Chapter 1

BOUNDEDLY RATIONAL AND EMOTIONAL AGENTS COOPERATION, TRUST AND RUMOR

Michael J. Prietula, Kathleen M. Carley
University of Florida,
Carnegie Mellon

Abstract Computer-based agents, in various forms, are becoming actively involved in our personal and professional decisions and deliberations. We interact with them; they interact with each other. In this paper we describe a broad Model Social Agent study where we explore how boundedly-rational agents with emotion behave across increasingly social contexts, and the impact of cooperation, trust, rumor, and deception within those contexts. We report the results of an initial set of small simulations that suggest the use of rumor has a beneficial group effect by attenuating the extreme reactions of overly trustful and overly distrustful groups in increasingly uncertain decision environments. This is accomplished by influencing the destruction or the preservation of the coalition.

Keywords: : computational organisational modelling, agent trust, rumor, bounded rationality, emotional agents

1. INTRODUCTION

Our work focuses on understanding the relationship among humans, agents, tasks, and the social situations in which they are engaged. From this, we seek to establish the elemental basis of social behavior and group phenomena, and make predictions about them. Our guide for this effort is the Model Social Agent matrix.

The Model Social Agent matrix (Carley & Newell 1994) provides a two-dimensional categorization scheme that specifies the kind of knowledge required by the agent(s) (in terms of increasingly complex social situations), and the kind of information processing capabilities required by the agent(s) to operate with that knowledge, in order to exhibit various kinds of individual and collective phenomena. The scheme is summarized (in slightly modified form as will be discussed) in Figure 1.

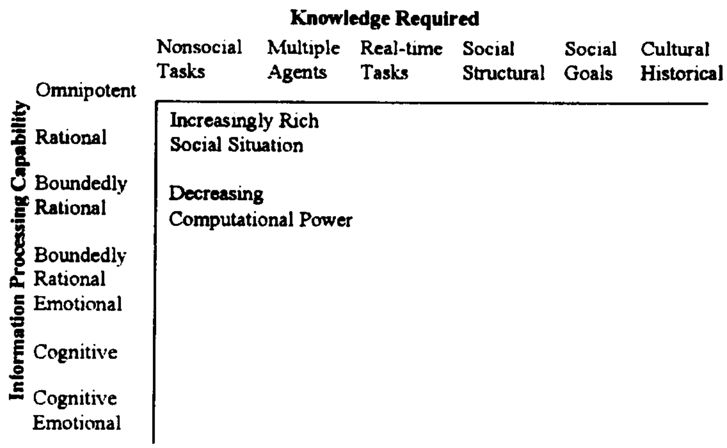

Figure 1.1 Modified Model Social Agent Matrix

Agent information processing capability range from an agent that is omnipotent as Laplace's demon (i.e., can process all knowledge relevant to the specific situation), to an agent that is both cognitively and emotionally defined (and thus constrained). The categories for this dimension are based on general differing critical assumptions of agent models and would certainly allow subtypes and variations. All agents have the capability to define and alter goals, and interact with other agents and objects. What differs is their capability to exploit the knowledge available within the processing (and sometimes emotional) constraints available. As agents' processing capacity diminishes, different, and perhaps more complex, behavior emerges.

The type of knowledge an agent has regarding social situations range from none (i.e., goals are defined solely in terms of self, any other agents are viewed as objects), to an agent that has a historical perspective and influence of culture, defining beliefs, norms and values. As the social situations an agent is presented with becomes increasingly complex, the richer (and more detailed) the knowledge has to be for the agent to successfully operate in the setting. Consequently, the behavioral repertoire increases as the complexity of the social situation increases. For example, a critical change occurs within the Social Structural setting. Prior situations require that an agent only needs self-defined (e.g., preservation, accomplishment-based) and task-defined (e.g., acquire item) goals. It is here when the agent is faced with goals that are externally (i.e., socially) defined.

As we have noted, this is a modified version where we have included Boundedly Rational Emotional Agents - the focus of this research. Note that Cognitive and Cognitive Emotional agents subsume bounded rationality. Overall, bounded rationality (in the strong sense) refers to a procedural rationality, that is, to the mechanisms of action and not the outcomes of action, at the individual or institutional level (Rowe 1989, Simon 1976b). These agents are restricted in both their processing capacity (what they can know) and how they use it. They rely on internal (perhaps flawed) representations of the social and task situations, but are goal-driven and bring to bear knowledge in service of those goals.

In this work we explore the impact of emotional components on (boundedly rational) agents in a social setting. Emotions serve as a nontrivial component of certain problem solving and decision making processes (Simon 1967). As one might well imagine, there are a host of definitions for "emotion" with recent attempts at incorporating various types of emotion in computational agents (e.g., Bates 1994, Morignot & Hayes-Roth 1996, Picard 1997), with simulations even approximating physiological components and influences (e.g., Canamero 1997).

