Attitudinal Polarization on Social Networks: A Cognitive Architecture Perspective

Mark Orr (mo6xj@virginia.edu)
Biocomplexity Institute, University of Virginia,
P.0. Box 400298, Charlottesville, VA 22904 USA

Andrea Stocco (stocco@uw.edu)
Department of Psychology, University of Washington
Campus Box 3515525, Seattle, WA 98195 USA

Christian Lebiere (cl@cmu.edu)
Don Morrison (dfm2@cmu.edu)
Michael Martin(mkm@cmu.edu)
Department of Psychology, Carnegie Mellon University
5000 Forbes Avenue, Pittsburgh, PA 15213

Abstract

Polarization of attitudes is an important, and often troubling or disruptive, effect of interest in many fields. We seek to shed some light on how polarization arises by applying cognitive architectures to the problem. We created a novel embedding of individual cognitive agents, using ACT-R's declarative memory model, into social networks, simulated them communicating over time, and observed the evolution of the agents' attitudes, both collectively and individually. The primary measures we use are both Shannon entropies, of the distribution of attitudes in the final configuration of the whole social network, and of the distributions of memory traces in the individual agents at the end of the simulation. Simulations were run over ten different network topologies, using three different distributions of narrative valences, and five different values of the agents' memory decay parameter. These simulations demonstrated that polarization can be understood from a social and cognitive perspective simultaneously, each providing insights into the system's behavior.

Keywords: Cognitive Modeling; Long-Term Memory; Resting-state fMRI; Functional Network, Attitudes, Polarization, Social Networks.

Introduction

Attitudinal polarization has a long history in political science, sociology, and social psychology. It is no less relevant today than it was 50 years ago. A seemingly obvious scientific question is to ask to what extent we can understand attitudinal polarization from the perspective of cognitive architectures. This question has interest beyond the purely scientific. Understanding the structure of attitudes (as relations among beliefs) and the dynamics of attitude change can yield actionable insight for applications in, for example, public health (Orr, Thrush & Plaut, 2013). Yet, our understanding of attitudinal polarization, from a cognitive perspective, relies nearly exclusively on work in social psychology, a discipline with little intersection with cognitive architectures.

Polarization in attitudes, typically, is a valenced affair, in which an object of contention is evaluated with respect to its goodness/badness, desirability/repulsiveness, approachability/avoidability. It is typically described as a distribution across individuals (humans, bots, agents, or even models). It is also naturally described as a networked phenomenon, where clusters of individuals can develop solidarity or polarization or other variants on a theme. Important questions include: What aspects of cognitive functioning are implicated by polarization? Do polarized minds lead to polarized social spaces (or vice versa)? Are there interesting threshold effects or other non-linear relations between mental and social scales?

If the answers to these kinds of questions seem obvious to you, then consider this. The famous Shelling segregation model (Shelling, 1971) provided somewhat shocking insight into the relation between mental states and social structure -- low degrees of individual-level preferences for segregation generated strong system-level segregation. We use this example to illustrate that the relation between levels of scale is not obvious. It must be investigated rigorously. Using the perspective of cognitive architectures, as a computational, mechanistic lens, should yield a set of novel insights into polarization and other social phenomena of interest to those working on public health, security, human rights and environmental issues.

The goal of this paper is to describe an approach for studying attitudinal polarization using cognitive architectures and to show its potential value. We do this in a stylized way, with an abstract social space and the co-opting, in a highly formal way, of a specification of attitudes from social psychology. The central question we pose, but do not yet answer, is this: Can we describe the conditions, initial or otherwise, of the mental and social systems that guarantee stability in both (or either) the mental or social systems. Stability is well-understood in real-valued or binary networked systems (e.g., Bhat and Redner, 2019). But these networks are a poor abstraction for human cognitive complexity and their organization in social structures. What about socio-cognitive systems?

