Advice, Trust, and Gossip Among Artificial Agents

Michael Prietula
Decision and Information Analysis
Goizueta Business School
Emory University

Chapter 5 in A. Lomi and E. Larsen (Eds.),
Dynamics of Organizations:
Computational Modeling and Organizational Theories
Cambridge, MA: AAAI Press/MIT Press, 2001.

Funding for this work is provided in part from:

National Science Foundation, Directorate for Computer, Information Science & Engineering
Computational Social Systems Program, Award IIS-0084508.
All errors and other burdens are the author's.

Advice, Trust, and Gossip Among Artificial Agents

Michael J. Prietula
Emory University

In a nutshell, I am suggesting that language evolved to allow us to gossip.
R. Dunbar (1996)

What is necessary to build a social agent?
K. Carley & A. Newell (1994)

Do you trust me?
(Aladdin to Jasmine)
Aladdin and the King of Thieves (Disney)

In the Foreword to this volume, Jim March (2001) presents an insightful excursion into the historical role that computer simulation has played (or, in most cases, has not played) in halls of organizational theory, and identifies two general theoretical problems that can be tackled by such simulations: ecological (contextual) complexity and historical (temporal) complexity. The former addresses the difficulties of making macro predictions (or explanations) involving certain types of interactions of micro-events (e.g., agent decisions and behaviors). The latter addresses the difficulties in making state predictions (or explanations) that are strongly derivative of prior events (especially if a component of those prior events involved exogenous random components). As the reader will discover, the topic of this chapter concerns an organizational system that exhibits both ecological and historical complexities. As the reader will judge, computational modeling provides a mechanism to embrace those complexities and generate insight into the organizational system of interest.

In this chapter, it is argued that the emergent advice coalitions, such as those on the Internet, form a new model of business that incorporates a fundamental human activity to control for sources of bad advice – gossip. In addition, it is argued that gossip and trust interact, but in a manner that actually insulates trust while altering advice-taking behaviors. How these interact is explored in a series of computer simulations. A series of computer simulations is created to explore the implications of these proposals. The underlying substrate for the arguments concerning gossip, social agent coalitions, and the methodology is given by weaving together several theoretical positions.

Gossip. The first theoretical position addresses the role of gossip and is articulated in the first quote above. Specifically, Dunbar (1996) argues that one of the central functions of human language (i.e., human communication) is to exchange information about other people's behavior in a group through means other than direct observation. This, in effect, is a form of social grooming that facilitates the definition and maintenance of group membership, and affords several advantages that work to permit the size of the group to be increased (i.e., by replacing a 1-to-1 physical-contact grooming constraint with 1-to-N verbal grooming) as well as to maintain and coordinate the group (i.e., information details on behavior of members observed by others via gossip). In fact, the theory argues that the complexity of the demands of social interaction is the primary factor in emergence of intelligence. This is closely related to the Machiavellian intelligence (e.g., Humphrey, 1976; Byrne, 1997; Whiten & Byrne, 1997) and the social intelligence hypotheses (Kummer, et al., 1997), where the evolutionary pressures that shaped the human mind addressed the relation of the individual to the social environment and less so towards the physical environment, though they are certainly intertwined. The primary reason is that the deliberation required for social interactions (social objects) are relatively more complicated (but of higher adaptive value) than deliberations required for physical objects (tools).

However, this chapter is not about evolutionary anthropology, but about interacting social groups, which leads to Dunbar's concluding speculations. Dunbar interprets his theory to current societal behaviors and influences. In particular, Dunbar argues that the impersonality of the "information superhighway " (i.e., the Internet) will inhibit socially-oriented behaviors and negotiations as people will "lose control that social interaction normally imposes in the interests of cooperation and bonding" (p. 205). According to Dunbar (1996), "suspicion of the unknown and the fear of being duped by untrustworthy strangers will continue to dictate our decisions" (p. 205), and "favours will only be done when there is a clear quid pro quo, an immediate return to the giver, rather than being a matter of communal obligation" (p. 206). Is this social extrapolation of an evolutionary theory correct?

It is proposed that Dunbar's *extrapolations* about Internet behaviors are not entirely plausible; however, the *behaviors* occurring on the Internet are actually supported by his theory.

With the development of email standards and access to – and growth of – the commodity Internet, organization and societal members have experienced substantial increases in connectivity. Furthermore, the growth of rapid communication is greatly facilitated by the Internet and linkages may expand well beyond the artificial or geographical boundaries of the organization, group, or coalition. This permits diverse sets of participants to engage in information exchanges and, depending on the particular context, commerce. Consequently, social behavior does occur as many forms of commerce-based coalitions arise opportunistically on the Internet, where advice from strangers is often sought and gossip seems an integral component of maintaining and communication behaviors about the behaviors of those dynamic and emergent coalitions. Though different than the original social groups defined by Dunbar, these are also social groups that cooperatively interact (i.e., it is not a requirement for any purchase or exchange) through online messages.

Consider a study by Constant, Sproull and Kiesler (1997) where they explored people's use of "weak ties" (i.e., relationships with strangers or mild acquaintances) to obtain technical advice in a larger international firm (Tandem Computers), by analyzing the results of broadcast requests over their internal computer network. They found that 81% of the information providers (responders) did not know the seekers (who posted the request) at all and that acquaintanceship was uncorrelated with the number of replies or the usefulness of the replies. In that corporate setting, advice was not a function of prior personal relationships. What mattered was the advice, not the strength of the personal relationship ties. In a related series of articles in the Harvard Business Review, Prietula and Simon (1989) describe the important role of "hidden experts" in an organization are argue that expertise is not necessarily located where managers think it may be, and the importance of managers to know where the expertise actually resides; Kelley and Caplan (1993) show that defining expertise in an organization is not as straight-forward as some novice employees (or experience managers) might think; and, Krackhardt and Hanson (1993) provide an insightful piece on the difficulty and importance of informal networks in an organization.

The research in this chapter resembles, in part, the situation addressed the Tandem example and by the "weak ties" theory of Granovetter (1973). However, that theory basically examines how the strength of dyadic ties relates to the degree of overlap of their friendship network, and the subsequent implications (such as the added value of dissimilar, relatively unattached constituents over more similar, strong-tie constituents that is based, in part, on the likelihood of dissimilar, but valuable, information from the weak-tie sources). The groups simulated in this chapter are indeed weak-tie groups (and some of the weak-tie theory implications certainly apply), but there is no alternative for comparison as these are individuals who are brought together opportunistically with no prior contacts. This lack of prior history and no current organizational alliance also differentiates the groups in this chapter from those that are opportunistic, but do have organizational alliances, such as in temporary work groups (e.g., Meyerson, Weick & Kramer, 1996).

Many of the connected forms, however, are collections of individuals who are gathered (virtually) in a common (virtual) place for a common purpose. Participation in these groups may be ephemeral (perhaps task- or event- based) and formal membership or extensive historical-social ties do not exist, as the set of participants may be complete (and even anonymous) strangers. What is found in these groups is that strangers exchange information (experience and rumors) about the quality of advice from other strangers – they gossip.

Allport and Postman (1965) provide a definition of rumor as "…a specific (or topical) proposition for belief, passed along from person to person, usually by word of mouth, without secure standards of evidence being present" (p. ix). They argued that there are two basic conditions for rumor: (1) the content involves something of "importance" to the speaker and listener, and (2) the truth is ambiguous. Rosnow and Fine (1976) offer a succinct definition of rumor as "a proposition that is unverified and in general circulation." Thus, the truth or falsity of a rumor is not the issue, for truth or falsity is unknown; rather, it is that truth or falsity is not immediately verifiable and that the proposition be dispersed. Rumors differ from other sorts of social story exchanges (e.g., legends) in that they address current events, are about specific facts with respect to those events, and are intended to be considered for belief (Kapferer 1990). The distinction between gossip and rumor is difficult at best in complex situations. For example, Koenig's (1985) study of commercial rumors considers them related, with gossip as "communication about people known to the persons involved in the communicating" (p. 2), and a form of social control, while rumor is a "story or report current within any known authority for its truth" (p. 2), thus the difference is less of kind and more of context or extent. Little research has been done in the context of the Internet, though rumor/gossip are known phenomena that have resulted in lawsuits (Associated Press, 1999a) and international incidents (Associated Press, 1999b). In this chapter, the context is simple so rumor and gossip will be considered equivalent. Gossip is taken to be *information conveyed about the quality of advice from another source.*

Social Knowledge and Information Processing. This leads to a second theoretical stance and addresses the second quote at the chapter's beginning. In accordance with the prior (evolutionary) theoretical stance, social interactions are adaptively beneficial, but require a richer cognitive apparatus (more "intelligence," richer language) than non-social interactions, and gossip plays an important role in those social situations. But to extrapolate from that theoretical stance to current social situations requires a more refined view of both the social interaction and the nature of the agents (the players) in those situations. It requires a perspective that can differentiate between the general categories of social problem solving that may be required and the type of problem solvers that are engaged in that social setting. Such a perspective can be found in the Model Social Agent matrix.

The Model Social Agent matrix (Carley & Newell, 1994) is a two-dimensional categorization scheme that specifies the kind of *knowledge* required by the agent to operate in a particular type of social setting (thus defining particular types of social behaviors). One axis defines increasingly complex social settings ranging from non-social (little knowledge required) to cultural-historical (much knowledge required). The other axis defines the particular type of information processing capability of the agents as specified by construction, calculation, or assumption, ranging from Omnipotent (infinite) to cognitive-emotional (quite restricted). The intersections of the matrix place a type of information processing agent in a particular type of social situation, thus defining the broad fundamental *assumptions* of the agent and the situation within which the agent is to behave. The behaviors that emerge depend on the task (i.e., the problem) given to the agent(s) in the context of its (their) information processing capabilities and the knowledge required to function in that social situation. As one can imagine, variance in behavior is accounted for by the capabilities of the agent, the knowledge possessed by the agent, and the demands of the task on the agent. It is from this matrix that the agents in this chapter are conceived.

In terms of Carley and Newell's (1994) Model Social Agent matrix, the agents in this chapter are boundedly-rational (in terms of information processing capability) in the social setting of Real-Time Interactive Situations. The agents in this chapter exchange information in real-time in a social setting where multiple agents exist, but no cultural or social goals exist, and no social structure is imposed. However, social sanctions are defined relative to the functioning of the agents in the group – the reduction in the provision of information. In this chapter a third-dimension is added that defines elements of the task (see Carley & Prietula, 1994). The specific architecture of the agents supports the requisite elements for knowledge required to operate the social situation for the task at hand as defined by the matrix.