However, we elect to incorporate a more symbolic form that can be situated in a broader context of problem solving and deliberation in specific settings. Specifically, we focus on the cognitive influences of agent affect, emotion, and behavior. Our view of emotion is most similar to that articulated by Ortony, Clore and Collins (1988) as "valenced reactions to events, agents, or objects, with their particular nature being determined by the way in which the eliciting situation is construed (p. 13)." In fact, their theory forms the basis for the elicitation structure of emotion in our model, so we refer to it as the ET component. Under certain conditions, Agents can have emotional responses to events. Basic affective reactions are differentiated with respect to cognitive constraints (as conditions) defining a fundamental set of emotional types.

The task we model explores the mutual effects of events, behaviors, and objects under varying agent properties, event types, group sizes, and task stability. It is a generic Drosophila task on which we can begin to tack our theory. The task defined within this model is simple:

Agents seek specific items in an external search space (e.g., the Internet, a warehouse) and may cooperate (give and receive advice) as to item locations or the quality of other Agents' advice.

What complicates matters is how task and agent properties impact choices, events, affect, and emotions. As this task is communication-based (i.e., all interaction between Agents are communications), then trust (in other Agents' advice) is a natural cognitive construct to explore within the setting. In our model, trust has both cognitive and emotion-based components. As this task is also about communicating collectives, then cooperation (giving advice about locations) and rumor (giving advice about other Agents' advice) are additional organizational constructs to explore within the setting. In our model, generating advice and rumor are behavioral responses that also have both cognitive and emotional components.

In this chapter we report on the results of an initial set of simulation experiments that explore how differential trust models, presence of deceptive agents, and rumor impact group performance under stable and unstable task coditions.

2. THEORETICAL COMPONENTS

The foundation of our efforts is a melding of two theories that address different, but complementary, components of emotion. The first theory is a treatment of how emotions can be represented and elicited under a cognitive structure. Ortony et al. (1988) present a detailed description of what could be described as a cognitive-based theory of emotional elicitation. We refer to their theory as an elicitation theory (ET) where they posit the overall macro cognitive structure (at an abstract level) linking eliciting conditions (event perceptions) to emotional types and the factors that can influence their intensity. In general, we often interact with events around us quite affectively (i.e., in general "positive" or "negative" type reactions). This theory makes the argument " ... whether or not affective reactions are experienced as emotions depends upon how intense they are" (p. 20). Thus, implicit in the eliciting condition is the notion of sufficient intensity in the context experienced.

Eliciting conditions describe the conditions (defined as valenced reactions to situations) under which an emotional type can be triggered. In ET, there are three basic situations that could be perceived (sometimes different from reality) as cueing such reactions: events, agents responsible for events, and objects (qua objects). Though emotional onsets require varying degrees of cognition, it is important to note that this does not always require deliberations that are conscious to the Agent. Rather, " ... they (emotions) are determined by the structure, content, and organization of knowledge representations and the processes that operate on them" (p. 4). In our work agents responsible for events (i.e., as advice) primarily cue the reactions.

In our model Agents experience events and, under certain conditions, emotional states are invoked (per ET). However, our model requires initial valued states (that permit analysis of threshold levels), adaptation of values with respect to experienced events (e.g., affective and cognitive impact of Agent messages), behavioral choices (e.g., to give or accept advice), and subsequent event experiences. Thus it is essential that our model address, for each event, the effects of the event on: (a) the underlying affect structure of the participating Agent, and (b) the behavioral choices of the Agent. This is handled by the second theoretical contribution: Affect Control Theory (Heise 1992, 1987, 1979).

Let events experienced by Agents (and descriptions of events via communication among Agents) be stipulated in ABO constituent form: Actor (Agent) - Behavior - Object (of the Behavior). For example, "The agent (A) deceived (B) the coworker (O)." In Affect Control Theory, each of the constituent elements of such an ABO event description has an associated 3-tuple set of EPA values (Evaluation, Potency, Activity) associated with it (from the perspective of the recipient of the message). Therefore, any given ABO event description has a 3x3 matrix of invoked values (whatever they might be). These are described as out-of-con EPA values and reflect the perceiving Agent's fundamental sentiments toward the elements in isolation, but perhaps defined within a general contextual frame (Carley 1986).

Event descriptions, however, describe agents and behaviors in-context and it is such events that alter affect, and consequently emotion and behavior. Affect Control Theory (Heise 1987) "shows how affective meanings of social identities and behaviors are maintained while they control interpersonal perception and social action" (p. 1). Specifically, it

- a) defines precisely the underlying scales representing the dimensional values, and does so in a manner that the dimensional scales are operationally meaningful - operators can incorporate multiple scales (quantification);
- b) specifies how the EPA values associated with the constituent elements (ABO) change in context (adaptation) through a set of normalized equations; and,
- c) describes how subsequent behavior choices are influenced (response) through a set of normalized equations.

