more likely to move toward ideas that are close to my own than toward those that are different. In real-world analysts, the correct model is likely to be a **mixture** of these two effects: interests too far from mine are threatening, and interests too close to mine are boring, so my response will be greatest somewhere in the middle.

We formalize these options by letting $x_i = |d_i - a_i|$, where d_i is the level of interest for the ith topic of the selected interaction partner and a_i is the level of interest for the same topic for the analyst. We can then parameterize the update rule with s as $x_i > s : (1 - x_i)/(1 - s)$ and $x_i < s : x_i/s$. The parameter s indicates the hypothesized difference that will lead to maximum movement. When $s = 0$, we recover homophily, while $s = 1$ yields curiosity.

To explore the effect of the update rule, we focus on two observable events in our model: the convergence of the interests of the two groups (a), and the response to a probe (b). We instantiate a large number of systems for each value of s and measure the probabilities $p(a)$ of convergence and $p(b)$ that the system responds to a probe (or in other words, is not collapsed). We hypothesize that as we increase s from 0

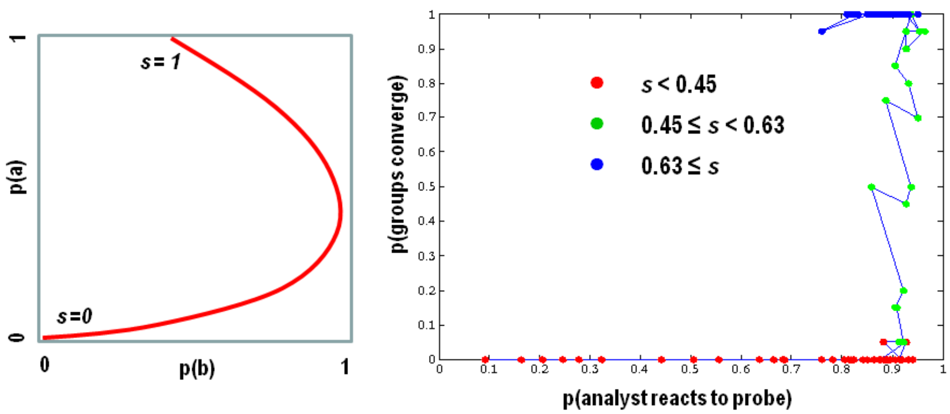

Fig. 5. Effect of update rule parameter. Left: hypothesized. Right: observed.

(homophily) to 1 (curiosity), these probabilities will vary qualitatively as sketched in Fig. 5 (left). At $s = 0$, homophily dominates and strongly limits interest changes of the analysts, so we expect convergence of the two groups to be unlikely ($p(a) \approx 0$) and the agents are likely to be collapsed within their own groups ($p(b) \approx 0$). At the other extreme ($s = 1$, high curiosity), the agents are highly mobile in topic space and thus are likely to converge on a common interest ($p(a) \approx 100\%$), but we still expect a significant risk for cognitive collapse ($p(b) < 1$) as new information may be drowned out by the wealth of stimuli already accessible to the agent. Intermediate values of s lead to intermediate levels of convergence, but with low risk of collapse ($p(b) \approx 1$).

The right side of Fig. 5 shows the results for low affinity and high query temperature (b in Fig. 4), qualitatively confirming our hypothesized curve. We again plot the average over 20 random seeds for each data point. Low s (homophily) yields no convergence between the groups but decreasing likelihood of collapse (increasing likelihood of non-collapse). The results show that group convergence has a critical threshold. Thus, instead of a gradual rise of the convergence probability as in our hypothesis, the plot remains at 0% until a critical value of s. Then the probability of all analysts converging on the same interest region rapidly increases, but the analysts remain equally open to new information (low probability of collapse). Finally, as curiosity dominates (large s, the TMV update behavior used in earlier examples), the likelihood of collapse begins to rise to the 20% level observed at Fig. 4b. The increase in risk of cognitive collapse towards the end of the sweep suggests another practical lesson for real-world knowledge workers: a mixed learning strategy that is most sensitive to information that is neither completely novel nor entirely familiar is less vulnerable to collapse than either extreme.

5. Conclusion

The opinion dynamics of knowledge workers are complex, often counter-intuitive, and yet critical for much collaborative work in the modern world, and enjoy the attention of a significant research community. Previous studies focus on the evolution of an initial distribution of opinions across agents. While suggestive, such studies do not account for two features of knowledge workers in the real world.

- Their social environment is highly clustered, and they are more likely to interact with another agent in their cluster than with an agent in another cluster.
- Their information environment includes not only analysts, but also exogenous knowledge sources (documents) that they query.

