Mnemonic Diffusion: An Agent-Based Modeling Investigation of Collective Memory

Christian C. Luhmann (christian.luhmann@stonybrook.edu)
Department of Psychology, Stony Brook University
Stony Brook, NY 11790 USA

Suparna Rajaram (suparna.rajaram@stonybrook.edu)
Department of Psychology, Stony Brook University
Stony Brook, NY 11790 USA

Abstract

Past research has uncovered evidence of social influences on a wide variety of behaviors. Everything from our choice of clothing to smoking appears to be shaped by the people we know. However, little is known about the mechanisms that underlie these influences. Here, we report a series of agent-based simulations demonstrating that information diffuses across social networks in much the same way that behavior diffuses. These findings lead us to conclude that many previously observed social influences on behavior likely rely on a substrate of information transmission and representation.

Keywords: learning, memory, collaboration, social network, agent-based modeling

Diffusion of Behavior

The idea of social influence has long been a topic of fascination for both scientists and the general public (Bartlett, 1932; Cialdini, 2001; Gladwell, 2000; Schelling, 2006). The general concept of social influence is an intuitive one. For instance, peer pressure is a factor in adolescents' tendency to drink, smoke, and engage in sexual behavior and some individuals slavishly follow the latest fashion trends, mimicking the styles seen on the runway or worn by celebrities. However, the intuition about social influence is far too narrow. The examples cited above are seen as exceptions; perhaps the susceptibility is restricted to a particularly impressionable population (e.g. adolescents) or perhaps influence is seen for relatively trivial behaviors (e.g., the clothes you wear). Such mindless imitation would not be seen, intuition suggests, concerning behaviors that are both deeply personal and of great consequence (e.g., suicide, how many children to have). Surprisingly, research has challenged this intuition, finding that social ties strongly influence a wide range of behaviors, including transformative life decisions (Christakis & Fowler, 2009; Watts, 2003). For example, a program of research by Christakis and Fowler (2007) has revealed the surprising "contagion" of health-related attributes such as obesity and smoking.

In psychology, there is a long history of work exploring the influence of social context on behavior. Early work focused on the potentially deleterious behavioral effects of social influence. For example, Milgram (Milgram, Bickman, & Berkowitz, 1969) examined social influences on arbitrary behavior as a function of group size. In this study, groups of between one and 15 confederates stood on a New York street corner looking upward at a window overhead. Passing pedestrians were likely to mimic some aspect of the observed behavior (e.g., looking up) and this tendency increased with group size. The classic studies of Asch (e.g., Asch, 1951) demonstrate the power of social influence even more starkly because his participants were asked about matters of objective fact (e.g., the length of lines). Despite being accurate when making judgments individually, participants placed with confederates tended to conform, producing substantial errors.

Though past work has revealed the presence of social influences on a variety of behaviors, we know little about the mechanisms that underlie these influences. For example, it has been suggested that, "Social networks function…by giving us access to what flows within them" (Christakis & Fowler, 2009, p. 91). But what does flow within our social networks that allows for these powerful influences on our behavior?

Existing Models of Social Influence

Several mathematical models of social influence have been proposed (Easley & Kleinberg, 2010; Jackson, 2008; Lopez-Pintado & Watts, 2008). Among the most influential are linear threshold models (Granovetter, 1978). Such models assume that each individual has two mutually exclusive and exhaustive behavioral options available. For example, in Granovetter's classic example, each individual chooses whether or not to join a riot. In addition, each individual is assumed to observe the behavior of all other individuals. The decision of the individual is then a function of their own idiosyncratic threshold and the behavior observed in the group. If the number of other people observed to be rioting does not exceed the individual's threshold, she remains a bystander. If this number exceeds the individual's threshold, she begins to riot.