Thus, given a presumed event (not yet occurred), the adaptation equations can be incorporated to predict the deflection of an ABO event/message. This is the important point: Affect Control Theory argues that choices of subsequent ABO events (to the extent they are available) are those that minimize the deflection from the fundamental sentiments. In effect, this is an effort to select events that confirm the fundamental sentiments.

The details of the model can be found in Prietula and Carley (1999) and is depicted in Figure 2.

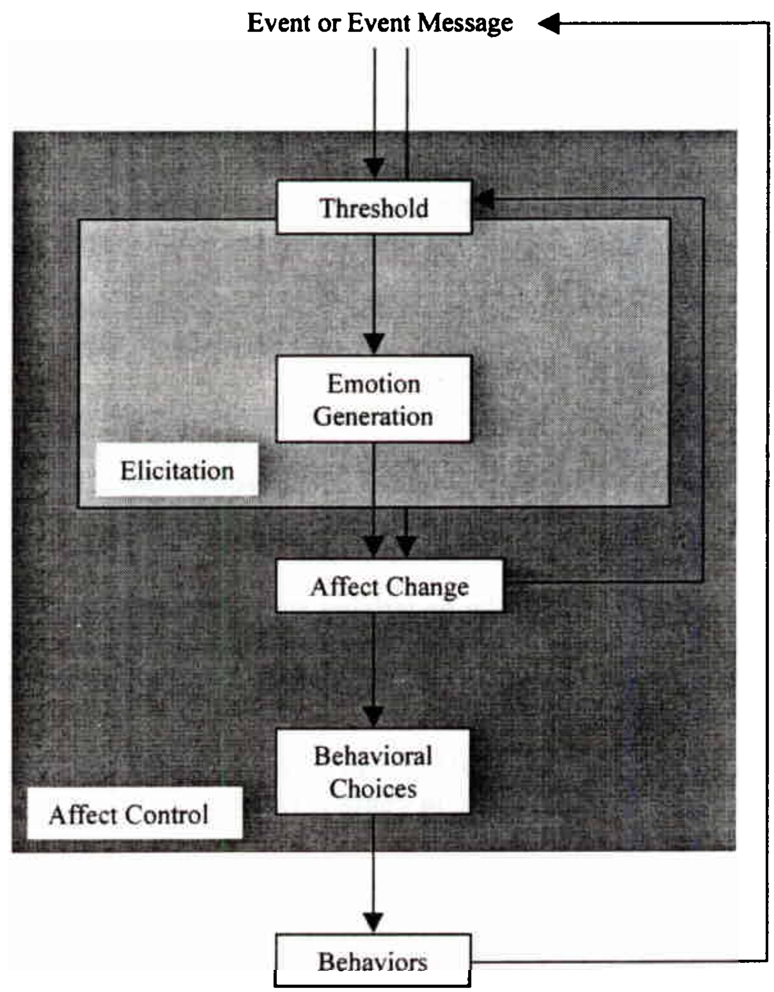

Figure 1.2 The EBR Agent Model

3. COOPERATION, RUMOR, AND TRUST

As we have noted, our models are characterizations of a multi-agent search task. Agent cooperation in these models is defined solely in terms of communication. The impact of emotion influences the choice of cooperation (i.e., communication) behaviors. As such, we explore the constructs of cooperation, trust, and rumor.

3.1 COOPERATION

In our model, agents can seek (and provide) cooperation in the form of providing advice of two types. First, agents can provide information to facilitate the search goal of another agent by communicating some of its experiential knowledge – whether the agent has seen the item (this requires a task memory for the agent). Second, an agent can provide information that characterizes the agent's experiences with another agent – whether the agent in question has provided helpful location advice (this requires a social memory for the agent). In general, advice is always in response to a question from another agent. An agent chooses to request advice as a search strategy, but an agent chooses to provide advice based on emotionally influenced issues regarding agent behaviors.

Because of the critical role of cooperation to an agent and an agent group, advice becomes an important individual and social construct that serve individual as well as social goals. Consequently, we include a deception construct for advice. Thus, agents can not only cooperate or not (i.e., give advice or not), but can also "anti-cooperate" (i.e., interfere) for either individual or social reasons. Accordingly, trust in advice must be considered.

3.2 TRUST

If cooperation in the form of advice is a critical component of behavior, then it is important and functional for agents to consider whether or not the source of advice is believable. In our model, this simply means that if Agent i trusts Agent j, then Agent i will believe any advice from Agent j.

Trust has a host of definitions. In our work we generically define (interpersonal) trust as the Agent's ability, given attributional information, to act on predictions and make predictions that other Agents will act in a cooperative manner. Trust is primarily a cognitive construct (based on deliberation) and secondarily an emotional one (Prietula & Carley, 1999). The cognitive components of trust is based on default assumptions (of other agents), and direct experience (with agent communications). For example, we might assume that, at least initially, all agents trust each other. If an agent does not provide correct advice, then that agent might not be trusted or might be suspect (depending on the trust-algorithm for the agent).