Trust. A component of the social situation defining a type of knowledge required for these agents involves a model of trust. Trust, in the restricted sense used in the chapter, refers to a theoretical framework from Rempel, Holmes and Zanna (1985) that has been explored and refined in a related series of experiments on agent advice (Lerch, Prietula & Kulik, 1997; Lerch, Prietula & Kim, 2000). This framework emphasizes the multidimensional nature of trust and suggests how dimensions relate to each other. The results suggest that trust has at least three dimensions (predictability, dependability and faith), but different types of informational messages and experience differentially impact the underlying dimensions. A judgment of *predictability* can be made whenever one observes consistent behavior over time in a stable environment. With this dimension of trust it is not necessary to "interpret" the person's behavior or to make attributions about the person. *Dependability* which reflects the most common definition of trust and is based largely on instances in which environmental factors are thought to be an insufficient explanation for the person's behavior. In Dependability, the person's behavior is attributed to internal characteristics of the person. *Faith* reflects an emotional security that goes beyond the available evidence and dispositional attributions.[1] Specifically, external attributions about an agent's performance only have an impact on predictability and do not influence the other two dimensions of trust, while direct and early experiences substantially impact all dimensions (and, consequently, comparative trust levels). This implies that direct experience has a larger impact on trust overall, but indirect experience, such as gossip, influences a singular component of trust – predictability.

Furthermore, being a multi-dimensional construct, trust levels can be equivalent over varying predictability levels (if augmented by different dispositional influences on other dimensions). As a consequence, it is possible to have a general disposition toward another with respect to trust, but have a situational alteration of behavior because of a particular component of that construct (predictability). To incorporate that model, the agents in this chapter (as will be explained) alter their trust judgments via direct experience only, but can situationally alter their behavior based on gossip. Direct experience (repeated cooperation in this model) increases the likelihood of trust (Burt, 1999a). However, gossip information is discounted structurally via the trust model.

Agent Models. A plausible approximation to these simplified social agents resides the form of computational modeling (Carley & Prietula, 1994; Prietula, Carley & Gasser, 1998; Prietula & Watson, 2000). Computer models of agents, with particular knowledge, in a particular social setting, and doing a particular task are defined and allowed to interact in well-defined and constrained manners. The interest resides not in the simplicity of the agent architectures or their group, but in the complexity of the behavior that emerges from their interactions (Holland & Miller, 1991; Lomi & Larsen, 1999). Thus, rather than working from group principles derived from macro performances in order to hypothesize plausible underlying agent architectures and behaviors, the plausible underlying agent architectures and behaviors are first established, and the subsequent macro performances they generate are observed and tested under varying manipulations (Burton & Obel, 1995; Carley, 1995; Carley & Prietula, 1994; Prietula & Carley, 2001).

This is accomplished by creating a simple, well-defined task with simple, well-defined computational agents, then systematically manipulating elements of the task and the agents to examine the emerging social behaviors. Consequently, this is detailed exploratory study of a computational model of group communication and performance in order to see "what matters and why" in this model. The computational model is, in fact, also a business model – one that which incorporates gossip as a mechanism to alter the advice network (Krackhardt & Hanson, 1993).

A simple search task is simulated, where computational agents seek to find objects (e.g., items, websites, patterns) in a search space. The concept of "search" is loosely defined and can be interpreted with either a more or less physical analogy (e.g., seeking items over a geographical area, or seeking relevant web sites on the Internet). For the agents in this chapter, the objects they seek are simply integers scattered around a block-metric set of locations. They

---

[1] Faith captures the essence of trust that is not securely rooted in past experience and is not incorporated

have a list of integers they are to locate, so they go about their task of searching the block-metric space for those integers.

The first set of simulations examines baseline behaviors of agents in the task without communication. These simulations hold the task constant, and vary the number of agents engaged on the task, as well as two key agent properties, where each agent engages in random searches for their objects, but does not communicate with other agents. In addition, beyond baselines, these simulations address versions of a question proposed by Cohen (1990) who investigated single versus multi-actor agents in a learning task – when and how do multi-actors do better than single actors on a learning task?

The second set of simulations elaborate on this baseline data by adding agent communication in the form of advice on object locations. For this set of simulations, a basic model of information sharing is used: if the agent can answer a request for advice, it answers. Of course, many studies have investigated the role of group information sharing under various manipulations and conditions (e.g., for a review, see Argote, Gruenfeld & Naquin, in press). The work in this chapter does not include the general factors typically explored in this type of research, as the current agent models are too simplistic to include the elements investigated using “real” groups, such as face-to-face meetings (Olivera & Argote, in press), shared ideas (Wittenbaum & Stasser, 1996), leadership (Larson, Christensen, Abbot & Franz, 1996), social loafing (Karau & Williams, 1993), gossip (Burt & Knez, 1997), or diversity (Williams & O’Reilly, 1998).

The third set of simulations broadens the agent architecture and adds simple trust models that vary in their tolerance for bad advice. Bad advice is generated by randomly disrupting the environment, so that advice provided by agents is invalidated. The results of the interactions of the factors are presented as a set of observations concerning impacts the factors on three dependent variables measuring organizational *effectiveness* (percent of the task that was completed) and *efficiency* (time to complete the task, total effort required by the group).

The fourth set of simulations keep the environment stable, but add agents that are purposefully deceptive, and provide false advice on locations. Thus, this examined how the three observations previously defined would hold over a different source of advice uncertainty.

---

in the model for these agents.

The final set of simulations adds the capability of generating and spreading gossip about particular agents as credible sources of advice, and examined the impact of gossip on the prior observations. As will be seen, gossip involve several specific decisions for agents that must be addressed by additional agent component models. The chapter concludes by interpreting the results and discussing the limits and generalities of the model and the study. An Appendix is provided with the descriptive details of the Agent models used.

TrustMe: A Simulation Model

A simulation model, called TrustMe, was written that consisted of a set of one or more *agents*, where an agent would search in a search space (an abstract set of 144 locations defined in a simple 12 x 12 block model) for a particular item pattern (an integer) located at a particular location. The default (base) search strategy for an agent is a random walk through the search space. When an agent finds an item, it will then search for the next item on its list, because in this set of simulations, an given agent searches for items on its list sequentially, not in parallel. Each agent, however, does have a *location memory* where the agent can recall the last *n* locations visited and the items located there at the time (where *n* = 14, or 10% of the search space, for these simulations). Furthermore, each agent has a *search limit* that places an upper bound on the number of locations that can be searched for any given item (where the search limit for these agents is 100% of the search space, 144 locations). If the search limit is exceeded, the agent "gives up" and proceeds to search for its next item.

Agent Parameters

The phenomena of interest defined in these simulations centers around information and communication – information about the task and agents in the task environment (Newell and Simon, 1972), and the communication flow that distributes or influences that information. Specifically, agents can provide *advice* about the location of items in the search space as well as advice on the experienced trustworthiness other agents.

Agent behavior is initially described according to two components: a trust model and a deceptiveness property. An agent's *trust model* defines how experience and advice influences an agent's propensity to provide or listen to advice from another agent. *Deceptiveness* reflects the propensity of an agent to provide deceptive advice (wrong item location information) in response to a request for information.

Task Parameters

The search task is defined by properties of the task as well as properties of the agent group. The set of parameters that define the task are: the total *number of agents*, the total *number of items* to be found (defining the number of items per agent), and the *stability* of the environment (i.e., location of the items). The set of parameters of the agent group are: the number of deceptive agents in the group and the specific trust models in the group.

Simulation 1: Independent Baselines

The first simulation determined the baseline behavior of agents acting without communication. For this simulation, the number of agents on the task was varied (1, 10, 20, 50, and 100) while holding the total number of items constant at 100, thus varying the number of items sought per agent (100, 10, 5, 2, and 1). In this simulation, no agents communicated with other agents, so no issues of trust or deception arose. Essentially, the baseline describes how groups of agents, working individually and in parallel, solve the organizational task.

One hundred replications were run for each condition with the initial placement of the agents and items randomized for each replication. Three dependent variables were examined. Two were measures of *task efficiency*: task time (where the time unit is the maximum agent visits required to complete the task, assuming all agents work in parallel) and organizational effort (in terms of total agent-visits). One was a measure of *task effectiveness*: items found (expressed as a percentage of total).

Results

The results are summarized in the first section (Constrained) of Table 1. Increasing the number of agents on the task can significantly reduce task time, but this also increases the total organizational effort. The average agent visit (effort) is reduced, as the total number of items to be found is held constant.

Insert Table 1 about here

However, note that as more agents participate in the task, less of the total task (items found) is actually completed. Adding agents reveal a hidden cost to the organizational task – the task terminates earlier, but actually less of the task gets accomplished. This cost exists because the agents simply fail to find the items under the search constraint of the agent architecture.

Recall that there are two components of the agent architecture that bound the search performance of the agent: location memory and search limit. Therefore, relaxing these constraints should reveal the nature of the impacts of these constraints on individual, and group, performance.

This, in fact, does occur. The results of a second set of 100 simulations where the agents' location memory was increased from recalling the content of the 14 past locations visited to recalling the content of the 144 past locations visited, and the agents' search limit was increased from 144 to 1000 visits. As can be seen from the second (Relaxed) section of Table 1, these two architectural modifications allowed the agents to complete the organizational task successfully (items found), but at the cost of increased task (task time), and collective organizational resources (organizational effort).

Which of the two architectural constraints of the agents accounted for most of the effect? A 2 x 2 ANOVA experiment was conducted on the baseline agent architecture, where two levels of location memory (original: 14, relaxed: 144) were crossed with two levels of search limit (original: 144, relaxed: 1000) and pooled across agent group size. A Tukey HSD post-hoc analysis revealed that *increasing the search limit* was the dominant factor in all dependent variable results, including average percent of items found (indicated by SL in the third section of Table 1). Additionally, significant search limit by location memory interactions were found (indicated by SL x LM in the third section of Table 1). In all three cases, although search limit was the major factor, location memory had a significantly greater effect on search limit when the search limit was 1000, in reducing total task time, reducing organizational effort, and increasing the percentage of items found.

The results of the baseline simulations establish reference levels of behavior for the task *sans* communication impacts and indicate that, for this task, there are two agent architectural components, search limit and location memory, that contribute both independently and, in some cases collectively, to individual and organizational performance. Overall, for this task *search limit imposed by the agent's architecture mattered more than the memory limit.*

For the remaining simulations described in this chapter the more restrictive constraints (search limit: 144, location memory: 14) were used in order to define a boundedly rational agent for the task in terms of the computational architecture's performance on a task (Simon, 1979). The subsequent manipulations explore some increasingly complex conditions under which communication among agents can (possibly) reduce task effort and offset their individual architectural constraints by altering their collective social behaviors.

Simulation 2: Advice Baselines

The next set of baseline simulations explored the impact of *advice* on the task in the context of alternative two group sizes. The two sizes of agent groups were the following: 20 agent groups (5 items per agent) and 50 agent groups (2 items per agent). Advice, of course, would facilitate both group and individual performance. Two group sizes were selected as they represented a tradeoff between the number of agents on the task (more agents benefiting the group) and communication densities over time (more items per agent imply more opportunities for communication and the development of a social history).