Toward this end, we provide a study of the relation among the distribution of attitudes and beliefs in a population and the social network structure of that population with respect to two outcomes, one in terms of external behavior and the other in terms of internal mental representation. The former is derived from the distribution of beliefs in the population and the latter is derived from the distribution of beliefs within individuals. Thus, we capture both the mental and the social in equal, symmetrical measures. In the results section, we will tie our work to future directions in the contexts of cognitive architectures, social psychology and sociology.

Simulating Social Networks

The spread of information across social networks is difficult to study experimentally. For this reason, researchers frequently make use of either large-scale, quasi-experimental data, such as analysis of large corpora of Twitter messages, or multi-agent simulations. Such techniques also are routinely used by social media companies.

In social network simulations, agents are modeled as nodes in a network whose edges are communication channels. Agents exchange information across these channels. The spreading of information is then studied as a function of factors such as network geometry (e.g., small world networks), agent goals (e.g., reaching consensus), and communication intent.

Simulating Plausible Cognitive-Social Agents

To reduce the complexity of the simulations, most computational social science efforts use relatively simple agents, often with little or no cognitive ability. This is sufficient to capture some network-level dynamics, such as those that lead to consensus within a group (Romero & Lebiere, 2014) or the production of original ideas in science. When the goal is to understand the interplay between social interactions (at the network level) and psychological constructs (at the agent level), it is warranted to imbue the agents with cognitively plausible assumptions about their thought processes. For example, Lindstrom et al. (2019) augmented social agents with reinforcement learning capabilities to successfully capture the addictive qualities of social media behavior.

Because we are interested in the interaction between network dynamics and internal beliefs, we endowed our agents with a realistic model of declarative memory. Specifically, we used Anderson's model of memory, reflecting the rational analysis of our environment reflected in our memory mechanisms (Anderson & Schooler, 1991). Those regularities, such as the power law of practice and forgetting, have also been observed in recently developed information environments such as social networks (e.g., Hubermann et al, 1998; Stanley & Byrne, 2016). In this model, the availability of memory m is related to its base-level activation function $B(m)$. Every time m is retrieved or re-encoded, a new trace is created. The final activation of m is the sum of the decaying activations of all its traces, with stochastic noise added to make the retrieval process stochastic:

$$B(m) = \log \sum_i t_i^{-d}$$

where t_i is the time elapsed since the creation of the i-th trace and d is the characteristic decay rate of an agent memory.

Connecting Memory and Social Behavior

In addition to receiving and internalizing information, an agent in a social network also sends messages across the network. The choice of which messages to spread is, ultimately, a problem of decision-making (Hackel et al., 2020.). To connect an agent's decisions to its memory, we used Gonzalez et al.'s (2003) instance-based learning framework (IBL). In IBL, agents select their next action by generating expectations reflecting previous experiences in memory that match the current context. This framework is particularly appealing because it meshes well with the ACT-R declarative model and has a long history of successful applications in decision-making (e.g., Erev et al, 2010). Furthermore, while rooted in declarative memory models, IBL gives predictions that are largely consistent with reinforcement learning (Chelian et al., 2015), another framework that has been successfully applied to social networks.

In IBL previous memories are aggregated through a mechanism called blending, which combines different outcomes in a weighted average, based on the probability P_i of retrieval of each memory reflecting its activation and similarities between the contents of the memories:

$$V = argmin \sum P_i \cdot \left(1 - Sim(V_i, v_i)\right)^2$$

Our model represented narratives and their associated valence (defined in a [-1, 1] interval) as chunks in memory. At each iteration, the model representing a node in the network would store in memory all the narratives received from its neighbors. It would then compute the node's attitude by performing a blended retrieval over all narrative chunks in memory, extracting a consensus valence. The model would then generate a narrative to spread to its neighbors by matching the node's attitude against the valence of the narrative chunks in memory. The resulting output reflects a combination of attitude of the node and popularity of narratives in its ego network.