Our model implements both of these features. The resulting system exhibits an interesting metastability that allows us to formulate an operational measure of cognitive collapse. A preliminary exploration of the parameter space of social affinity and query precision with this measure yields two (very provisional) practical lessons.

(1) Query precision has much more influence on collapse than does social affinity. If one wishes to modulate the convergence of a community of analysts, managing the amount of noise added to their queries is a more promising method than changing their group membership.
(2) Motivating analysts to avoid opinions that are completely new or completely familiar leads to more robust convergence without collapse than the extremes.

6. Acknowledgment

Supported by the Intelligence Advanced Research Projects Activity (IARPA) via US Army Research Office contract number W911NF-10-C-0108. The U.S. Government is authorized to reproduce and distribute reprints for Governmental purposes notwithstanding any copyright annotation thereon. Disclaimer: The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of IARPA, ARO, or the U.S. Government.

References

[1] Axelrod, R., The dissemination of culture: A model with local convergence and global polarization, *Journal of Conflict Resolution* **41** (1997) 203–226.
[2] Bednar, J., Page, S., Bramson, A., and Jones-Rooy, A., Conformity, consistency, and cultural heterogeneity, in *Annual Meeting of the American Political Science Association* (American Political Science Association, Marriott, Loews Philadelphia, and the Pennsylvania Convention Center, Philadelphia, PA, 2006).
[3] Blei, D. M., Ng, A. Y., and Jordan, M. I., Latent dirichlet allocation, *Journal of Machine Learning Research* **3** (2003) 993–1022.
[4] Deerwester, S., Dumais, S. T., Furnas, G. W., Landauer, T. K., and Harshman, R., Indexing by latent semantic analysis, *Journal of the American Society for Information Science* **41** (1990) 391–407.

[5] Deffuant, G., Neau, D., Amblard, F., and Weisbuch, G., Mixing beliefs among interacting agents, *Advances in Complex Systems* **3** (2000) 87–98.
[6] French, J., A formal theory of social power, *Psychological Review* **63** (1956) 181–194.
[7] Hegselmann, R. and Krause, U., Opinion dynamics and bounded confidence models, analysis, and simulation, *Journal of Artifical Societies and Social Simulation* **5** (2002) none.
[8] Krause, U., A discrete nonlinear and non-autonomous model of consensus formation, in *Communications in Difference Equations*, eds. Elaydi, S., Ladas, G., Popenda, J., and Rakowski, J. (Gordon and Breach Science Publ., 2000).
[9] Lakkaraju, K. and Speed, A., A cognitive-consistency based model of population wide attitude change (2010).
[10] Long, G., Yun-feng, C., and Xu, C., The evolution of opinions on scale-free networks, *Frontiers of Physics in China* **1** (2006) 506–509.
[11] Malinchik, S. and Rosenbluth, D., Paradoxical dynamics of population opinion in response to influence of moderate leader (2011).
[12] Mosteller, F. and Rourke, R. E., *Sturdy Statistics: Nonparametrics and Order Statistics* (Addison-Wesley, Reading, MA, 1973).
[13] Parunak, H. V. D., A survey of environments and mechanisms for human-human stigmergy, in *Proceedings of E4MAS 2005*, eds. Weyns, D., Michel, F., and Parunak, H. V. D., *Lecture Notes on AI*, Vol. LNAI 3830 (Springer, 2006), pp. 163–186.
[14] Parunak, H. V. D., A mathematical analysis of collective cognitive convergence, in *the Eighth International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS09)*, eds. Decker, K., Sichman, J., Sierra, C., and Castelfranchi, C. (Budapest, Hungary, 2009), pp. 473–480.
[15] Parunak, H. V. D., Belding, T., Hilscher, R., and Brueckner, S., Understanding collective cognitive convergence, in *Multi-Agent-Based Simulation IX: International Workshop, MABS 2008, Estoril, Portugal, May 2008, Selected Papers*, eds. David, N. and Sichman, J. S., Vol. LNAI 5269 (Springer, Berlin, 2009), pp. 46–59.
[16] Parunak, H. V. D. and Brueckner, S., Entropy and self-organization in multi-agent systems, in *The Fifth International Conference on Autonomous Agents (Agents 2001)* (ACM, Montreal, Canada, 2001), pp. 124–130.
[17] Parunak, H. V. D. and Brueckner, S. A., Engineering swarming systems, in *Methodologies and Software Engineering for Agent Systems*, eds. Bergenti, F., Gleizes, M.-P., and Zambonelli, F. (Kluwer, 2004), pp. 341–376.
[18] Sznajd-Weron, K. and Sznajd, J. O., Opinion evolution in closed community, *International Journal of Modern Physics C* **11** (2000) 1157–1165.