When humans interact with agents, the picture becomes more complex. In particular, when advice is characterized from a computer agent, humans react differently in their construction of trust judgements (Lerch, Prietula & Kulik, 1997). This appears to be based on a quickly emergent series of attributional effects centered on the attribution of knowledge to the system, and can even impact a construct of "machine faith" (Lerch, Prietula & Kim, 1999). Specifically, in Lerch, Prietula and Kulik (1997), a series of human experiments revealed how trust in a computer-based agent could be defined and influenced. In Lerch, Prietula and Kim (1999) a follow-up series of experiments are described where they develop an instrument to assess and predict trust responses in a computer-based agent.

Emotion impacts trust when affect responses exceed thresholds under certain circumstances. Extending the above example, if an agent does not provide correct advice, then an emotional response could be activated if either the importance/significance of the event is large, or if the expectations of the agent are radically violated. On the other hand, an assumption of trust might be made as an emotional response to advice from a trusted source that performed a significantly approvable action (e.g., good advice). Cognition and emotion work together in effecting the three components of trust (Remple, Holmes & Zanna 1985): predictability, dependability, and faith. Predictability refers to the most concrete and observable dimension of trust – predicable agent behavior over time in a presumably stable environment. Therefore, informing an agent on the stability or instability of the environment could impact these attributions.

Emotion and trust are highly contextual and may have both individual and social consequences. The contextuality is (properly) suggested by the definition provided by Ortony et al. (1988). Thus, the exact nature of how emotion and trust unfold depends on the task, agents, and social situation.

3.3 RUMOR

We include rumor in this model as both a behavioral choice and as an organization construct as it reflects a very real social mechanism of communication. Current communication technologies alter the speed and scope of their distribution and (presumably) their organizational effect. Once whispered or dispensed on the telephone, rumors now are broadcast real-time on firm's email or the nation-world Internet in seconds.

In their classic study of wartime rumors, 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).

Our interpretation of rumor, in fact, overlaps with what is called "gossip." Rosnow and Fine (1976) make a distinction between rumor and gossip in that gossip involves "less important," more personal issues, that can be either true or false, with a shorter communicative life-cycle, that serves more ego-based and social-exchange functions. Rumor, on the other hand, involves some sort of closure-seeking, is not verified, often concerns more "significant matters," and even serve to function as a mechanism for group problem solving (Shibutani 1966). We see both types occurring in a work setting. A rumor can exist regarding possible lay-offs (facts are unknown, but the matter is significant) while gossip can surround the discussion (e.g., Who is spreading the rumor? How reliable are they?).

In our model, rumors are communications of two types: information about a location (similar to the definitions of rumor above), but also information about other agents (similar to the definition of gossip above). Our interpretation of rumor subsumes that of task-related "institutional gossip" as we see this as an important component of defining and maintaining informational coalitions.

4. A SIMULATION EXPERIMENT

We describe a simple simulation experiment where we examine the behavior of emotional and boundedly rational agents under varying task conditions in specific social settings.

4.1 THE TASK

The task involved a set of agents, where each agent sequentially acquires an order for an item at a particular location (the order stack), then searchers for that item in a series of other locations (in the warehouse), retrieves that item when found, places the item at another specific location (conveyor), and then proceeds back to the order stack for the next request. The general form of this is shown in Figure 3.

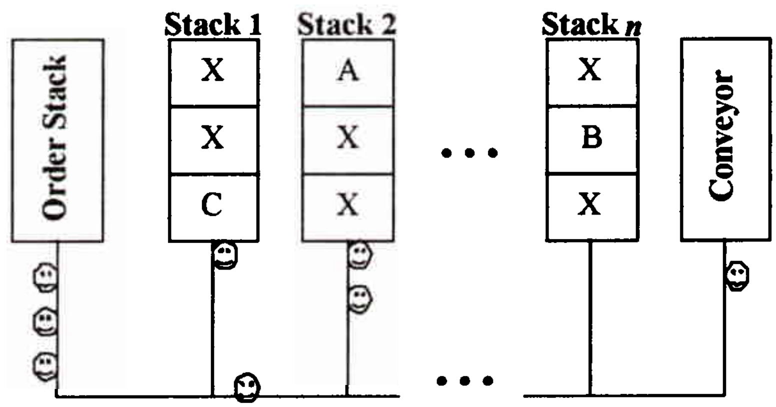

Figure 1.3 Structure of the Warehouse Task

An agent can determine if an item is in a location (stack) only by going to that particular stack's queue. If the queue is empty (i.e., no other agents are there), then the agent can access the stack and retrieve the item. If the queue is not empty, the agent must wait until it is before the stack can be accessed.