Information Entropy as a Common Measure of Cognitive and Behavioral Dynamics

Because our goal is to measure changes in social behavior and in agent cognition at the same time, it is useful to have a common metric that applies to both.

To do so, we used Shannon's information entropy H (Shannon, 1948), which can be defined over the set of beliefs S (pairs of narrative and valence):

$$H = -\sum_{i \in S} P(i) \log P(i)$$

where $P(i)$ is the probability of encountering the i-th belief, Although the definition of H is the same, the interpretation is different within a social context (between agents) and a cognitive context (within a single agent's memory).

Social Entropy. Social entropy is a measure of uncertainty or consensus of the narratives that were propagated by all agents in the network during the final time step of a simulation. The probability $P(i)$ of the i-th belief is defined as the proportion of times it is propagated over the network in a given interval time. Thus, social entropy reflects the order or disorder of each simulation's final state. Roughly speaking, 0 bits of social entropy indicates consensus, i.e. all agents propagating the same narrative with the same valence, whereas fragmentation (a diversity of opinions) is indicated by 2 or more bits of social entropy. One bit of social entropy indicates polarization.

Cognitive Entropy. Within a single agent, entropy is defined by the activation of beliefs in memory. Because the combination of narrative/valence pairs are encoded as chunks in ACT-R, entropy is calculated from the probability that a given chunk in declarative memory (DM) will be retrieved and spread over the network. In turn, the retrieval probability of a chunk i is related to a memory's base-level activation by the function:

$$P(i) = e^{B(i)} \;/\; \textstyle\sum_{j\,\in\,DM} e^{B(j)}$$

Thus defined, Shannon's entropy captures the degree of the internal organization of memories in a given agent and captures the agent's need to allocate cognitive resources to the different narratives. In this sense, information entropy has been previously used, for example, to derive predictions about the size of the hippocampus in humans (Smith et al., in press).

Materials and Methods

We conducted a 5 (Memory Decay Rate) x 10 (Network Topology) x 3 (Narrative Valence) simulation-based experiment with 10 replications per cell. The Memory Decay Rate manipulation varied the architectural decay rate parameter to address the general question of whether memory matters for simulations of information diffusion. The Network Topology and Narrative Valence manipulations addressed general questions about the effects of social context and the types of messages exchanged.

During each of the 1500 simulations, a connected social network of 200 cognitive agents exchanged a set of 10 narratives over a period of 100 ticks. During each tick, an agent encoded the narratives conveyed by all alters in its ego-network, decided on a narrative to convey, and then conveyed that narrative to all neighbors in its ego-network at the next tick. Agent behavior thus arose from a combination of neighbors' opinions and ego's evolving attitude in a closed-loop system defined by a simulation's initial state.

Initialization of social structure and cognitive agents proceeded as follows.

Network Topology: Agents were embedded in one of 10 network structures, all of which were based on a classic "caveman" graph (e.g., Watts, 1999). In our caveman networks, agents are divided into 10 "caves" of 20 agents each. All agents within each cave are fully connected with each other, except for two agents, each of which communicates with one other cave (see Figure 1).

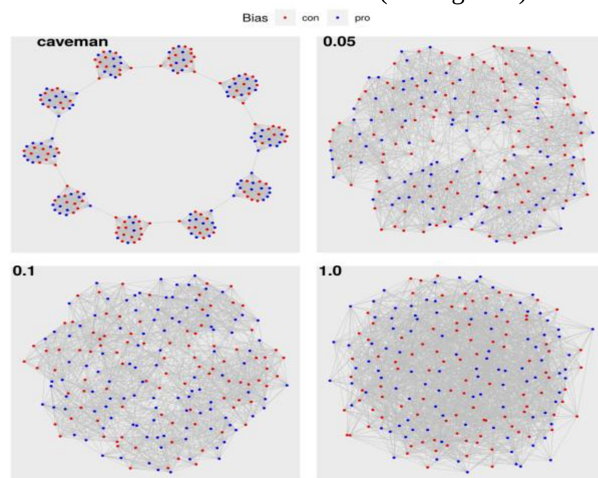

Figure 1: Different network topologies used in this study.