The primary manipulation for this simulation was the role that communication could play among agents, where communication consisted of solicitations and provisions of advice. Advice consisted of a possible location for an item pattern and was obtained and generated as follows. All agents had a preference to ask other agents for advice on the possible location of the item that they were currently seeking. A given agent A_i sent a message to all other agents requesting location information for some item, I_j. A_i always began with its default search method, random walk, but before it made a move it checked for a response. If some agent A_k responded, A_i accepted the advice and proceeded to the location. If multiple agents responded, A_i randomly selected the advice.

Results

All agents trusted other agents by default, and there were no deceptive agents. One hundred replications were run and the results of the primary dependent variables are shown in last section, Constrained (Advice), of Table 1. An analysis of variance revealed the significance of the effects between group size. In general, the group with the larger numbers of (all communicating) agents benefits – the task takes less time, with less organizational effort, but no significant difference in percentage of items found (both groups found over 97% of the items). Further, comparing these to the Constrained results in the first section of the table, the advantage of communication and advice in this task is clearly demonstrated. Overall for the task, *larger groups cooperating via communication perform better than non-cooperating groups, and larger groups are more resource efficient than smaller groups*.

The primary mechanism for this success can be found by analyzing the nature of the communications themselves. An analysis of the advice revealed that the small and large groups averaged 90.3 and 92.8 good advice communications respectively with no bad advice; therefore, these agents solved over 90% of the task on the basis of advice from other agents. Individuals and the group both benefited.

When Good Advice Goes Bad

As more communication ensues on the task (i.e., agents seeking and providing advice), a simple "social history" emerges among agents based on their communication experiences. Accordingly, each agent was given a *social memory* structure where the results of the advice from other agents were retained. Using social memories, each agent realized a "trusted information structure" that defined its advice network (Krackhardt & Hanson, 1993) – a list of agents whose trust level exceeded the threshold defined by its particular trust model. As Wellman (1997) notes, "The Internet encourages specialized relationships because it supports a market approach to finding social resources through online relationships" (p. 197). The specialized relationships here are the context-specific trusted information structures defining trustworthy sources of advice for the task at hand. In the sense used here, trust refers to an agent's "belief" that another agent will fulfill its obligation to provide truthful advice in its network exchange relation. This embodies a simple, but basic, concept of focusing on the likelihood of future events (Blau, 1964; Haas & Deseran, 1981) and is qualitatively related to "reputation effects" in economics (Axelrod, 1984; Kreps, Milgrom, Roberts & Wilson, 1982). In Simulation 2, all of these encounters were positive, as the advice presented was accurate. However, what happens when advice is not so accurate?

To explore this, the accuracy of the advice was manipulated in two ways: the accommodation of the task (to be stable) and the intentions of the agents (to be truthful).

Simulation 3: Good Agents, Turbulent Environment.

The accuracy of the advice of the agents in this simulation was caused by keeping the agents truthful, but disrupting the stability of the task environment. The environment of the task was disrupted in the following manner. At random times during the simulation, the location of each item was moved from where it was to a new, random location (The average number of locations disrupted per simulation run was 3,444). This, in effect, negated some of the advantage of the agents' location memory and could have two impacts, depending on the specific state of the agent. First, it could cause the agent's memory-based search for an item to fail, so an agent proceeded to some location would not find it there. Second, it could cause the advice provided by an agent (based on its memory of an item location) to be false.

Advice was obtained and generated as was described in Simulation 2 with the following difference. If some agent A_k responded to a request for advice from agent A_i, A_i decided whether to accept the advice or not, based on its *trust model*. Furthermore, if some agent A_k received a request for advice, it chose to respond or not based on the same trust model. Thus, an agent's trust model defined how an agent determined whether or not to provide advice or to trust the advice of another agent.

As the impact of the trust model could be significant, two simple, but quite different, trust models were defined in terms of state transitions based on the quality of advice (good, bad) sequences received (as revealed by the social memory) from other agents.

The first trust model was an *unforgiving* model that viewed all agents as initially trustworthy; however, it took only one piece of bad advice for an agent to be judged as untrustworthy. When an agent received bad advice (i.e., it searched at a particular location based on another agent's communication), it reposted the request for advice back to the group. Both models reflect different extents of generalized reciprocity that seems to underlie online group support norms (Wellman, 1997).

The second trust model was a *forgiving* model that similarly viewed all agents as trustworthy. However, forgiving agents required four incidents of bad advice to make a judgment of untrustworthy. For both models, once an agent was deemed untrustworthy, that judgment could not be redeemed.

There were two behavioral implications for a judgment of untrustworthy: (1) an agent that received advice from an untrustworthy agent would not accept that advice, and (2) an agent that received a request for advice from an untrustworthy agent would elect not to provide advice to that agent (assuming it knew the answer).

Trustworthiness, then, involved agent-specific judgments based on direct experience (the result of acting on advice) and thus influenced social activity (to accept or to provide advice) in the task. The questions of how trust models would underlie individual and group performances, and how those performances would vary in turbulent task environments were addressed in this simulation.

A 2 x 2 x 2 analysis of variance simulation study was conducted that varied the Group Size (small = 20 agents, large = 50 agents), Trust Model (unforgiving, forgiving), and Environmental Stability (stable, turbulent). One hundred replications were run per condition with the initial placement of the items and agent randomized.

Results

The results of the analysis are summarized in Table 2. The first sections of Table 2 present the main effects and 2-way interactions for the analysis, while the third section (last three rows) presents the 3-way interaction results.[2]

Insert Table 2 about here

Main Effects. Overall, there was a significant *group size effect* (A), where the larger number of agents on the task resulted in shorter task times, less organizational effort, and a greater percentage of items found. There was a significant *trust model effect* (M), where groups comprised of more forgiving agents completed the task in less time, with less organizational effort, and found a higher percentage of items. There was also a significant *environmental stability effect* (E), where turbulent environments resulted in significantly longer task times, increased organizational effort, and a lower percentage of items found.

Consequently (and intuitively), larger and more forgiving groups are more effective and efficient at the task, while environmental turbulence can significantly disrupt these efforts. But do these effects hold over the varying conditions of interaction? Although the main effects have been reviewed individually under environmental turbulence, it is important to the entire story of the behaviors of the agents to determine if the main effects or interactions regarding those main effects — Group Size, A by Trust Model, M — are preserved or altered over the environmental conditions varied (stable, turbulent) in the simulation set.

[2]In Table 2, significant 2-way interactions are represented as $Factor_1.direction \rightarrow Factor_2$, which describes how $Factor_1$ varies (.inc = increases, .dec = decreases) over the conditions of $Factor_2$, where $Factor_i$ are the manipulations of agent group size (A), trust models (M), or environmental stability (E). For example, consider the agent group size (A) by trust model (M) interaction effect on Agent Effort in Table 2. The entry M.dec → A

Group Size x Trust Model x Environmental Stability. The A x M x E row of Table 2 indicates whether or not the interactions were significant across environmental conditions and, if so, the nature of the significant interaction. The entries refer to the significance of the A x M interactions in the "stable → turbulent" conditions. Thus, for example, there is a significant 3-way interaction (A x M x E) regarding task time. The entry "*ns* → sig" indicates that the A x M interaction is not significant under stable conditions, but becomes significant under turbulent conditions. The particular nature of interaction is discussed below. The last two rows represent the dominant effect(s), if any, accounting for the best score (as determined by Tukey HSD analyses). These results are also summarized verbally in the first row of Table 3.

Insert Table 3 about here

Task time. There is no A x M interaction with respect to total task time under stable conditions, but under turbulent environments that interaction was significant, where forgiving agents (in the smaller group) had significantly lower task times than unforgiving agents in the same condition. The larger group obtained the minimum task times across models and environmental conditions (Best Stable and Best Turbulent both A_L).

Organizational Effort. The significant A x M interaction with respect to organizational effort occurred under turbulent conditions, and was attenuated (but still significant) under stable conditions. Total organizational effort was minimized across environmental conditions by using larger agent groups of forgiving agents (Best Stable and Best Turbulent both $A_L M_F$).

Items Found. The significant A x M interaction with respect to percent of items found occurred under turbulent conditions, but disappeared under stable conditions. The percent of items found under stable environmental conditions was not sensitive to any of the manipulations (all yield equivalent performance); however, under turbulent conditions the percent of items found was maximized by a larger group of forgiving agents (Best Stable, *ns*; Best Turbulent $A_L M_F$).

---

indicates that the main effect of trust model, M (favoring forgiving agents over unforgiving agents), is strongest at the smaller group size of A, but that effect declines (M.dec) under the larger group size of A.

Simulation 4: Deceptive Agents, Stable Environments.

The accuracy of advice in Simulation 3 was manipulated by altering the stability of the task environment, thus reducing the accuracy of the (well intentioned) statements of the agents providing advice. Simulation 4 replicates the Group Size and Trust components, but replaces the Environmental Turbulence manipulation with a Deception manipulation. Specifically, modifying the intentions of some of the participating agents to be deceptive was the cause of the uncertainty of the advice. A deceptive agent was architecturally and behaviorally equivalent to the prior agents with one modification – they would supply incorrect advice (i.e., locations of an item) at every opportunity (i.e., request from another agent).

Deception in this model can be interpreted as a betrayal of expectations, as all agents in all conditions have a default presumption of honest cooperation from other agents. Elangovan and Shapiro (1998) make distinctions among *accidental* betrayal (betrayal event, but absence of intent), *intentional* betrayal (betrayal event, intent present), and *opportunistic* betrayal (betrayal event, intent present but arises situationally and is justified in context). For the model in this chapter, accidental betrayal can arise when honest agents provided advice in turbulent tasks (Simulation 3). Intentional betrayal occurs when the default cooperation choice of an agent is deception (Simulation 4). Opportunistic betrayal occurs as an emotional choice to change from an honest to a deceptive agent under certain trust events (not included in this chapter).

The manipulation to the group was made on two levels: no deceptive agents and 50% deceptive agents. The same levels of Group Size (20, 50 agents) and Trust Model (forgiving, unforgiving) were crossed with the two levels of Deception. For each cell, one hundred replications were run with initial placement randomization as previously described.

Results

The results are summarized in Table 4. The interactions are read the same way as in Table 2, with the differences in findings between the tables indicated by a delta (Δ).

Insert Table 4 about here

Main Effects. As can be seen from the table, the overall effects of number of agents (A) and trust model (M) generally replicated those found in prior Simulation 3. Larger groups of agents (A) took less time to complete the task, with less organizational effort, and found a higher percentage of items. Differing from the prior results, there was no main effect for trust model (M) on time to complete the task. Similar to the prior results, groups of more forgiving agents took less organizational effort and found a higher percentage of items. The deception manipulation (D) revealed significant main effects for all dependent variables and directionally paralleled those found in Environmental Stability. That is, groups comprised of 50% deceptive agents took significantly more time to complete the task, with more organizational effort, and found a lower percentage of the items.