The warehouse task was implemented with ten locations and thirty unique items that were systematically assigned in decreasing modulo fashion starting at Stack 1 (with Item30, Item20, and Item10) through Stack 19 (with Item21, Item11, and Item1). This assignment was constant for all simulations. One critical task property was manipulated, environmental stability, as follows. If a task environment was stable, agents could access items in stacks without disrupting other items in the stack. If a task environment was turbulent, agents accessing a stack with multiple items would cause disruption by random placement of non-retrieved items to other stacks.

4.2 SOCIAL SITUATION

The social situation was a Real-Time Interactive setting where agents could communicate with other agents (see Figure 1). No particular social structure is defined, so judgements are based on direct or communicated experience. Rumor structures are supported as a communication structure, but the lack of social structure imposes no specific group constraints or effect. Task and individual goals exist, but no social goals exist.

4.3 AGENTS

Five agents were used in all simulations, and all simulations began with the agents queued at the OrderStack. For all simulations an agent could take only one order at a time. Three (indeed interrelated) manipulations were made to the agent models: Agent cooperation, emotionally-influenced trust models, and propensity to assert and listen to rumor.

Cooperation. In this simulation two types cooperative agents were defined: honest or deceptive. Agents remain that way throughout the simulation. We also vary the number of deceptive agents in the group (for the other manipulations), ranging from zero (no deceptive agents) to five (all agents are deceptive).

Trust. The trust models used in the simulation are the following:

1. *Model 1 – Trusting.* These agents do not see events as sufficiently important to generate an emotional response to alter trust behavior. The always view other agents as honest and trust them. For these agents, emotional states (and thus responses) are not relevant.

2. *Model 2 – Forgiving.* These agents' exhibit emotional responses that are based on bad advice, or rumor, about other agents. For these agents, it takes two pieces of bad advice in a row from an agent before that agent is judged as deceptive. On the other hand, two pieces of good advice in a row can redeem the agent (and be once again viewed a trustworthy).

3. *Model 3 – Reactive.* These agents are similar to the Reactive agents in that only one piece of bad advice causes them to make judgements of deception. Once judged, however, these agents do not forgive. Emotional reactions to events are quickly translated into enduring dispositional judgements of agents.

4. *Model 4 - Distrusting.* These agents do not trust any other agent.

Rumor. As noted, rumors can play an important part in a group or organizational setting as surrogate experiences and attributional influences. In this simulation, rumors are incorporated as follows:

- If some Agent i received bad advice from some Agent j, and Agent i is a gossip, then Agent i starts a rumor that Agent j is a liar.
- Any Agent k (i.e., in all above models) that detects a rumor about another Agent j, will choose not to provide advice to Agent j.

Rumors have an immediate effect. Two levels of this variable are incorporated as a property of the agents: assert rumors, do not assert rumors. Perceived rumors do not change an Agent's judgement (about the rumored Agent), but only whether to provide advice or not, or to accept advice, or not. However, the Agent originating the rumor has experiences that can alter trust judgements (based on the particular Trust model).

4.4 DESIGN SUMMARY

Manipulations were made at the group level (i.e., homogeneous Agents within manipulation). The manipulations were:

- Task. Two levels of Environmental Stability were examined (stable environment, turbulent environment).
- Agents. Four Trust models (Trusting, Forgiving, Reactive, and Distrusting), Two levels of Rumor (asserts, does not assert), and six levels of Deceptive Agents (0. 1, 2, 3, 4, and 5) in any 5-agent group.

We report on the following sample of dependent variables:

- Organizational effort (total number of effort by the group);
- Information Withheld (an agent knows the answer to a question from another agent, but does not answer because there is a rumor about the questioning agent, or the agent is judged to be deceptive);
- Structural Duration (average duration of positive trust relations expressed as percentage of maximum cycles);
- Conflict (relations that have been judged as deceptive as a percentage of the total coalition).

5. RESULTS AND DISCUSSION

If we make the (admittedly broad) assumption that the set of models and task/agent properties we have selected are somewhat representative of a set of random variables taken from a population of models and properties, we can take a high-level view of main effects and interactions across that set. Table 1 presents the main effects obtained.

Table 1. Main Effects

	Effort	Info Withheld	Duration	Conflict
Trust	↑*	↓***	↑***	↓***
Turbulence	↑***	ns	ns	ns
Deception	↑***	↑**	ns	ns
Rumor	ns	ns	↑***	↓***

*p < .05, **p < .01, ***p < .001, ns = not significant

Increasing the trust models of the group results in subsequent increases in overall effort and more durable information coalitions, and subsequent decreases in the information withheld and conflict. Increasing the turbulence of the environment has a single effect of increasing effort. As more deceptive agents are added to a group, both effect and information withheld increases. Finally, if rumors are incorporated, the duration of the information coalitions and decreases the conflict of the group.

Overall, most of the effects are what one may expect, but bearing in mind the nested effects, we can explore further the interactions. However, limits in the values (variance) in some conditions restrict the explorable set of interactions. Consider, for example, the "extreme" behaviours of the Trusting and Distrusting agents - the former is always trustful while the latter is never trustful - so Duration and Conflict values cannot be used. In this analysis, we disregard these extreme models (Trusting, Distrusting) and explore the interesting interactions with Rumor for the Forgiving and Reactive models only.