We manipulated the dense clustering of social interactions within caves by randomly replacing in-cave connections with new between-cave connections with probability p_{rewire}. Ten levels of p_{rewire} were used to transition from regularity to randomness: 0, 0.025, 0.05, 0.1, 0.2, 0.3, 0.4, 0.5, 0.75, and 1.0.

Agent Initialization. As described above, an agent's attitude is a consensus valence produced by blended retrievals of belief chunks from declarative memory. Each belief is a combination of a narrative with an associated *valence*, a number representing the extremity of the belief conveyed by the narrative. For instance, a neutral narrative would have a value close to zero, while extreme narratives would have values close to +1.0 or -1.0 (e.g., Eagly & Chaiken, 1993).

To establish an initial attitude for each agent, we seeded declarative memory with 100 belief chunks. The procedure to generate and allocate these agent belief histories involved three steps. First, we generated a population of notional attitudes by drawing 200 values from a truncated normal distribution (mean = 0, standard deviation = 0.25, minimum = -1, maximum = +1).

Each notional attitude was then combined with the valences associated with 10 narratives to determine the probability that a belief chunk would appear in an agent's history. Specifically, each agent's history was generated by drawing 100 samples from a discrete, truncated normal distribution (mean = notional attitude, standard deviation = 0.8, minimum = -1, maximum = +1).

For our experiment, we created three types of narrative-valence associations: polarized, centrist, and linear.

Polarized narrative valences represent strong, extremist message content (5 narratives with valences of -1 and 5 with valences of +1). Centrist narrative valences represent the use of moderate language, where valences for narratives were drawn from the same truncated normal distribution used to generate the notional attitudes. Linear narrative valences represent well-defined narratives that convey attitudes which span the valence spectrum (-1, -0.8, -0.6, -0.4, -0.2, 0.2, 0.4, 0.6, 0.8, 1).

Finally, we randomly allocated notional attitudes and their histories to agents in the network. The red and blue node colors in Figure 1 illustrate the distribution of negative and positive notional attitudes, respectively.

Results

Social networks have an inherent duality. They can be described by focusing on (a) global properties at a network Level Of Analysis (LOA), or (b) local properties at a node LOA. Our initial analyses, reported below, reflect this duality in separate ANOVAs: one concerned with entropy at a network (i.e., social) LOA, the other concerned with entropy at a node (i.e., cognitive) LOA.

A Memory Decay Rate x Rewiring Probability x Narrative Valence ANOVA of social entropy yielded a 3-way interaction, $F(72, 1350) = 1.37$, $p < .05$. Figure 2 shows the means and 95% confidence intervals for each experimental condition in the Decay x Rewiring x Valence interaction.

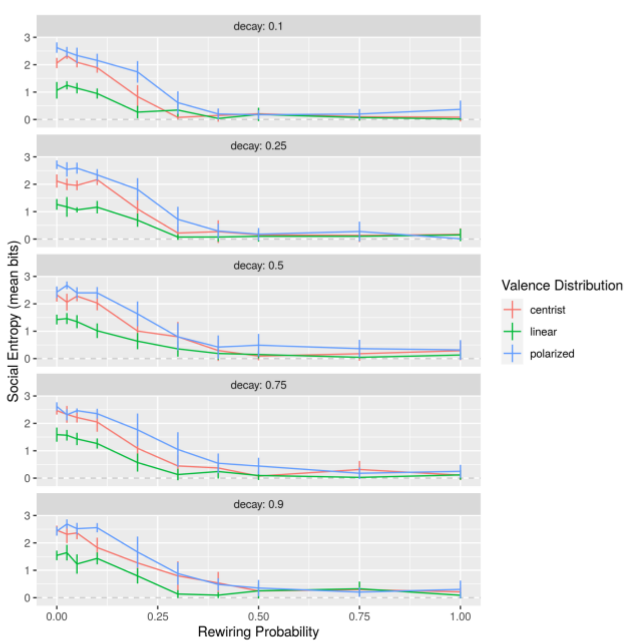

Figure 2: Effects of network topology on social entropy.