Group Size x Trust Model x Deception. This is the most important component of the analysis and is compared with those findings of the Environmental Stability of the prior sections. Similarly, A x M x D results are described in the bottom rows of Table 4. The results are summarized verbally in the second row of Table 3 and the details of the interactions are discussed below.

Task time. Differing from the prior results, there was no A x M interaction with respect to task time, as only group size accounted for lower task times. Larger agent groups obtained the minimum task times across deceptive conditions (Best Honest and Best Deceptive both A_L).

Organizational Effort. The significant A x M interaction with respect to organizational effort occurred under turbulent conditions, but disappeared under stable conditions. Organizational effort was minimized by using larger groups when the agents were honest (trust model did not matter); with deceptive agents, organizational effort was minimized with larger groups of forgiving agents (Best Honest, A_L; Best Deceptive, A_LM_F).

Items Found. Similar to the results of Environmental Stability, a significant A x M interaction with respect to percent of items found (as well as increases in both A and M main effects) occurred with deceptive agents, but disappeared when all of the agents were honest. The percent of items found using honest agents was not influenced by either group or trust model (all yielded equivalent performances); however, with deceptive agents, the percent of items found was maximized by larger groups of forgiving agents (Best Honest, *ns*; Best Deceptive A_LM_F).

Discussion

The systematic manipulation of an artificial set of agents on an abstract search task was conducted to obtain specific baseline behaviors arising from agent communication and parallel task execution. The certainty of the advice was then manipulated by introducing environmental turbulence and agent deception, in the presence of two simple trust models that varied in their tolerance for bad advice. For each source of uncertainty (turbulence or deception), the results were cast as a set of observations, one for each of the dependent variables measuring effectiveness (percent of items found) and efficiency (task time, organizational effort). As the results indicate, the observations based on the two different sources of advice uncertainty were quite similar.

Regarding the measure of *effectiveness*, it is apparent that in more certain environments (stable locations, honest agents), neither the trust model nor the number of agents on the task (group size) mattered. In these cases, events did not unfold that would engage differentiating behaviors from the individual agents. However, the uncertainty of the advice increased (either by location disruption or agent deception), there was a significant interaction between the two factors of group size and trust model —larger groups of forgiving agents were the most effective at the task.

Regarding the measures of *efficiency*, both sources of advice uncertainty resulted in the same conclusion with respect to task time — group size mattered most in reducing task time. Similarly, both sources of advice uncertainty resulted in the same conclusions regarding organizational effort — larger groups of forgiving agents were the most efficient in reducing organizational effort. With honest agents, only group size mattered.

The measures of effectiveness and efficiency indicate the strong influence of larger groups (over the trust models selected) and the conditional influence of forgiving agents. Why do larger groups and forgiving agents matter when advice is uncertain?

The reasons underlying the performances can be found by examining the role of advice, the stability of the advice network, and the likelihood of ignoring advice. Consider the turbulence of the environment. As noted, turbulent environments invalidate the accuracy of the agents' location memory and transform agents into inadvertent liars – but not all of the agents and not all of the time. There is good advice mixed in with the bad. From the perspective of the agent receiving the advice, turbulent environments and deceptive agents are not that different.

Individual properties (trust model) and organizational properties (group size) interact to determine the performance of the group. Forgiving agents are more tolerant to bad advice; therefore, they are likely to receive more bad advice than less tolerant agents. However, forgiving agents may also receive more good advice (if it is available). Smaller groups of agents have more opportunities to interact (communicate) than larger groups of agents and, consequently, develop more of a social history that is affected by the particular trust model, which perturbs the trusted information structures to a greater extent.

The relative value of the trust model (or impact of the group or source of certainty) can be seen by comparing the "Best Stable, Best Honest" and "Best Turbulent, Best Deceptive" values shown in the last two rows of Tables 2 and 4. Deceptive agents in a group are more disruptive (in this task) than a turbulent environment. There is a mechanism that may address the issues of deception and uncertainty in advice from agents, gossip.

Simulation 5: Gossip – Rumors of Uncertainty

Gossip was defined in this chapter as information conveyed about the quality of advice from another source. This is a slightly restricted version and intentionally so, to begin to establish additional baselines for communication and performance on this artificial task. However, implementing a gossip mechanism is slightly more complicated than it would first appear. Accordingly, the architecture of the agent was extended via three component models supporting gossip.

First, an agent's *assertion model* defined under what circumstances an agent would initiate gossip in the form of a rumor about the deceptiveness of another agent. Thus, gossip was the process and rumors were the objects that realized (i.e., were instances of) the gossip process. For the agents in this simulation, gossip was only based on direct experience; that is, an agent could not initiate gossip unless it had received bad advice.[3] Gossip was initiated as an instantiated rumor posted to a globally accessible (viewable) *rumor-board* as a 3-tuple, <rumored-agent, rumor-count, asserting-agent>, and all other agents could check the number of rumors asserted (rumor-count) about any given agent (rumored-agent) and which agent initiated the rumor (asserting-agent).

Second, an agent's *belief model* described the circumstances under which the agent believed the gossip. This was a *count-based* mechanism where each agent had some integer limit of rumors above which "gossip mattered" for that given agent as reported on the rumor-board. All agents in this simulation incorporated a count-based mechanism with the limit being set at one rumor.

[3] This, of course, is a simplification of the common contagion view of gossip as being routed through several intermediary sources ("A heard it from B, who heard it from C,…"). Recall that the specific form of the simulated

Finally, an agent's *impact model* described the effect of gossip on an agent's behavior and trust mechanisms. For the agents in this simulation, when gossip was believed (according to the belief model) about some agent A_i , then any advice supplied by A_i was ignored. In addition, if there was gossip about some agent A_i , then requests for advice from that agent were ignored. The impact model did not allow gossip to directly impact the trust model (as described in the introductory paragraphs), but it did impact behavior (i.e., ignoring of advice). As a consequence, rumors can alter behavior while maintaining levels of trust.

How would gossip impact task effectiveness and efficiency in turbulent environments? How would gossip impact task effectiveness and efficiency in deceptive environments? A simulation study was conducted exploring the impact of gossip on the results of the 3-way interactions (thus testing how gossip augments uncertainty conditions) described in the Simulation 4 and Simulation 5 sections, regarding the measures of task effectiveness (percent items found) and task efficiency (total task time and organizational effort). Two sets of simulations were run. The first set replicated the Group Size by Trust Model by Environment Stability study of Simulation 3 (A x M x E) in both Gossip and non-Gossip conditions, and each of the cells had 100 simulations. The second similarly replicated the Group Size by Trust Model by Deception study of Simulation 4 (A x M x D) in both Gossip and non-Gossip conditions, and each of the cells had 100 simulations.

Results

The results of the analysis are summarized in Table 5. Table 5 includes the best prior results (the interactions or main effects that account for the best scores, as given in the last two rows of Tables 2 and 4) and the mean scores for comparative purposes. Finally, the conditions with the best performing values (for task time, organizational effort, and percent of items found) under the pooled sources of uncertainty conditions (Environmental Stability, Deception) are listed for Gossip and non-Gossip conditions, as determined by a post-hoc Tukey HSD analysis (last two rows of the table). These results are verbally summarized in the third row of Table 3.

Insert Table 5 about here

----

task was a chat room, where gossip is essentially a public event, but is checked only under specific conditions

As can be seen from the Table 5, the main effects and interactions without gossip are generally preserved under gossip conditions; only the score is altered and, in most conditions, worsens. Gossip helps (and mildly) only in conditions of advice certainty (stable locations, honest agents) in reducing task time and organizational effort (which, in fact, are the best scores under both uncertainty conditions).

There is an overall *negative* gossip effect across environmental stability conditions where gossip increases task time, increases organizational effort, and decreases the percent of items found. There is also an overall *negative* gossip effect across deception conditions where rumor increases organizational effort and decreases the percent of items found but, unlike the prior results, does not significantly alter task time.

Consider the impact of gossip on *task time* across uncertainty conditions (column 1 in Table 3). When using gossip, the general results mitigate the differences obtained with trust models under turbulent conditions and are similar (in effect) with deceptive conditions without gossip. Next, consider *organizational effort* across uncertainty conditions and in comparison with the prior observations (column 2 in Table 3). When using gossip, the overall effect of trust models is mitigated across environmental stability conditions and deception conditions, where the group size effect is also mitigated. Finally, consider the *percent of items found* across uncertainty conditions and in comparison with the prior observations (column 3 in Table 3). Gossip did <u>not</u> alter the results of either the environmental stability conditions or the deception conditions.

Conclusions

In general, two comments can be made regarding gossip on this task. First, gossip generally made things worse in terms of both efficiency and effectiveness measures used when advice is uncertain. Agent groups using gossip took *more* time to complete the task, took *more* collective organizational effort to complete the task, and completed *less* of the total task than agents not using rumors. Second, gossip had an impact on three of the six original overall results (as defined in the first two rows of Table 3) where the impact involved either attenuating or removing the effects of trust model differences under conditions of uncertain advice.

---

related to the acceptance or provision of advice.

However, there seemed to be a similar pattern of the impact of gossip not only in the three results, but also across all six results – the differing effects of the trust models diminish under conditions of advice uncertainty when gossip is used. Subsequent analysis supported this hunch and revealed the nature of the pattern. Most of the underlying effect was caused by the more significant changes in the forgiving model with small groups, moving from certain to uncertain conditions with gossip (recall that larger groups were typically dominated by the group size effect). Consequently, *gossip (in uncertain advice conditions) reduces performance differences afforded by the different (forgiving, unforgiving) trust models.*

Thus, gossip somehow induced the differences between trust models (as determined by the effectiveness and efficiency measures) to diminish in conditions of uncertain advice. What about the nature (direction) of those differences? Recall that, ignoring other effects, in this task forgiving trust models generally do significantly better than unforgiving trust models on all three organizational measures. A post-hoc analysis indicated that most of the movement could be accounted for by values of the forgiving trust models moving toward the values of the unforgiving models. Therefore, under conditions of uncertain advice, *gossip "causes" the groups of forgiving agents to degrade their performance levels to those obtained by groups comprised of unforgiving agents.*

How does this occur? The answer actually lies not in the source of the advice uncertainty (deception or location disruption), but in the behavior of the trust and gossip models in *response* to that uncertainty. This is revealed by examining the underlying communications of the agents and how those communications interact with agent and task properties. Accordingly, the sources of the uncertainty were pooled for the analysis. Significant results of this analysis are all $p < .001$.