Figure 4 and 5 show the interaction graphs for the Forgiving and Reactive models respectively, exploring how Deception, Environmental Stability, and Rumor impact organisational Effort.

Consider the left two panels (with no Deceptive Agents and one Deceptive Agent) of Figure 4. Turbulent (uncertain) environments increase the organizational effort. Rumors have little effect in turbulent environ-

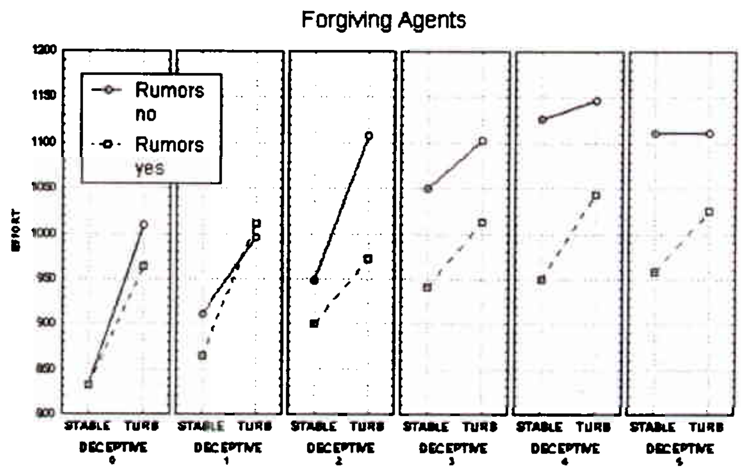

Figure 1.4 Interactions for Forgiving Agents on Effort

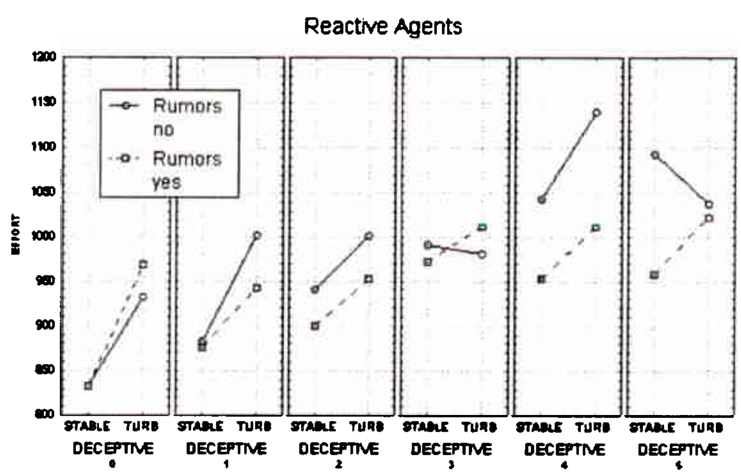

Figure 1.5 Interactions for Reactive Agents on Effort

ments, but begin to have some effect in stable ones (with one Deceptive Agent). As Deceptive Agents are added to the mix, benefits from rumors are reaped in stable environments and, to a lesser extent, in turbulent ones. Turning to Reactive Agents (Figure 5), the benefits from rumor are accrued but the benefits are more discontinuous. In addition, the general level of effort is attenuated as the number of Deceptive Agents increase. When the number of Deceptive Agents is three or five, there seems to be a slight Environment by Rumor interaction. Thus, Forgiving Agents benefit more from rumor in Deceptive and Turbulent environments than do Reactive Agents, as Reactive Agents not only actively spread rumor, but are more immediate in their behavioral changes with respect to advice as the likelihood of correct advice decreases (via Turbulence or Deception).

Figures 6 and 7 show the interaction graphs for the Forgiving and Reactive models respectively, exploring how Deception, Environmental Stability, and Rumor impact Information Withheld.