Several key details of the current crop of mathematical models should be noted. First, this work typically assumes "zero-intelligence agents" (Gode & Sunder, 1993) that can do nothing but copy the behavior of their neighbors with some probability (e.g., Granovetter, 1978). This is undoubtedly convenient, but represents a substantial simplification, at least when attempting to model human behavior. Though some behaviors may be the result of innate imitative mechanisms, many more behaviors are deliberative, relying on individuals' beliefs, goals, and desires. Second, many of the computational models of behavioral diffusion assume that individuals can occupy one of a small number of behavioral "states". For example, in the rioting example discussed above, individuals are assumed to either be rioting or not rioting, another vast over-simplification. Finally, because of these simplifications, previous models have largely avoided questions about the mechanisms by which behaviors are transferred between individuals. Indeed, these models expressly omit such mechanisms by assuming that mimicry is the critical basis of diffusion. To address the mechanisms themselves, the current study takes an agent-based modeling (ABM) approach (Smith & Conrey, 2007; Stasser, 1988; Carley, Martin, & Hirshman, 2009; Parunak, Belding, Hilscher, & Brueckner, 2009; Coman, Kolling, Lewis, & Hirst, 2012) in which agents are information processing units capable of representing information and learning.

The current study investigates the representation and transmission of information within social networks as fundamental mechanisms underlying these potent influences on our behavior. Specifically, we investigate how information is represented by individuals within a larger network and how the nature of social interactions shape the information as it flows through the network.

Collaborative Memory in Small Groups

Recent behavioral work on the social transmission of memory in small groups has identified several key mechanisms that facilitate or inhibit information transmission in small groups, and how the interaction among these mechanisms shapes convergence amongst group members (what is referred to as collective memory). The collaborative memory paradigm provides a robust method for measuring the transmission of information in small groups of two or three participants (Rajaram, 2011). In this paradigm, each participant is first exposed to experimenter-provided stimuli (words, pictures, narratives). Participants then form groups and recall items collaboratively. Finally, participants recall items individually to assess the post-collaborative representations retained by each participant.

The consequence of collaboration on group memory is counterintuitive. Though a collaborating group recalls more than a given individual, the group recalls significantly less than its potential, a phenomenon called collaborative inhibition (Weldon & Bellinger, 1997). To estimate the group's potential, performance is compared to that of a nominal group: the total, nonredundant recall of an equal number of participants who recalled individually (Blumen & Rajaram, 2008; Weldon & Bellinger, 1997). Although it seems reasonable to assume participants perform suboptimally because they feel less accountable while working in groups (social loafing; Latane, Williams, & Harkins, 1979), experimental findings shows this is not the case (Weldon, Blair, & Huebsch, 2000).

Mechanisms Involved During Collaboration

The suboptimal performance of collaborative groups has been attributed to the retrieval disruption process where the output of one participant's recall disrupts other participants' attempts at recall, and as a result lowers the latter participants' output (B.H. Basden, Basden, Bryner, & Thomas, 1997). Because each individual recalls less than her potential during collaboration, researchers have asked whether their post-collaborative representations would continue to exhibit this deficit. Though some forgetting does occur (Cuc, Koppel, & Hirst, 2007) two mechanisms usually enhance the quantity and accuracy of post-collaborative representations; one, items not recalled during collaboration bounce back post-collaboratively (rebound) and, two, collaboration acts to expose each participant to items she might not have remembered otherwise (re-exposure, Blumen & Rajaram, 2008, 2009; Congleton & Rajaram, 2011).

Several individual- and interaction-based properties influence these collaborative effects. One such change of note relates to increase in memory errors. As one example, social contagion errors arise when the stimuli activate plausible items for recall that were never presented (B.H. Basden et al., 2002; French, Gary,&Mori, 2008; Reysen, 2007; Roediger, Meade, & Bergman, 2001). Such contagion has been demonstrated in collaborative studies (B.H. Basden et al., 2002) using DRM stimuli (Roediger & McDermott, 1995) in which a list of associatively-related words such as *bed, rest, awake, tired, dream, wake, snooze, blanket, etc.* leads participants to recall the never-presented lure (*sleep*) with great confidence. Propagation of memory errors in the real-world has been an enduring concern of cognitive scientists (e.g., Bartlett, 1932) but empirical investigations have remained elusive due to feasibility.