First, consider the total amount of communication in the task under (combined) advice uncertainty conditions, <u>without</u> gossip. As one might expect, forgiving agent groups have significantly higher communication levels than less forgiving agent groups, with larger groups of forgiving agents having the highest communication levels. Their forgiving nature allows them to receive more good advice as well as more bad advice, because these agents ignore significantly less advice overall. Why advice is ignored less is directly traceable to the components of their trust model that allow significantly less conflict in their trusted information structures. They have a very resilient advice network.

Gossip, as defined in this chapter, serves as a basic “defense mechanism” against bad advice. Bad advice generally leads to additional search, both individualistically (by definition, agents that received bad advice did unnecessary search) and cooperatively (gossip about deceptive agents to the coalition causes them to ignore advice and, possibly, search via random walk lacking other advice).

Next, consider the impact of adding gossip to the combined uncertainty conditions and analyzing the resulting interactions. What happens when gossip enters the mix of uncertainty? First, there is a significant drop in the overall communication level with the forgiving agents and the differences in communication levels between the two trust models diminish. This has a secondary impact as forgiving agent values approach unforgiving agent values in the amount of good advice received, bad advice received, advice taken, and advice ignored. Essentially, the use of gossip resulted in the two trust models to communicate (i.e., behave) quite similarly.

It is also insightful to examine the reason for the agents ignoring advice under gossip conditions. Gossip, as noted, originates in direct experience, but has the collective effect of quashing a source of advice and inhibiting the alteration (via rumors) other trust judgments. In a sense, gossip serves as a buffer against direct bad experiences (via following bad advice) and, consequently, as a buffer against changes in trust judgments (and trusted information structures in the advice network). If an agent A_i did not follow bad advice from some trusted agent A_j (because of gossip from some other trusted agent, A_k), then it could not alter its basic trustworthiness of A_j. As Hardin (undated) puts it, “A consequence of distrust is that, if I generally distrust people, I am likely to take few risk of cooperating with others and I will therefore acquire little information about their trustworthiness” (p. 2). In this model Hardin’s “general distrust” (i.e., not experientially based) is similar to the decisions not to cooperate by belief in gossip.

For both forgiving and unforgiving agents, gossip inhibited agents from having bad experiences directly. However, as the two trust models vary in their tolerance for bad advice, this variance (in the conditions of the task examined) resulted in different reasons for ignoring advice. Although the overall levels diminish in the uncertain conditions (because of gossip), unforgiving agents ignore significantly more advice because of trust judgments than forgiving agents (see Figure 1). Therefore, the overall levels of conflict (explicit distrust) are reduced by gossip (trust judgments are being inhibited), but forgiving agents have significantly more of a reduction (implying a more intact trusted information structure defining their advice network) than unforgiving agents (see Figure 2). Consequently, *forgiving agents ignore advice because primarily because of gossip while unforgiving agents ignore advice primarily because of trust judgments.*

Insert Figures 1 and 2 about here

The implication of this is that the advice coalition is preserved at the expense of organizational performance. The agent groups collectively have an organizational memory of 280 (with 20 agents) and 700 (with 50 agents) "slots" respectively. Organizational memory and knowledge are critical components organizational performance, but have many forms and interpretations (Argote, 1999). Withholding information was a component of both the trust model and gossip models that reflected individual behavioral responses to bad advice about a location (direct experience) or responses to communications about another agent's encounters (indirect experience). However functional as a tit-for-tat individual strategy, it has organizational ramifications. Overall, *gossip causes more information to be withheld in the accomplishment of the task over combined uncertainty conditions*.

The summary implications are presented in Table 6.

Insert Table 6 about here

Implications and Thoughts…

In this chapter a small exploratory study was conducted as a series of increasingly elaborated computer simulations to explore the effects of varying individual and task properties on organizational measures of effectiveness and efficiency on an artificial task where cooperation (in the form of advice) mattered. The study was crafted as a set of computational agents, where agents varied in their propensity to trust the advice from other agents (Do trust models matter?) and their propensity to generate or accept gossip in the form of rumors (Does gossip about other agent's matter?). In addition, the certainty of the advice available to these agents was varied by manipulating their propensity to provide false advice (Does deception matter?) and by altering the stability of the task environment (Does environmental turbulence matter?). Accordingly, the detailed underlying behaviors of the agents on the task were assessed (How do these things matter?).

The generality of the results from this chapter is restricted, necessarily, to the context of the computational study and the extent to which it captures the abstractions of relevance. The context of the study was also simple, but a reasonable (albeit abstract) interpretation of events unfolding in the real world. Imagine a set of agents that can do a task individually (e.g., find something on the Internet, purchase an item) but may benefit from advice of others (e.g., a possible website). Additionally, imagine that these agents do not know each other personally, but are brought together opportunistically (e.g., a chat room, response from a posted corporate bulletin board query). Also imagine that there is a collective interpretation to the task; that is, the task, although individually defined, has a finite beginning, ending, and membership (the number of agents participating), thus has meaningful and derivable metrics of performance (effectiveness, efficiency). Finally, imagine that the only communications that ensue among agents are simple messages (email).

The results of the simulations revealed the conditions under which the architectural constraints of the agent (as a boundedly-rational participant) interacted with agent properties (trust model, gossip model) and advice uncertainty. The baseline simulations revealed that without cooperation (communication), adding agents to the task reduced one aspect of task efficiency (task time), but increased another (organizational effort) and reduced the effectiveness of the task (percent items found). Under these same conditions, cooperation among agents increased all efficiency and effectiveness measures.

Adding gossip to the situations generally made things worse and surely made things more complicated for the agents (i.e., the integration of that type of behavior and information in their trust and behavioral models). Recall that gossip was only about agents that had provided bad advice, and gossip was generated, as rumors, only through direct experience. Intuitively, gossip should serve to help the group – agents are sharing their knowledge of a "faulty part" to be avoided by the others in the group. However, this raises a key issue of differentiating between what gossip <u>is</u> (information that a source of advice is uncertain) versus what gossip <u>does</u>, or rather, what an agent <u>does</u> <u>with</u> gossip.

In this chapter, gossip (from trusted agents) resulted in the subsequent ignoring of the advice from that agent and the refusal to submit advice to that agent. This had two effects. First, it inhibited subsequent direct experience with that agent that could lead to alterations in fundamental trust judgments (as defined in this chapter). Consequently, gossip *stabilized the trusted information structure* of the group, but altered agent behavior. In a sense, this could be interpreted as preserving the coalition at the expense of performance. Second, the behavior of the agents that had a forgiving trust model *functionally approximated* that of the agents of the unforgiving model. The component gossip model yielded behaviors in the forgiving agents that, to an "outside observer," looked exactly like unforgiving agents. Thus, by considering trust as a multidimensional construct and differentiating it situationally from over behaviors, apparent differences in effects may be more readily explained, such as the problems in resolving the "gossip" and "cohesion" predictions (c.f., Burt, 1999b).

Finally, consider one of the primary impacts of both the trust models and gossip: *withholding of information*. This could be interpreted as a metaphorical difference between declarative and procedural organizational knowledge and learning (though popular, these terms are actually quite inadequate at the organizational level). Imagine that a type of organizational declarative knowledge resides in the collective memories of the agents (where items are located) and a type of procedural knowledge (access and performance) are realized by the emergent and current states of the trust models and gossip models of the agents. As the task unfolds, the models and social memories dynamically adjust to avoid communication contact with certain agents and prefer communication contact with others, thus learning what component of the organization are more reliable than others, and acting on that knowledge. Note that the uncertainty of the environment, whether interpreted broadly as turbulence (Mintzberg, 1979) or narrowly as fitness landscapes (Levinthal, 1997), to which one has to adapt (from which one has to learn) affects organizations as it does people. Learning is the exploitation of regularities in order to reduce subsequent resource expenditures. An individual cannot learn if the environment does not afford learning in relation to the learning mechanisms available (Gibson, 1979), and that argument holds for organizations as well. This indeed demonstrates how "individual memories interact to shape the character of organizational routine" (Cohen & Bacdayan, 1994) and illustrates the critical role that individual learning plays in organizational learning (Simon, 1991) as well as how social and organizational components impact knowledge accumulation (Carley, 1986).

These findings should be considered in view of the work by Carley and Hill (2001) in this volume, as both studies explored the implications of knowledge exchange and organizational learning as derivative of agent states, behaviors and linkages. In their work, the likelihood of agent interaction is defined in terms of their relative knowledge as specified by two strategies – active and passive. This has implications, as the information shared and diffused is relevant for the accomplishment of an organizational task to be done by the agents. Although the substance of the two studies varies, the general form is similar as are the results: properties of individual agent behaviors (in their case, style of interaction) can significantly impact organizational structure, performance, and learning. Furthermore, interesting interactions occurs between group size and agent properties. It would be interesting to see how a gossip structure imposed on these agents would alter behaviors in turbulent information environments. Correspondingly, it would be interesting to modify TrustMe to perform their information diffusion task. This certainly calls for a docking of the two models (Axtell, Axelrod, Epstein & Cohen, 1996).

Consider also the work of Miller (2001) also in this volume. Miller explores how simple adaptive mechanisms can generate superior organizational structures. The random firing of nodes could be interpreted as turbulence, and turbulence impacts organizational structure and performance. The evolved (or adapted) organizational structure is a barometer to the environments it has encountered. Similarly, the advice network (though defined within an organization) in this chapter reflects an evolved (or adapted) structure reflecting the turbulence and uncertainty it has encountered. As Miller notes, the type of structure that emerges in indeed closely linked to the underlying environmental factors.

Finally, consider the contributions in this volume by Loch, Huberman and Ulku (2001). Their work presents a clever implementation and exploration of how status competition impacts performance. The inclusion of status concepts would be an interesting elaboration to TrustMe, and begins to increase the complexity, and hence the social knowledge required, of the type of Social Situation in which the agents behave (in terms of Carley and Newell's Matrix). Another docking opportunity.

A final discussion can be made with respect to TrustMe as a business model. With the emergence of the commodity Internet, general connectivity is increasing as is the advice flowing over that connectivity. eMarketer (www.emarketer.com) reports that "by the end of this year, the total Internet population in the United States, including children of all ages, will reach 80.8

million in 2000, a penetration rate of 30% of the projected U.S. population of 273 million." Over 80% of the primary reason for being online includes gathering information and using e-mail, with about 30-40+% engaging in chat rooms (which can actually include gathering information and electronic communications). People are connected to, and are receiving advice from, individuals with whom they have never met and outside of the structure of typical organizational walls and forms.

Though the model in this chapter is simple, it does afford approximations to business situations in existence. PGP Security, a Network Associates company, offers the well-known Pretty Good Privacy public-key encryption and digital signature system developed by Phil Zimmermann for e-mail systems.[4] A component of that system is a user's list of signed public-keys to which the particular user defines whether or not (and to what extent) they are trusted. Furthermore, users generate and distribute their own public key and, as use unfolds, sign each other's public keys, creating a "web of trust." Certification of keys is not done via a central authority, but through the events that unfold for a particular user's public-key ring. For example, if you know and trust Harry (in the informal sense), you can give Harry your public-key and Harry signs and returns it (keeping a copy for himself). If you want to communicate with Laura, you send a copy of your signed public-key (indicating that Harry trust you). If Laura trusts Harry, then Laura probably trusts you because she trusts Harry to certify others' keys. The term "probably" is used because it is up to the particular user to define what constitutes sufficient trust (to be extremely cautious or not). The system simply conveys the information.