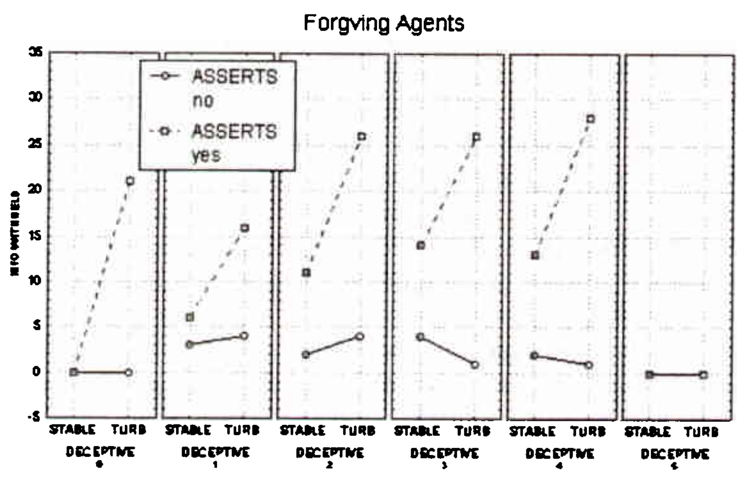

Figure 1.6 Forgiving Agents and Information Withheld

Note that information withheld means that a given Agent i perceives a request for advice from another Agent k and Agent i knows (or presumably knows) information of value to the questioning Agent k. If there is a Rumor about Agent k, the advice is withheld by Agent i. If Agent i has previously made a judgement about Agent k as being untrustworthy, the advice is again withheld by Agent i. On the other hand, recall that Deceptive Agents do not withhold information (though they are susceptible to bad advice, make judgements of other agents, and gener-

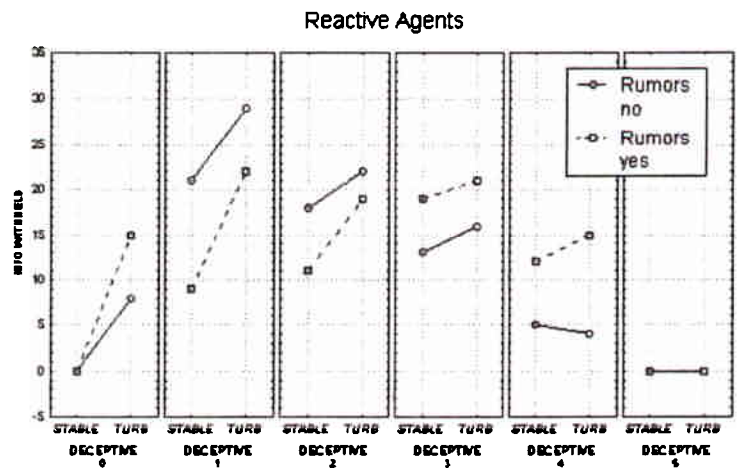

Figure 1.7 Reactive Agents and Information Withheld

ate rumors). Thus, Information Withheld is a direct function of rumor about a given Agent for non-Deceptive Agents.

First, we can examine the behavior without Rumors. In Figure 6, the stability of the environment has little effect on Information Withheld. Without rumors, Forgiving Agents must rely on direct experience, and that experience includes higher tolerance for bad advice, thus a general reduction in Information Withheld. Additionally, as the number of Deceptive Agents grows, the number of Agents withholding Information decreases (recall Deceptive Agents do not withhold Information). The graph for the Reactive Agents (Figure 7) is somewhat more involved, but the same phenomena hold. Like the Forgiving Agents, Reactive Agents respond to advice and when that advice is bad, they alter their trust judgements. If such Agents are judged as untrustworthy, then they will not receive responses to their requests for advice (except by Deceptive Agents). Unlike Forgiving Agents, Reactive Agents are less forgiving and judge more Agents as untrustworthy; therefore, they withhold information when asked. Without rumors, these Agents rely on experience and their experience generates judgements and withholding of information. However, as the number of Deceptive Agents grows, the number of Agents withholding Information decreases, causing the overall decreasing level.

In order to examine the cases with Rumor, it is necessary to recall that information is withheld (by non-Deceptive Agents) because (1) the questioning Agent has been judged untrustworthy (via direct individual experience), or (2) there is a rumor that the questioning Agent is untrustworthy, where condition (1) dominates condition (2). Furthermore, as rumors inhibit direct experience, rumors inhibit judgement changes. For Forgiving Agents, their judgements are generally positive (Figure 6), so condition (1) above is generally not met. However, when Rumor is introduced, condition (2) above is met and persists as the number of Deceptive Agents increases. In fact, as the number of Deceptive Agents increases, the decreasing number of non-Deceptive Agents accounts for Information Withheld. For Reactive Agents, low numbers of Deceptive Agents result in judgement changes that dominate the effects of Rumor. As direct experience causes more negative judgements, less Agents are likely to listen to advice, experience bad events, and spread rumor. Recall that Rumors are derivative of following bad advice, so when less Agents follow less advice, less rumors are spread and, therefore, less rumors are taken.

Figures 8 and 9 show the interaction graphs for the Forgiving and Reactive models respectively, exploring how Deception, Environmental Stability, and Rumor impact the average Duration (as a fraction of maximum simulation cycles) of the information coalition members. For example, in Stable environments with no Deceptive Agents, the Information Coalition is maximally intact at 1.0.