Memory representations in small groups are also characterized by the frequency with which information is processed before and during collaboration. For instance, the individual who dominates the collaborative discussion benefits most from rehearsal (Rajaram & Pereira-Pasarin, 2010) and has the largest influence on the post-collaborative representations of other group members (Cuc, Ozuru, Manier, & Hirst, 2006). Conversely, post-collaborative memory deficits occur for information not discussed during collaboration, either through omission (Cuc et al., 2007), rejection of correct responses (Merckelbach, van Roermund, & Candel, 2007), or group conformity to incorrect responses (Reysen, 2005). We have further shown that frequency of discussion prior to or during collaboration changes both collaborative group recall and post-collaborative memory; when information is repeatedly processed prior to collaboration, it can reduce or even eliminate collaborative inhibition in group recall (Congleton & Rajaram, 2011; Pereira-Pasarin & Rajaram, 2011), and improve post-collaborative memory (Congleton & Rajaram, 2011). Just as interestingly, when groups are given the opportunity to discuss more frequently this too reduces collaborative inhibition in group recall (B.H. Basden et al., 2000; Blumen

& Rajaram, 2008). These behavioral outcomes raise intriguing questions about how frequency of discussion influences group-level representations in large social networks.

Yet another intriguing finding concerns the effects of group size. Even within small groups, research shows that as group size increases (from 2 to 3 or 4 members) collaborative inhibition increases with group size (B.H. Basden et al., 2000; Thorley & Dewhurst, 2007). This raises the question about whether larger social networks would display an exaggerated version of this decline or whether the complex interplay of mechanisms would completely change how the group-level representation evolves.

Current Approach

In the current study we investigate the processes that shape the transmission of information during both the collaborative remembering in the laboratory paradigm and more realistic social contexts. We take an agent-based modeling approach in which individuals are represented by computational agents and allowed to interact much as human subjects interact in the collaborative memory paradigm. The agents are endowed with simplified memory models capable of storing a set of N items (e.g., words). The memory model consists of two separate representations. First, agents represent a set of inter-item associations that exist prior to any social interaction (a matrix denoted S). These associations represent pre-experimental knowledge such as the semantic associations between words. For the sake of simplicity, these inter-item associations were assigned random values (within the range [0,1]) in the current simulations. More systematic prior knowledge could obviously be constructed, particularly if such factors were important for specific research questions. For example, lists of categorized words can be simulated by constructing high within-category associations and low between-category associations, a strategy we have successfully used in recent modeling (Luhmann, Congleton, Zhou, & Rajaram, 2013). The second representation is a set of N activations (a vector denoted A, with elements bound to the range [0, 1]), which allow for learning to occur during the experimental experience itself. For example, these activations capture recent experience studying experimenter-provided stimuli (e.g., word lists), items generated by collaborative partners, and even items generated by the agent itself.

Agents have two behaviors. First, they may encode a presented item by increasing the activation associated with the presented item (i.e., $\Delta A_i = \alpha[1 - A_i]$, where A_i is the activation of the item and α is a learning rate). This encoding occurs when items are presented by the experimenter (i.e., during the collaborative memory paradigm's initial study phase) and when agents are exposed to the items retrieved by other agents (e.g., during collaborative recall). Second, agents can retrieve an item. This is done by randomly selecting an item in proportion to the activation levels in A (i.e., more active items are more likely to be generated). If the activity of the candidate item is above the agent's recall threshold (γ) and has not yet been generated by the group, then this item is successfully retrieved and generated (e.g., spoken out loud). Finally, associates of the retrieved item (from S) have their activations decreased (i.e., $\Delta A_j = -\beta S_{ij} A_i$ where A_i is the retrieved item, S_{ij} is the strength of the association between items i and j, and β is a forgetting rate).

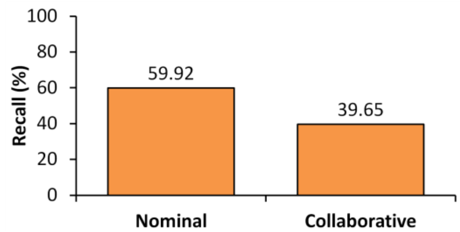

Figure 1 – Simulations replicate the collaborative inhibition effect

Simulation 1: Collaborative Inhibition

Our first investigation is of the most surprising finding to come out of the collaborative memory paradigm: collaborative inhibition. This was done for two reasons. First, the finding is elicited from a fairly simple experimental paradigm making these initial simulations relatively straightforward to construct. Second, to the degree we are capable of successfully replicating the least intuitive aspect of the empirical data, we can proceed with somewhat more confidence that our formalism has not overly simplified the cognitive processes involved.