The pattern of interaction indicates that networks tend toward consensus (i.e., social entropy near 0) at rewiring probabilities of 0.4 and greater, independently of memory decay rate and narrative valence. Thus, in networks that are poor imitations of real-world social networks (i.e., those lacking local coherence), memory and message content have minimal effects on social entropy. This tendency toward consensus replicates agent-based simulation studies demonstrating assimilation with simplified agent models.

At rewiring probabilities that produce networks more similar to real-world social networks (i.e., those exhibiting dense clusters of peers), the networks tend toward polarization (i.e., entropy near 1[1]) or fragmentation (i.e., entropy near 2 or more bits), depending on memory and narrative valence. In the context of messages with a linear valence distribution, polarization is more likely to occur at reasonable values of memory decay (i.e., near the default ACT-R value of 0.25 for the transient activation noise parameter). As decay rate increases toward unrealistic values, social entropy increases and the networks tend toward fragmentation (narrative diversity) as patterns of narratives fluctuate across the network without the damping effect of memory to stabilize them. Messages with polarized and centrist valence distributions tend to produce fragmentation regardless of memory decay rate.

Generally, these results indicate that polarization is a relatively infrequent phenomenon that arises when narratives of a particular type are exchanged in realistic social networks by agents who act in a manner that is congruent with memory (e.g., strong, stable attitudes). The narratives that lead to polarization are distinct from one another (linear). Narratives that are more easily confused with one another (polarized, centrist) lead to a diversity of opinions.

Furthermore, it was puzzling that polarization, when viewed across the complete network, was rare even when we tried to force the issue by using extreme values of belief valence for the initial conditions (e.g., in the polarized condition). This may seem paradoxical, but what may explain it is that, within each cave, there existed low social entropy, due to the strong effect of the polarized initial condition. When aggregating across caves, however, the entropy is naturally larger as each cave has settled on a local set of beliefs that are uncorrelated with other caves and this is reflected as more equal probabilities for each of the 10 narrative beliefs (this would be especially true of the zero-rewire condition). A prediction, for cognitive entropy (which we explore next), is that the differences across the three valence distribution conditions will be much less than we see in social entropy, especially when the rewiring probability is zero or low.

[1] One bit of entropy is an indication that two narratives dominate a network, not necessarily that narratives representing two opposing attitudes dominate the network.

To study cognitive entropy, we conducted a 5 (memory decay rate) x 3 (narrative valence distribution) x 10 (rewiring probability) ANOVA yielded 3 2-way interactions of $p < .05$: Decay x Rewire, $F(36, 1350) = 1.86$; Decay x Valence, $F(8, 1350) = 2.15$; Rewire x Valence, $F(18, 1350) = 15.99$.

Figure 3 shows how cognitive entropy changes as a function of social context and the memory decay rate. Cognitive Entropy (and variance in cognitive entropy) generally increases (up to a certain level) as the local coherence of networks decreases (i.e., as rewiring probability increases). The relatively homogeneous social contexts provided by locally coherent networks minimize the effect of decay rate on entropy. As local coherence decreases, the decay-rate effect grows more pronounced. Thus, networks that tend toward a social consensus produce more cognitive entropy than do networks tend toward polarization or fragmentation. Our memories help reduce the degree of cognitive entropy experienced from social pressures to conform in contexts that lack the redundancy of cliquish peers.