Consider the online bookstore component of amazon.com. For each book, there is an opportunity to post reviews of the book in their Customer Review section. Furthermore, they have recently added the ability to "review the reviews" where site visitors can rate the helpfulness of the posted advice (i.e., "Was it helpful to you?" yes/no). The results are posted next to the review (e.g., "7 people found this review helpful.").

Lastly, consider a new Internet-enabled business model – the eBay™ auction site (www.ebay.com). eBay brings together registered buyers and sellers to bid and purchase items. However, since the buyers and the sellers do not know each other, and are essentially unknown

[4] Public-key encryption systems consist of a unique key-pair for each individual (e.g., Schneier, 1996). One key is public (i.e., held on a public directory and associated with the owners name) and is used to encrypt messages sent to

to eBay, how can the legitimacy of the transaction or the participants be judged? eBay offers a set of what they call Safe Harbor™ services such as online verification of user information (partnering with Equifax Secure, Inc.), an escrow service for payments (partnering with i-Escrow.com), basic insurance (via Lloyd's of London) with a fraud reporting system, and suggestions on hiring third-party appraisers. A final component of their Safe Harbor™ services is a chat room called Feedback Forum where members (registered buyers or sellers) can leave positive or negative reviews or comments. Members have a "feedback profile" that is simply an integer referring to the net number of positive comments, so "Sofia (112)" means that there was a net of 112 positive comments regarding transactions with Sofia. Someone who has had "direct experience" can only supply negative comments; that is, negative comments are generated only from the winning high bidder or seller in an auction. As eBay states on their website (www.ebay.com):

> **If you're a buyer**, checking a seller's Feedback Profile before you make a bid is one of the smartest and safest moves you can make. This Feedback Profile answers many questions about how a seller does business. Is she highly recommended by other buyers? Does he sell quality merchandise?
>
> **If you're a seller**, reviewing Feedback Profiles can be helpful, too. Find out if a buyer is known as a great customer who provides fast payment. Or you can also see what bidders are looking for in a good seller.

The Feedback Forum is essentially a "kindness of strangers" advice that is a formal component of the business model.

Thus, advice from strangers is not only likely to be helpful (e.g., Tandem Computers), but can be a formal component of the business model (e.g., amazon.com, eBay.com). What is interesting about the latter two business models is that they support others rating their advice or behavior, either indirectly through a reported vote and/or directly through verbal messages concerning specific sources. All of the examples are email or chat based – simply words on a screen from strangers.

This brings it all back to the Foreward. This volume addresses a type of organizational research and theorizing that incorporates computational models. In Jim March's words,

---

that person by others. The other key is private (i.e., not revealed) and is the only key that can decrypt those messages encrypted with individual's public key.

“Simulation represents an approach that appears both to match the phenomena of interest and to provide some analytical power.” This chapter, in part, addressed a type of organizational phenomena that is based, in part, on computers. That is, the computational models were the phenomena of interest. Perhaps the tool and the phenomena are merging.

<u>References</u>

Allport, G. & Postman, L. (1965). *The Psychology of Rumor.* New York, NY: Russell & Russell.

Argote, L. (1999). *Organizational Learning: Creating, Retaining and Transferring Knowledge.* Boston, MA: Kluwer.

Argote, L., Gruenfeld, D. & Naquin, C. (in press). Group learning in organizations. In M. Turner (Ed.), *Groups at Work: Advances in Theory and Research*. Mahwah, NJ: Lawrence Erlbaum.

Associated Press. (1999a). Bank sues over Internet rumors. *AP Online*, Wednesday July 28.

Associated Press. (1999b). Taiwan cites China Internet rumors. *AP Online*, Saturday, August 7.

Axelrod, R. (1984). *The Evolution of Cooperation*. New York, NY: Basic Books.

Axtell, R., Axelrod, R., Epstein, J. & Cohen, M. (1996). Aligning simulation models: A case study and results. *Computational and Mathematical Organization Theory,* 1(2), 123-142.

Blau, P. (1964). *Exchange and Power in Social Life.* New York, NY: Wiley.

Burt, R. (1999a). *Evidence of density, distrust, and character assassination*. Working Paper. Graduate School of Business, University of Chicago, Illinois.

Burt, R. (1999b). *Trust and gossip on the path to equilibrium*. Chapter Draft. Graduate School of Business, University of Chicago, Illinois.

Burt, R. & Knez, M. (1997). Trust and third-party gossip. In S. Kiesler (Ed.), *Culture of the Internet*. Mahwah, NJ: Lawrence Erlbaum.

Burton, R. & Obel, B. (1995). The validity of computational models in organizational science: From model realism to purpose of the model. *Computational and Mathematical Organization Theory*, 1(1), 57-71.

Byrne, R. (1997). Machiavellian intelligence. *Evolutionary Anthropology*, 5, 172-180

Carley, K. (1995). Computational and mathematical organization theory: Perspective and directions. *Computational and Mathematical Organization Theory*, 1(1), 39-56.

Carley, K. (1986). Knowledge acquisition as a social phenomenon. *Instructional Science,* 14, 382-438.

Carley, K. & Hill, V. (2001). Structural change and learning within organizations. In A. Lomi & E. Larsen (Eds.), *Dynamics of Organizations: Computational Modeling and Organizational Theories*, Cambridge, MA: AAAI Press/MIT Press (this volume).

Carley, K. & Newell, A. (1994). The nature of the social agent. *Journal of Mathematical Sociology*, 19(4), 221-262.

Carley, K. & Prietula, M. (1994). ACTS theory: Extending the model of bounded rationality. In K. Carley & M. Prietula (Eds.), *Computational Organization Theory*. Hillsdale, NJ: Lawrence Erlbaum.

Cohen, M. (1990). When can two heads learn better than one? Results from a computer model of organizational learning. Paper presented at the *Amsterdam Conference on Computer Modeling of Organizations*.

Cohen, M. & Bacdayan, P. (1994). Organizational routines are stored as procedural memory: Evidence from a laboratory study. *Organizational Science*, 5(4), 554-568.

Constant, D., Sproull, L. & Kiesler, S. (1997). The kindness of strangers: On the usefulness of electronic weak ties for technical advice. In S. Kiesler (Ed.), *Culture of the Internet*. Mahwah, NJ: Lawrence Erlbaum.

Doran, J. (1998). Simulating collective misbelief. *Journal of Artificial Societies and Social Simulation*, 1(1), <http://www.soc.surrey.ac.uk/JASSS/1/1/3.html>

Dunbar, R. (1996). *Grooming, Gossip, and the Evolution of Language*. Cambridge, England: Harvard University Press.

Elangovan, A. & Shapiro, D. (1998). Betrayal of trust in organizations. *Academy of Management Review*, 23(3), 547-566.

Gibson, J. J. (1979). *The Ecological Approach to Visual Perception.* Boston, MA: Houghton Mifflin.

Granovetter, M. (1973). The strength of weak ties. *American Journal of Sociology*, 78(6), 1360-1380.

Haas, D. & Deseran, F. (1981). Trust and symbolic exchange. *Social Psychology Quarterly*, 44(1), 3-13.

Hardin, R. (undated). Distrust. Paper from the *Working Group on the Construction and Maintenance of Trust*. New York, NY: Russell Sage Foundation.

Holland, J. & Miller, J. (1991). Artificial adaptive agents and economic theory. *American Economic Review, Papers and Proceedings*, 81, 365-370.

Humphrey, N. (1976). The social function of intellect. In P.G. Bateson and R. Hinde (Eds.) *Growing points in ethology*. Cambridge, MA: Cambridge University Press.

Kapferer, J-N. (1990). *Rumors: Uses, Interpretations, and Images.* New Brunswick, CT: Transaction Publishers.

Karau, S. & Williams, K. (1993). Social loafing: A meta-analytic review and theoretical integration. *Journal of Personality and Social Psychology*, 65, 681-706.

Kelley, R. & Caplan, J. (1993). How Bell Labs creates star performers. *Harvard Business Review*, July-August, 128-139.

Koenig, F. (1985). *Rumor in the Marketplace.* Dover, MA: Auburn House.

Krackhardt, D. & Hanson, J. (1993). Informal networks: The company. *Harvard Business Review*, July-August, 104-111.

Kreps, D., Milgrom, P., Roberts, J. & Wilson, R. (1982). Rational cooperation in the finitely repeated Prisoner's Dilemma. *Journal of Economic Theory*, 27, 245-252.

Kummer, H., Daston, L., Gigerenzer, G. & Silk, J. (1997). The social intelligence hypothesis. In P. Weingart, P. Richerson, S. Mitchell & S. Maasen (Eds.), *Human by Nature: Between Biology and the Social Sciences.* Hillsdale, NJ: Lawrence Erlbaum.

Larson, J., Christensen, C., Abbott, A. & Franz, T. (1996). Diagnosing groups: Charting the flow of information in medical decision-making groups. *Journal of Personality and Social Psychology*, 71, 315-330.

Lerch, J., Prietula, M. & Kulik, C. (1997). The Turing Effect: The nature of trust in expert system advice. In P. Feltovich, K. Ford & R. Hoffman (Eds.), *Expertise in Context: Human and Machine.* Cambridge, MA: AAAI/MIT Press.

Lerch, J., Prietula, M. & Kim, J. (2000). *Unraveling the Turing Effect: Measuring trust in machine advice.* Working Paper, Graduate School of Industrial Administration, Carnegie Mellon University, Pittsburgh PA.

Levinthal, D. (1997). Adaptation on rugged landscapes. *Management Sciences*, 43(7), 934-950.

Loch, C., Huberman, B. & Ülkü, S. (2001). Multi-dimensional status competition and group performance. In A. Lomi & E. Larsen (Eds.), *Dynamics of Organizations: Computational Modeling and Organizational Theories*, Cambridge, MA: AAAI Press/MIT Press (this volume).

Lomi, A. & Larsen, E. (1999). Evolutionary models of local interaction. In J. Baum & B. McKelvey (Eds.), *Variations in Organizational Science: In Honor of Donald T. Campbell*. Thousand Oaks, CA: Sage.

March, J. (2001). Foreward. In A. Lomi & E. Larsen (Eds.), *Dynamics of Organizations: Computational Modeling and Organizational Theories*, Cambridge, MA: AAAI Press/MIT Press (this volume).

Miller, J. (2001). Evolving information processing organizations. *Dynamics of Organizations: Computational Modeling and Organizational Theories*, Cambridge, MA: AAAI Press/MIT Press (this volume).