To simulate the collaborative memory paradigm, we first presented the entire list of N items to each agent in a random order. Three agents were then allowed to interact with one another. The interaction was structured such that each agent was given an opportunity to retrieve an item on each round. If an agent successfully retrieved an item, the retrieved item was encoded by the other two agents. Figure 1 illustrates the results of 1000 simulations of a collaborative condition and 1000 simulations of a nominal condition (i.e., total, nonredundant recall of three agents recalling individually) evaluated exactly as in the behavioral studies described above. As can be seen, the simulation results reproduce the collaborative inhibition findings describe above**.** This result is likely due to the fact that each agent is endowed with an idiosyncratic set of activations during the initial, individual study phase but then learns the contents of their peer's activations during the collaborative phase. Thus, the interaction between agents tends to increase the similarity of the agents' representations and minimize the idiosyncrasies that make the nominal groups more successful in generating greater quantity. Furthermore, exploratory simulations suggest that having individual agents repeatedly engage in isolated retrieval does not diminish the performance of the nominal group, suggesting that the collaborative inhibition is due to the social interaction (e.g., retrieval disruption, B.H. Basden et al., 1997) rather than repeated retrieval.

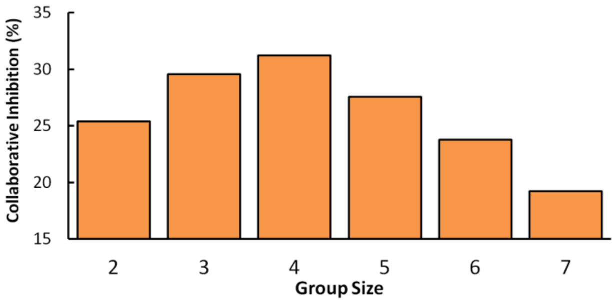

Figure 2 – Influence of group size on collaborative inhibition

Simulation 2: Group Size

As mentioned above, prior work with the collaborative memory paradigm has manipulated a variety of different factors. One factor that has received surprisingly little attention is the size of the collaborating group despite its obvious real-world relevance. Only two studies have manipulated group size (B.H. Basden, Basden, and Henry, 2000; Thorley and Dewhurst, 2007), and both concluded that increasing group sizes produce more detrimental collaborative effects. Details of these studies limit interpretation, however. For example, Basden et al. (2000) tested 1-, 2-, and 4-person groups. Though the 1-person groups recalled more than the 4-person groups, the 2-person groups were different from neither. Thorley and Dewhurst (2007) used DRM stimuli and were specifically interested in groups' tendency to falsely recall the lure items (rather than recall per se). Further, the group size tested was small in these studies (2-4 person groups) again limiting interpretation.

Here, we systematically manipulate group size and investigate the influence of this factor on collaborative inhibition. We simulated collaborative groups ranging in size from two to seven as well as nominal groups consisting of agents retrieving in isolation. As before, the entire list of N items was first presented to each agent. The agents within a group were then allowed to interact with one another with the interaction structured as described above.

Figure 2 illustrates the results of 1000 simulations for each group size. The standard collaborative inhibition effect (measured here as Nominal - Collaborative) was found for all group sizes. However, the relationship between group size and collaborative inhibition was not entirely straightforward. As groups grew from two to four, collaborative inhibition increased (replicating Basden et al., 2000). However, as group size increased further, collaborative inhibition decreased. This non-monotonic relationship appears to be driven by the relative balance between the facilitative effects offered by collaboration (i.e., more agents increase the probability that the group will retrieve a given item) and the detrimental effects of retrieval disruption (i.e., more collaborators means more opportunities to be disrupted).

Simulation 3: Diffusion of Collective Memory

The collaborative memory paradigms represent a realistic, real-world social network that is amenable to experimental study. However, the size of groups involved in this paradigm places obvious restrictions on the research questions that may be asked. The current simulation seeks to achieve substantially greater realism than the more traditional laboratory paradigms allow. Specifically, we wish to explore how the information represented by and shared between individuals makes its way through larger social networks.