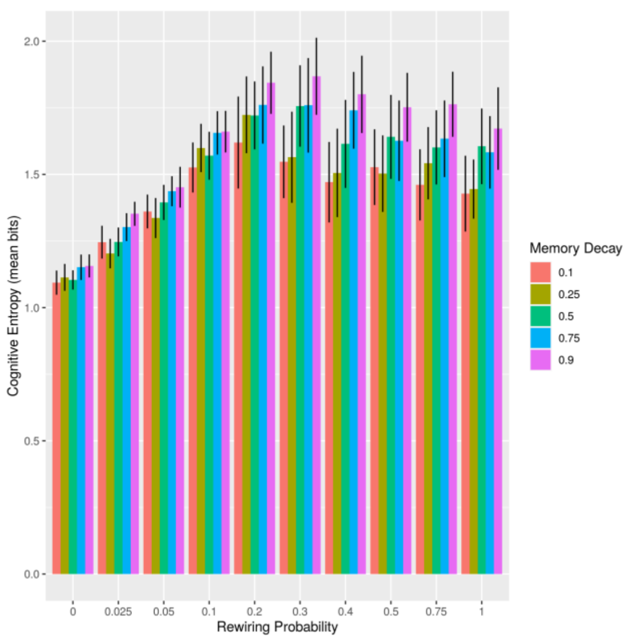

Figure 3: Effects of network topology and decay rate on cognitive entropy.

Our memories also help reduce the entropy experienced from narratives we encounter in social environments -- if the narratives can be distinguished from another. As can be seen in the left panel of Figure 4, the degree of cognitive entropy generally increases as the distinctiveness of narratives decreases. Thus a linear valence distribution (with clearly differentiable narratives) generally produces less entropy than a centrist valence distribution (with narratives that are more similar to one another), and centrist narratives produce less entropy than a polarized distribution in which everyone is using strong language which essentially conveys attitudes for or against some issue.

Interestingly, as shown in the right panel of Figure 4, heterogeneous social environments maintain the general effects of narratives on cognitive entropy: polarized > centrist > linear. In more cliquish environments, the effects of narrative valences on entropy are very similar (especially for zero rewiring probability), with the difference between polarized and centrist narratives being the most similar.

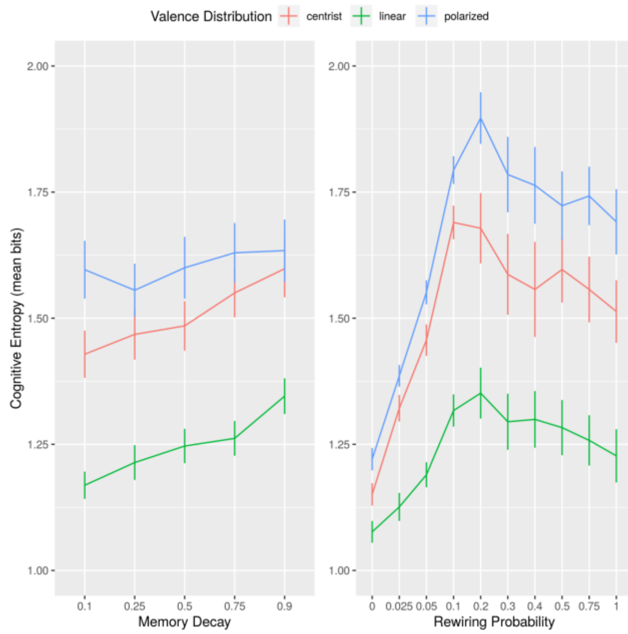

Figure 4: Effects of memory decay (left) and network topology (right) on memory entropy, divided by the distribute of narrative valence.

Discussion

We set out to show the potential value in exploring social kinds of phenomena from the perspective of cognitive architecture. At a high-level, this meant developing an understanding of the relation between internal mental representations and social structure using a single construct, information entropy. We demonstrated the ability to manipulate both cognitive and social entropy via both distribution of valence as an initial condition and network structure as a static condition.