Mintzberg, H. (1979). *The Structuring of Organizations*. Englewood Cliffs, NJ: Prentice-Hall.

Myerson, D., Weick, K. & Kramer, R. (1996). Swift trust and temporary groups. In R. Kramer & T. Tyler (Eds.), *Trust in Organizations: Frontiers of Theory and Research*. Thousand Oaks, CA: Sage.

Newell, A. & Simon, H. (1972). *Human Problem Solving*. Englewood Cliffs, NJ: Prentice-Hall.

Olivera, F. & Argote, L. (in press). Organizational learning and new product development: CORE processes. In L. Thompson, D. Messick & J. Levine (Eds.), *Shared Knowledge in Organizations*. Mahwah, NJ: Lawrence Erlbaum.

Prietula, M. & Carley, K. (2001). Boundedly rational and emotional agents: Cooperation, trust, and rumor. C. Castelfranchi & Y-H. Tan (Eds.), *Deception, Fraud and Trust in Virtual Societies.* Norwell, MA: Kluwer.

Prietula, M. & Carley, K. (1999) Simulating emotion and trust in groups of artificial agents. *INFORMS Workshop on Computational and Mathematical Organization Theory*, May.

Prietula, M. & Carley, K. (1998) A Computational Model of Emotion, Trust, and Deception. *Proceedings of the 1998 AAAI Fall Symposium Series, Emotional and Intelligent: The Tangled Knot of Cognition,* Orlando FL, Oct 23-25.

Prietula, M., Carley, K. & Gasser, L. (Eds.) (1998). *Simulating Organizations: Computational Models of Institutions and Groups.* Cambridge, MA: AAAI/MIT Press.

Prietula, M. & Simon, H. (1989). The experts in your midst. *Harvard Business Review,* Jan-Feb, 120-124.

Prietula, M. & Watson, H. (2000). A behavioral theory of the firm: Organizational and economic perspectives, *Organizational Science,* 11(5), 565-585.

Rempel, J., Holmes, J. & Zanna, M. (1985).. Trust in close relationships. *Journal of Personality and Social Psychology,* 49 (1), 95-112.

Rosnow, R. & Fine, G. (1976). *Rumor and Gossip: The Social Psychology of Hearsay*. New York, NY: Elsevier.

Schneier, B. (1996). *Applied Cryptography, 2nd Edition*. New York, NY: Wiley.

Simon, H. (1991). Bounded rationality and organizational learning. *Organizational Science*, 2(1), 125-134.

Simon, H. (1979). Rational decision making in business organizations. *The American Economic Review*, 69(4), 493-513.

Wellman, B. (1997). An electronic group is a virtually social network. In S. Kiesler (Ed.), *Culture of the Internet*. Mahwah, NJ: Lawrence Erlbaum.

Whiten, A. & Byrne, R. (1997). Machiavellian Intelligence II: Extensions and Evaluations. Cambridge, England: Cambridge University Press.

Williams, K. & O'Reilly, C. (1998). Demography and diversity in organizations: A review of 40 years of research. *Research in Organizational Behavior*, 20, 77-140.

Wittenbaum, G. & Stasser, G. (1996). Management of information in small groups. In J. Nye & A. Bower (Eds.), *What's Social About Social Cognition? Research on socially shared cognition in small groups*. Thousand Oaks, CA: Sage.

Appendix

Here are pseudocode representations of the agent algorithms are presented for each of the four agent simulations.

Simulation 1: Independent Baselines. These agents do not communicate and simply search the space independently. In simulation, the values of two architectural constraints were varied, *location_memory* and *search_limit*. For some agent *i* holding a list of items and is currently at some location *k*,

```
Agent(i):
while items on the list
begin {agent gets an item}
    get_next item
    repeat  {agent searchers for item}
        case
        1. item at location k → found(item) = true
        2. item not at location k → found(item) = false
            if item in location_memory then k = item location
                                       else k = new_random_location
            move_to location k
            add items at location k to location_memory
        end case
    until (found(item) = true) or (search_count > search_limit)
end.
```

Simulation 2. Advice Baselines. The agents in this group were similar to the Simulation 1 agents, but had a preference (after checking its own *location_memory*) to ask for advice from other agents (*post_question*, *check_ advice*). Agents would also check to see if they could answer any posted questions when their *location_memory* was updated (*check_questions, post_ advice*).

```
Agent(i):
while items on the list
begin {agent gets an item}
   get_next item
   post_question(Agent_i,item)
   repeat  {agent searchers for item}
      case
      1. item at location k → found(item) = true
      2. item not at location k → found(item) = false
         if item in location_memory then k = item location
            elseif check_advice (item) = true then k = advised location
                                              else k = new_random_location
```

```
        move_to location k
        add items at location k to location_memory
        if check_questions (location_memory)= true    {Question from some Agent_j}
            then post_advice (Agent_i,item, location)
      end case
   until (found(item) = true) or (search_count > search_limit)
end.
```

<u>Simulation 3. Advice, Trust, and Turbulence</u>. The agents in this group were similar to the Simulation 2 agents, but included a trust model that was influenced (positively or negatively) by believed advice (*adjust_trust*) from another Agent_j and served as a basis for believing advice from (*check_trust*), or answering posted questions by, some other Agent_j. Agents receiving bad advice or advice from untrusted agents would repost their question (*repost_question*).

```
Agent(i):
while items on the list
begin {agent gets an item}
   get_next item
   post_question(Agent_i,item)
   repeat  {agent searchers for item}
      case
      1. item at location k → found(item) = true
         if at location k via advice from some Agent_j then adjust_trust(Agent_j)
      2. item not at location k → found(item) = false
         if at location k via advice from some Agent_j then adjust_trust(Agent_j)
                                                       repost_question(Agent_i,item)
         if item in location_memory then k = item location
            elseif check_advice (item) = true then     {Answer posted by some Agent_j}
                  if check_trust(Agent_j) = ok then k = advised location
                                              else repost_question(Agent_i,item)
                                                   k = new_random_location
         move_to location k,
         add items at location k to location_memory
         if check_questions (location_memory)= true          {Question from some Agent_j}
             then if check_trust(Agent_j) = ok then post_advice (Agent_i,item, location)
      end case
   until (found(item) = true) or (search_count > search_limit)
end.
```

<u>Simulation 4. Advice, Trust, and Deception</u>. The agents in this group were similar to the Simulation 3 agents, but a set of these agents was deceptive in the following sense. These agents would post an incorrect answer (*post_deceptive_advice*) to a question from another agent

regardless of trust levels (i.e., no *check_trust* was performed) or whether or not it actually knew the answer.

```
Deceptive_Agent(i):
while items on the list
begin {agent gets an item}
   get_next item
   post_question(Agent_i,item)
   repeat  {agent searchers for item}
      case
      1.  item at location k → found(item) = true
         if at location k via advice from some Agent_j then adjust_trust(Agent_j)
      2. item not at location k → found(item) = false
         if at location k via advice from some Agent_j then adjust_trust(Agent_j)
         if item in location_memory then k = item location
            elseif check_advice (item) = true then     {Answer posted by some Agent_j}
                     if check_trust(Agent_j) = ok then k = advised location
                                                   else k = new_random_location
         move_to location k,
         add items at location k to location_memory
         if check_questions (location_memory)= true          {Question from some Agent_j}
               then post_deceptive_advice (Agent_i,item, location)
      end case
   until (found(item) = true) or (search_count > search_limit)
end.
```

Simulation 5. Gossip. The agents in this group were similar to the Simulation 3 and 4 agents, but incorporated gossip (i.e., negative rumors about agents providing bad advice). Gossip involved asserting rumors (*assert_rumor*) and incorporating their influence in believing advice or answering questions from other agents (*check_rumors*). Deceptive agents used rumors as did honest agents, but did not check for rumors in their decision to provide deceptive advice.

```
Agent(i):
while items on the list
begin {agent gets an item}
   get_next item
   post_question(Agent_i,item)
   repeat  {agent searchers for item}
      case
      1.  item at location k → found(item) = true
         if at location k via advice from some Agent_j then adjust_trust(Agent_j)
      2. item not at location k → found(item) = false
         if at location k via advice from some Agent_j then       assert_rumor(Agent_j)
                                                                  adjust_trust(Agent_j)
```

```
                                                            repost_question(Agent_i,item)
      if item in location_memory then k = item location
        elseif check_advice (item) = true then      {Answer posted by some Agent_j}
                if (check_trust(Agent_j) = ok) and (check_rumors(Agent_j) = ok)
                    then   k = advised location
                    else   repost_question(Agent_i,item)
                           k = new_random_location
      move_to location k,
      add items at location k to location_memory
      if check_questions (location_memory)= true          {Question from some Agent_j}
           then if (check_trust(Agent_j) = ok0 and (check_rumors(Agent_j) = ok)
                then post_advice (Agent_i,item, location)
    end case
  until (found(item) = true) or (search_count > search_limit)
end.
```

Table 1.
Summary of Baseline Simulations

Number of Agents	Task Time	Organizational Effort	Items Found
Constrained			
1	8,914.88	8,914.88	66.8%
10	1,210.35	9,436.25	60.8%
20	705.89	9,995.11	54.6%
50	292.00	10,557.61	47.8%
100	146.00	10,992.42	44.0%
Relaxed			
1	15,681.63	15,681.63	100.0%
10	14,881.97	14,881.97	99.7%
20	2,201.08	17,787.41	98.6%
50	1,673.35	20,533.99	96.8%
100	1,002.00	19,941.47	95.8%
SL x LM*	SL, SL x LM	SL,SL x LM	SL
Constrained (Advice)			
20	363.72	1,405.18	97.5%
50	238.43	1,218.52	97.2%

* Number of agents pooled. All reported effects significant at $p < .001$
SL = search limit, LM = location memory

Table 2
Summary of Advice x Trust x Environmental Stability

Number of Agents	Task Time	Organizational Effort	Items Found
Agents (A)	20 > 50	20 > 50	20 < 50
Trust Model (M)	F < U	F < U	F > U
Env. Stability (E)	S < T	S < T	S > T
A x M	*ns*	M.dec→ A	M.dec→ A
A x E	E.dec→ A	E.dec→ A	E.dec→ A
M x E	M.inc→ E	M.inc→ E	M.inc→ E
A x M x E	*ns* → sig	sig → inc	*ns* → sig
Best Stable:	A_L	$A_L M_F$	*ns*
Best Turbulent:	A_L	$A_L M_F$	$A_L M_F$

All reported effects significant at $p < .01$.
F = forgiving trust model, U = unforgiving trust model
S = stable environment, T = turbulent environment
A_L = large agent group, A_S = small agent group
M_F = forgiving trust model