To explore true social networks, we employed a larger population (60 computational agents of the kind described above), each of which was placed into a larger network structure. Though there are many potentially interesting network structures, we are most interested in those related to real world social networks. For this reason, the current simulation employs a so-called small world network (Watts & Strogatz, 1998). Within such a network, the shortest distance between two nodes is short on average despite the network itself being relatively sparse (most nodes are not neighbors). These features give rise to the well-known "six degrees of separation" phenomenon. We further chose to set the average degree to 2 (meaning that agents were, on average, connected to 2 other agents).

As in the simulations reported above, each simulation began by presenting the entire list of 40 items to each agent individually. All subsequent interaction between agents occurred over this network. On each epoch of the simulation, a random agent was selected along with one of that agent's randomly selected neighbors. This pair was then allowed to interact just as the collaborative groups simulated above (each taking a turn to retrieve, etc.). To assess the diffusion of information across the network, we computed the similarities between pairs of agents' representations (i.e., the correlation between activation vectors, A) after the simulations were completed. This measure of similarity goes up with the overlap (both in what they represent strongly and what they have forgotten, i.e., collective memory) and goes down when one agent has forgotten an item that the other agent still remembers (i.e., is strongly active in A). We computed the similarity between all pairs of agents (i.e., both neighboring pairs and non-neighboring pairs) sorting these similarities on the basis of how close the two agents in each pair were to each other within the network (i.e., minimum distance). Neighboring agents would have a distance of 1. Two non-neighboring agents that shared a common neighbor would have a distance of 2 and so on.

Figure 3 illustrates the results of a 60-node small world network that was allowed to run for 1000 epochs. As can be seen, neighboring agents acquired very similar representations. This is not particularly surprising since neighboring agents will have interacted with each other and learned the contents of their neighbors' representations. What is surprising is that agents at a distance of two are highly similar as well. These agents never interacted with one another, so direct communication cannot explain this similarity. Instead, the two agents' common neighbor presumably acted as a conduit through which information diffused, indirectly connecting the non-neighbors. Even more surprising then, is the fact that agents at distance three, separated by two intermediate agents, are also somewhat similar. After this point, the similarity between agents levels off, reflecting the boundaries of collective memory in large networks.

This similarity between non-neighboring agents is a phenomenon that has been observed in a variety of real-world social networks and is known as hyperdyadic spread (Christakis & Fowler, 2009). For example, previous work has shown that people are 57% more likely to become obese if a peer (e.g., friend) becomes obese and 20% more likely to become obese if a peer of a peer (e.g., a friend of a friend) becomes obese. Furthermore, in many of the behaviors studied within social networks, hyperdyadic spread from a given node in the social network has been found to extend to three "hops" from that node (e.g., to the friend of a friend of a friend) but not beyond – what Christakis and Fowler (2009) have termed the three degrees of influence rule. The fact that our simulations comply with this rule is interesting because the standard finding of hyperdyadic spread concerns the spread of behavior whereas the current results reflect the spread of information. Exploratory simulations employing other network structures (e.g., chains, trees) have either failed to uncover strong hyperdyadic spreading or failed to conform to the three degrees of influence rule. This suggests that this class of phenomena may be jointly driven by both the details of the social networks in which we live (e.g., small-world networks) and the constraints of human learning and memory.

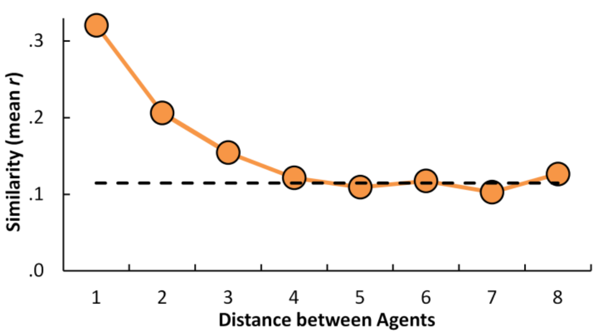

Figure 3 – Hyperdyadic spread of collective memory in our simulation results

Discussion

Broadly speaking, the goal of the current study has been to investigate social influence in real world social networks. Prior research has developed formal models to capture how behaviors diffuse amongst large groups, but these formalisms have been relatively agnostic about the underlying psychological mechanisms, instead modeling such behaviors as being literally contagious. Here, we have argued that the cognitive processes governing learning and memory are likely candidates for such mechanisms as they are in prime position to influence the representations individuals hold and transmission of information between f