Relations Between Cognitive & Social Entropy

According to our definition, the distribution of activation of the chunks in declarative memory determines cognitive entropy. If we imagine the hippocampus (a wetware component of declarative memory) as a communication channel between stimulus environment and response, the interpretation of cognitive entropy seems straightforward. Cognitive entropy describes the resources (i.e., channel capacity, aka attention) required to encode future events (which is compatible with the biological interpretation of Smith et al., in press).

In our simulations, agents with low cognitive entropy exist in predictable (orderly) social environments. Opinions from the neighbors of such agents provide little information for responding (i.e., propagating particular narratives). Responses of these agents thus may be driven more by expectations (cf. attitudes) than by social environments. When cognitive entropy is less than 1 bit, for example, agents "could" choose to propagate narratives that are "socially appropriate" without bothering to encode narratives received from their neighbors.

High cognitive entropy, on the other hand, indicates that social context requires responses that are more data-driven than conceptually driven. This implies that attitude strength, in some sense, should decrease as cognitive entropy increases. It also implies that agents with high cognitive entropy can be more easily influenced than those with low entropy. Furthermore, it implies that agents with low cognitive entropy may be difficult to influence, not because they harbor strong attitudes, but because experience limits their capacity for effectively encoding more complex messages; they simply do not have the bandwidth required to carry the information in complex messages that is relevant for accurate comprehension. They overgeneralize (and communication fails) because they have learned to attend to a non-discriminating subset of the features of meaning underlying the narratives of their neighbors.

Limitations

A number of limitations need to be acknowledged. This work used a highly-stylized social system to explore polarization. These results were not designed to provide insight into real-world social network dynamics, but to illustrate the approach. Another limitation is that the attitude of each belief was simulated at a purely symbolic level, without any connection to the possible effects of valence in cognition. These effects, instead, are well documented in the literature and have been incorporated into ACT-R agents in the past (Juvina et al., 2018; Smith et al., 2021). Future studies should aim to remove these limitations and test our findings in simulations with a greater degree of realism.

Implication for Polarization

These limitations notwithstanding, we believe that our results entail a number of interesting implications. First, these results might also shed light on a related, but different, problem in the social sciences: the fact that individuals who hold one extreme belief tend to harbor other extreme ones (Wood et al., 2012). In a striking example, individuals who believe in one conspiracy theory (i.e., “Princess Diana faked her own death to escape the Crown”) were also found to believe in other, incompatible ones (i.e., “Princess Diana was murdered by the Crown”).

It is highly unlikely that multiple radical beliefs spontaneously arise within a single person, as individuals Instead, extreme beliefs likely spread from person to person across social networks. This hypothesis is confirmed by the fact that the widespread use of social media in first-world countries, which amplify the reach and exposure of information, has been linked to increased partisanship, radicalization, and the spreading of fake news (Bail et al., 2018). In our results, entropy within a single agent likely tracks entropy in belief systems, and the rise of entropy in proportion to the polarization of narratives is consistent with such findings.

Finally, the finding that the network structure affects the inconsistency of beliefs has important applications for balancing policy and regulation of social media. The cognitive perspective may yield insights that are hard-gotten otherwise--understanding the micro-structure of the dynamics of social change, e.g., information operation campaigns and public health messaging, may provide the levers needed for beneficial social change.

References

Bhat, D, & Redner, S. (2019). Nonuniversal opinion dynamics driven by opposing external influences. Physical Review E., 100, 050301(R).

Bail, C. A., Argyle, L. P., Brown, T. W., Bumpus, J. P., Chen, H., Hunzaker, M. F., ... & Volfovsky, A. (2018). Exposure to opposing views on social media can increase political polarization. Proceedings of the National Academy of Sciences, 115(37), 9216-9221.