Table 3

Summary results of 3-way interactions

Group Size x Trust x	Total Task Time	Organizational Effort	Items Found
Environ-mental Stability	In *stable* and *turbulent* environments, larger groups (independent of trust model) perform best at reducing task time. (Trust models matter only with small groups in turbulent environments, and forgiving agents perform best)	In *stable* and *turbulent* environments, larger groups of forgiving agents perform best at reducing organizational effort.	In *stable* environments, neither trust model nor group size influences the percent of items found; in *turbulent* environments, larger groups of forgiving agents perform best.
Deceptive Agents	With *honest* and *deceptive* agents, larger groups (independent of trust model) perform best at reducing task time.$^{\Delta}$	With *honest* agents, neither trust model nor group size matters at reducing organizational effort; with *deceptive* agents, larger groups of forgiving agents perform best at reducing organizational effort.$^{\Delta}$	With *honest* agents, neither trust model nor group size influences the percent of items found; with *deceptive* agents, larger groups of forgiving agents performing best.
Gossip	In *stable* and *turbulent* environments, larger groups (independent of trust model) perform best at reducing task time. $^{\Delta}$ With *honest* and *deceptive* agents, larger groups (independent of trust model) perform best at reducing task time.	In *stable* and *turbulent* environments, larger groups (independent of trust model) perform best at reducing organizational effort. (Trust models only matter with large groups in turbulent environments, with larger groups forgiving agents performing best.) $^{\Delta}$ With *honest* agents, neither trust model nor group size matters at reducing organizational effort; with *deceptive* agents, forgiving agents perform best.$^{\Delta}$	In *stable* environments, neither trust model nor group size influences the percent of items found; in *turbulent* environments, larger groups of forgiving agents perform best. With *honest* agents, neither trust model nor group size influences the percent of items found; with *deceptive* agents, larger groups of forgiving agents perform best.

$^{\Delta}$ indicates a change in the findings from the prior results.

Table 4
Summary of Advice x Trust x Deception

Number of Agents	Task Time	Organizational Effort	Items Found
Agents (A)	20 > 50	20 > 50	20 < 50
Trust Model (M)	*ns*$^{\Delta}$	F < U	F > U
Deception (D)	H < D	H < D	H > D
A x M	*ns*	M.inc→ A$^{\Delta}$	M.dec→ A
A x D	D.dec→ A	D.dec→ A	D.dec→ A
M x D	M.inc→ D	M. inc→ D	M.inc→ D
A x M x D	*ns*$^{\Delta}$	*ns* → sig$^{\Delta}$	*ns* → sig
Best Honest:	A_L	$A_L$$^{\Delta}$	*ns*
Best Deceptive:	A_L	A_LM_F	A_LM_F

$^{\Delta}$ Result different from (Table 2). All reported effects significant at $p < .01$.
F = forgiving trust model, U = unforgiving trust model
H = honest agents, D = deceptive agents
A_L = large agent group; M_F = forgiving trust model

Table 5
Summary of gossip results

	Task Time	Organizational Effort	Percent Items Found
Environmental Stability			
Best Stable (No Gossip)	223.19 A_L	1,093.74 $A_L M_F$	97.6% *ns*
Best Turbulent (No Gossip)	223.39 A_L	1,156.98 $A_L M_F$	97.6% $A_L M_F$
Best Stable (Gossip)	213.32 A_L	1,075.34 $A_L M_F$	97.5% *ns*
Best Turbulent (Gossip)	256.02 A_L	1,701.97 $A_L M_F$	93.7% $A_L M_F$
Deception			
Best Honest (No Gossip)	223.21 A_L	1,093.73 A_L	97.6% *ns*
Best Deceptive (No Gossip)	286.94 A_L	4,290.20 $A_L M_F$	84.8% $A_L M_F$
Best Honest (Gossip)	213.43 A_L	1,075.33 A_L	97.4% *ns*
Best Deceptive (Gossip)	292.00 A_L	9,181.52 $A_L M_F$	60.7% $A_L M_F$
Pooled Conditions			
Best Turbulent & Deceptive (No Gossip)	476.93 *ns*	7,725.99***	68.6%***
Best Turbulent & Deceptive (Gossip)	483.88	9,210.80	61.1%

All reported effects are significant at $p < .01$ unless otherwise noted.
A_L = Group size, larger group dominates
$A_L M_F$ = Group Size x Trust Model interaction, larger group of forgiving agents dominate
*** = $p < .001$
ns = no significant effects detected

Table 6
<u>The summary impacts of gossip</u>

Under conditions of advice uncertainty…

Gossip degrades the performance of the group in terms of organizational effectives (Percent Items Found) and efficiency (Task Time, Organizational Effort) measures.

This degradation occurs because…

Gossip reduces performance *differences* afforded by the different (forgiving, unforgiving) trust models

such that…

Groups of forgiving agents *reduce* their performance levels to those obtained by groups comprised of unforgiving agents.

This performance reductions occurs because…

Gossip causes the communication levels (reducing good and bad advice overall) of forgiving agents to decline to levels obtained by unforgiving agents,

but as there was good advice mixed in with the bad…

The comparative advantage that the forgiving model had for this task (i.e., the tolerance for bad advice in order to obtain some good advice led to better performance) was lost because of ignored advice.

Although gossip induced the behaviors of the two trust models to be similar (i.e., to ignore advice), the reasons for ignoring advice differed as…

Forgiving agents ignored advice because primarily because of gossip, while unforgiving agents ignore advice primarily because of trust judgments

so…

Gossip serves to insulate negative impacts on trust based on direct experience (for forgiving agent trust models).

Thus, in a sense, gossip serves to preserve the coalition (advise network) at the expense of organizational performance as…

Gossip causes more information to be withheld in the accomplishment of the task over combined uncertainty conditions.

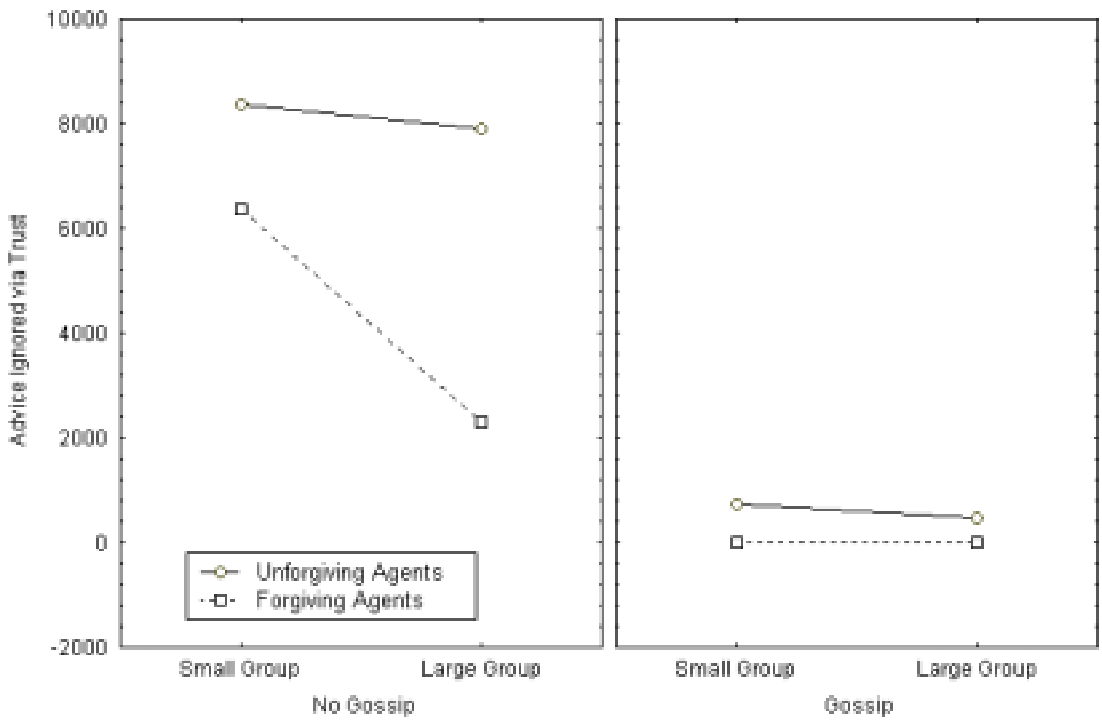

Figure 1
Advise ignored because of trust judgements

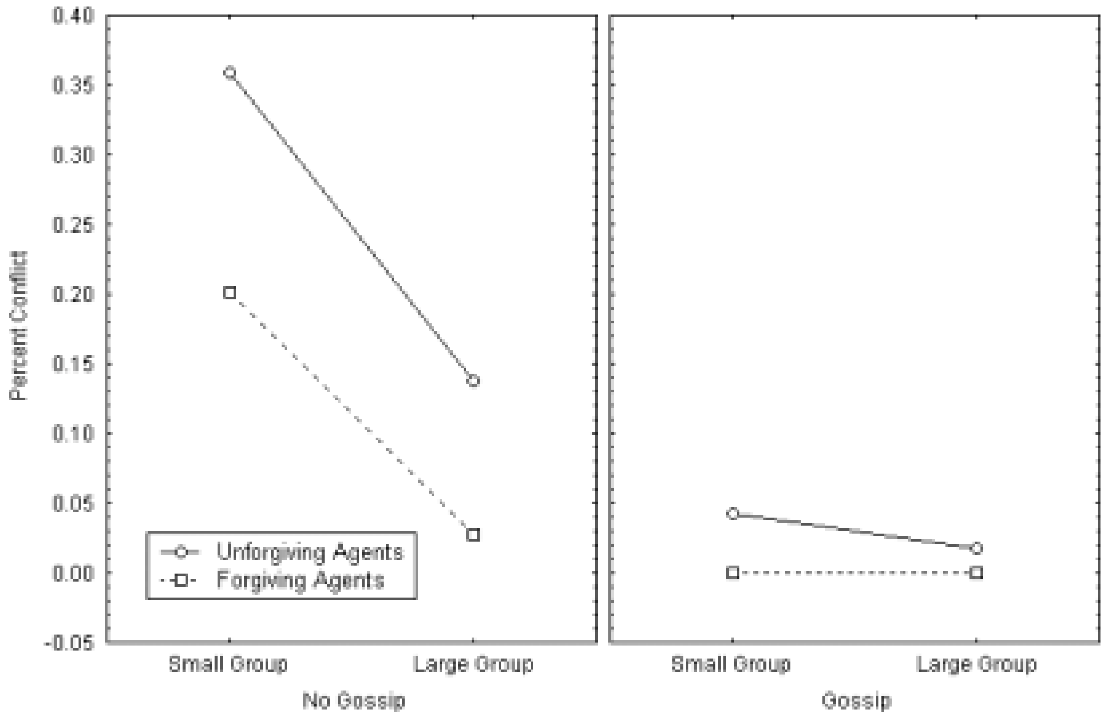

Figure 2
<u>Percent of conflict (distrust) in agent advice networks</u>