Past work using the collaborative memory paradigm has provided useful empirical data with which we can begin to explore our proposal. In this paradigm, groups of individuals first study items in isolation before collaborating in groups to recall these same items. Despite the practical limitations posed by this paradigm (e.g., small group sizes), the literature has provided a wealth of insights into the social influences on learning and memory. These insights include the role of retrieval disruption, re-exposure, and error correction, the influence of group size, and phenomena such as collaborative inhibition and error propagation.

In the current study, we have taken an agent-based modeling approach, simulating individuals as relatively simple information processing units capable of representing information, learning from experience, and interacting with other agents. In order to evaluate our proposal, we selected three different phenomena to explore. We first investigated the robust collaborative inhibition effect. Our simulations replicate the standard pattern of results, with collaborative groups under-performing relative their controls. We next simulated the somewhat less thoroughly studied role of group size on the collaborative inhibition effect. Here, we found that our simulations were capable of replicating the effects observed in the literature (increasing collaborative inhibition with increasing group size) but also made predictions about the boundary conditions of these effects. Finally, we extended our findings beyond the collaborative memory paradigm to investigate agents in a larger social network. Here, we found that our simulations exhibited hyperdyadic spread, a standard empirical finding in the diffusion of behavior across social networks.

We take the success of the current simulations as evidence in favor of our proposal. Our simulations demonstrate that the spread of information across connections in a social network mirrors the way in which behavior spreads across those same connections. Thus, it seems likely that social influence, and particularly the diffusion of behavior, relies on a substrate of information transmission and representation.

References

Asch, S. E. (1951). Effects of group pressure upon the modification and distortion of judgments. In H. Guetzkow (Ed.), *Groups, Leadership, and Men* (pp. 177-190). Pittsburgh: Carnegie Press.

Bartlett, Sir F. C. (1932). *Remembering: A Study in Experimental and Social Psychology*. London: Cambridge University Press.

Basden, B. H., Basden, D. R., & Henry, S. (2000). Costs and benefits of collaborative remembering. *Applied Cognitive Psychology*, *14*(6), 497-507.

Basden, B. H., Reysen, M. B., & Basden, D. R. (2002). Transmitting false memories in social groups. *American Journal of Psychology*, *115*, 211-231.

Basden, B. H., Basden, D. R., Bryner, S., & Thomas, R. L., III (1997). A comparison of group and individual remembering: Does collaboration disrupt retrieval strategies? *Journal of Experimental Psychology: Learning, Memory, and Cognition, 23,* 1176-1189.

Blumen, H. M., & Rajaram, S. (2008). Influence of re-exposure and retrieval disruption during group collaboration on later individual recall. *Memory, 16,* 231-244.

Blumen, H. M., & Rajaram, S. (2009). Effects of repeated collaborative retrieval on individual memory vary as a function of recall versus recognition tasks. *Memory, 17*, 840-846.

Carley, K. M., Martin, M. K., & Hirshman, B. R. (2009). The etiology of social change. *Topics in Cognitive Science, 1*(4), 621-650.

Christakis, N. A., & Fowler, J. H. (2007). The spread of obesity in a large social network over 32 years. *New England Journal of Medicine*, 357, 370-379.

Christakis, N. A., & Fowler, J. H. (2009). *Connected: The Surprising Power of our Social Networks and How they Shape our Lives*. New York: Little, Brown and Company.

Congleton, A. R. & Rajaram, S. (2011). The influence of learning methods on collaboration: Prior repeated retrieval enhances retrieval organization, abolishes collaborative inhibition, and promotes post-collaborative memory. *Journal of Experimental Psychology: General*, *140*, 535-551.