Chelian, S. E., Paik, J., Pirolli, P., Lebiere, C. & Bhattacharyya, R. (2015). Reinforcement learning and instance-based learning approaches to modeling human decision making in a prognostic foraging task. *Joint IEEE International Conference on Development and Learning and Epigenetic Robotics* (ICDL-EpiRob), pp. 116-122, doi: 10.1109/DEVLRN.2015.7346127.

Eagly, A. & Chaiken, S. (1993). The psychology of attitudes. Harcourt Brace: New York.

Elliot, A. J., & Devine, P. G. (1994). On the motivational nature of cognitive dissonance: Dissonance as psychological discomfort. Journal of personality and social psychology, 67(3), 382.

Erev, I., Ert, E., Roth, A. E., Haruvy, E., Herzog, S., Hau, R., Hertwig, R., Stewart, T., West, R., Lebiere, C. (2010). A choice prediction competition, for choices from experience and from description. *Journal of Behavioral Decision Making* 23(1): 15-47.

Gonzalez, C., Lerch, J. F., & Lebiere, C. (2003). Instance-based learning in dynamic decision making. *Cognitive Science, 27*(4), 591-635.

Hackel, L. M., Berg, J. J., Lindström, B. R., & Amodio, D. M. (2019). Model-based and model-free social cognition: investigating the role of habit in social attitude formation and choice. *Frontiers in psychology, 10*, 2592.

Huberman, Bernardo & A., Bernardo & Pirolli, Peter & T., Peter & Pitkow, & E., James & Lukose, & M., Rajan. (1998). Strong Regularities in World Wide Web Surfing. *Science*. 280. 95-. 10.1126/science.280.5360.95.

Irwin, H. J., Dagnall, N., & Drinkwater, K. (2015). Belief inconsistency in conspiracy theorists. *Comprehensive Psychology*, 4, 17-CP.

Lindström, B., Bellander, M., Chang, A., Tobler, P. N., & Amodio, D. M. (2019). A computational reinforcement learning account of social media engagement. *PsyArXiv*, (78mh5).

Orr, M., Thrush, R. & Plaut, D. C. (2013). The Theory of Reasoned Action as Parallel Constraint Satisfaction: Towards a Dynamic Computational Model of Health Behavior. Plos One, https://doi.org/10.1371/journal.pone.0062490.

Romero, O., & Lebiere, C. (2014). Simulating Network Behavioral Dynamics by using a Multi-agent approach driven by ACT-R Cognitive Architecture. In *Proceedings of the Behavior Representation in Modeling and Simulation Conference* (BRIMS-2014). Washington, DC, April 2014.

Shannon, C. E. (1948). A mathematical theory of communication. *Bell System Technical Journal*, 27(3), 379-423.

Schelling, Thomas C. "Dynamic models of segregation." *Journal of Mathematical Sociology* 1.2 (1971): 143-186.

Smith, B. M., Chiu, M., Yang, Y. C., Sibert, C., & Stocco, A. (in press). When fear shrinks the brain: A computational model of the effects of post-traumatic stress on hippocampal volume. *Topics in Cognitive Science.*

Stanley, Clayton & Byrne, Michael. (2016). Comparing vector-based and Bayesian memory models using large-scale datasets: User-generated hashtag and tag prediction on Twitter and Stack Overflow. *Psychological Methods*. 21. 542-565. 10.1037/met0000098.

Thagard, P., & Verbeurgt, K. (1998). Coherence as constraint satisfaction. Cognitive Science, 22(1), 1-24.

Van Prooijen, J. W., Krouwel, A. P., & Pollet, T. V. (2015). Political extremism predicts belief in conspiracy theories. *Social Psychological and Personality Science*, 6(5), 570-578.

Watts, D. J. (1999). Networks, dynamics, and the small-world phenomenon. *American Journal of Sociology*, 105(2), 493-527.

Wood, M. J., Douglas, K. M., & Sutton, R. M. (2012). Dead and alive: Beliefs in contradictory conspiracy theories. *Social psychological and personality science*, 3(6), 767-773.