Coman, A., Kolling, A., Lewis, A., & Hirst, W. (2012). Mnemonic convergence: From empirical data to large-scale Dynamics, *Social Computing, Behavioral-Cultural Modeling and Prediction*, 256-265.

Cuc, A., Koppel, J., & Hirst, W. (2007). Silence is not golden: A case for socially shared retrieval-induced forgetting. *Psychological Science*, *18*, 727-733.

Easley, D., & Kleinberg, J. (2010). *Networks, Crowds, and Markets: Reasoning About a Highly Connected World*. Cambridge: Cambridge University Press.

French, L., Gary, M., & Mori, K. (2008). You say tomato? Collaborative remembering leads to more false memories for intimate couples than for strangers. *Memory*, *16*, 262–273.

Gladwell, M. (2000). *The Tipping Point: How Little Things Can Make a Big Difference*. New York: Little, Brown & Company.

Granovetter, M. (1978). Threshold models of collective behavior. *American Journal of Sociology, 83*, 1420-1443.

Gode, D. K., & Sunder, S. (1993). Allocative efficiency of markets with zero-intelligence traders: Market as a partial substitute for individual rationality. *Journal of Political Economy*, *101*, 119-137.

Jackson, M. O. (2008). *Social and Economic Networks*. Princeton: Princeton University Press.

Latane, B., Williams, K., & Harkins, S. (1979). Many hands make light the work: The causes and consequences of social loafing. *Journal of Personality and Social Psychology*, *37*, 822-832.

Luhmann, C.C., Congleton, A.C., Zhou, X., & Rajaram, S. (2013). *When less is more: Retrieval-refined representations produce the testing effect*. Manuscript submitted for publication.

Merckelbach, H., van Roermund, H., & Candel, I. (2007). Effects of collaborative recall: Denying true information is as powerful as suggesting misinformation. *Psychology, Crime, & Law*, *13*, 573–581.

Milgram, S., Bickman, L., & Berkowitz, L. (1969). Note on the drawing power of crowds of different size. *Journal of Personality and Social Psychology, 13*, 79-82.

Parunak, H. V., Belding, T. C., Hilscher, R., & Brueckner, S. (2009). Understanding collective cognitive convergence *Multi-Agent-Based Simulation IX* (pp. 46-59): Springer.

Pereira-Pasarin, L., & Rajaram, S. (2011). Study repetition and divided attention: Effects of encoding manipulations on collaborative inhibition in group recall. *Memory & Cognition, 39*, 968-976.

Rajaram, S. (2011). Collaboration both hurts and helps memory: A cognitive perspective. *Current Directions in Psychological Science*, *20,* 76-81.

Rajaram, S., & Pereira-Pasarin, L. (2010). Collaborative memory: Cognitive research and theory. *Perspectives on Psychological Science, 5, 649-663.*

Reysen, M. B. (2005). The effects of social pressure on false memories. *Dissertation Abstracts International: Section B: The Sciences and Engineering, 65*, 6066.

Roediger, H. L., & McDermott, K. B. (1995). Creating false memories: Remembering words not presented in lists. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 21*, 803-814.

Roediger, H.L., III, Meade, M.L., & Bergman, E.T. (2001). Social contagion of memory. *Psychonomic Bulletin & Review, 8*, 365–371.

Smith, E. R., & Conrey, F. R. (2007). Agent-based modeling: A new approach for theory building in social psychology. *Personality and Social Psychology Review*, *11*, 87-104.

Stasser, G. (1988). Computer simulation as a research tool: The DISCUSS model of group decision making. *Journal of Experimental Social Psychology, 24*, 393-422.

Thorley, C. & Dewhurst, S. A. (2007). Collaborative false recall in the DRM procedure: Effects of group size and group pressure. *European Journal of Cognitive Psychology, 19*, 867-881.

Watts, D. J., & Strogatz, S. H. (1998). Collective dynamics of “small-world” networks. *Nature, 393*, 440-442.

Weldon, M. S., & Bellinger, K. D. (1997). Collective memory: Collaborative and individual processes in remembering. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 23,* 1160-1175.

Weldon, M. S., Blair, C., & Huebsch, D. (2000). Group remembering: Does social loafing underlie collaborative inhibition? *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *26*, 1568-1577.