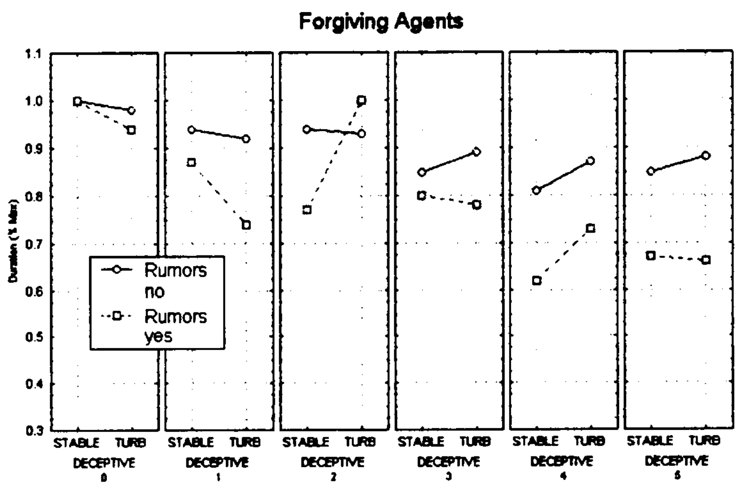

Figure 1.8 Forgiving Agents and Duration of Information Coalitions

In Figure 8, it can be seen that as the number of Deceptive Agents increases, the general duration of an information coalition of Forgiving Agents does not decline appreciably, and Turbulence does not seem to be a factor. Rumors have a general effect of reducing the Duration

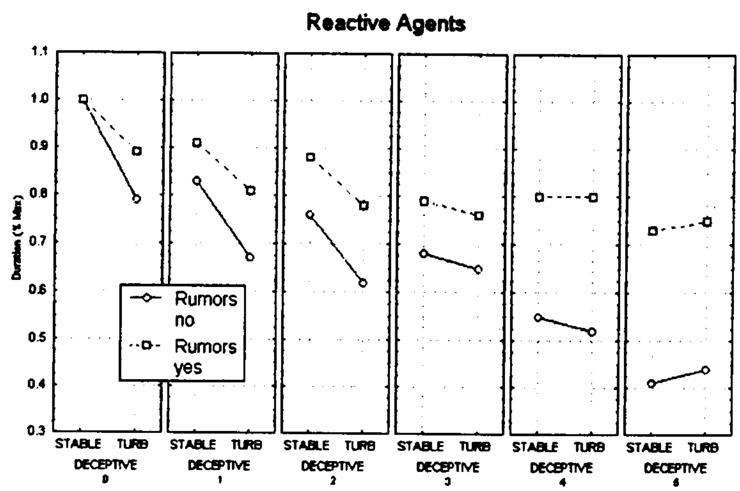

Figure 1.9 Reactive Agents and Duration of Information Coalitions

situationally. On the other hand, Reactive Agents (Figure 9) that do not use rumor have a general reduction in the Duration of the coalitions as the number of Deceptive Agents increases. For Reactive Agents, rumor sustains the Duration over increasing Deceptive Agents and turbulent conditions.

Thus, rumors have almost the opposite effects for the different Agent trust types. The reason for this is, again, the differential roles and effects of experience and rumor. Agent trust models are based on experience; that is, trust adjustments are based solely on advice from other Agents. Rumors, as noted, are weaker experience surrogates that inhibit the generation of advice (as cooperation), but attenuate the alteration of judgements. The Forgiving Agents of Figure 8, under Rumor conditions, do not follow advice (because of rumor) and thus cannot forgive the Agents. Rumor disrupts the coalitions. The Reactive Agents of Figure 9, on the other hand, are inhibited from imposing their negative assessments and rumor serves to preserve coalitions.

Finally, Figures 10 and 11 show the interaction graphs for the Forgiving and Reactive models respectively, exploring how Deception, Environmental Stability, and Rumor impact the Conflict (fraction of the information coalition disrupted).

Conflict is closely related to the Duration measure, so the explanations are equivalent. Left to their own calibration devices, Forgiving Agents (Figure 10) will augment their judgements through interaction with the Agents, but this is prevented when rumor structures are present. Simi-

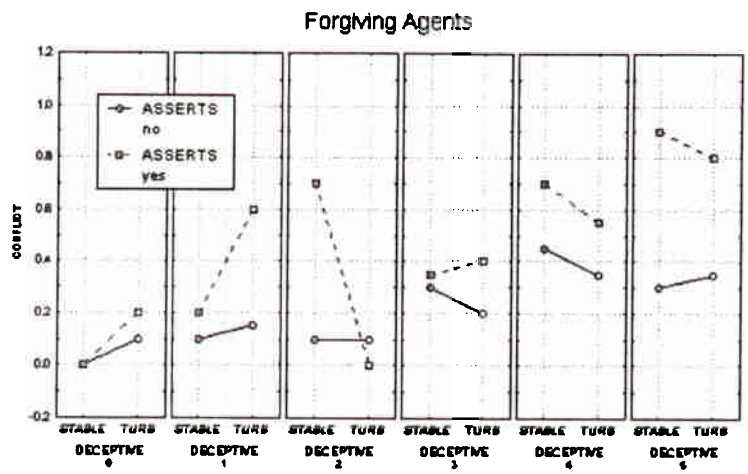

Figure 1.10 Forgiving Agents and Conflict

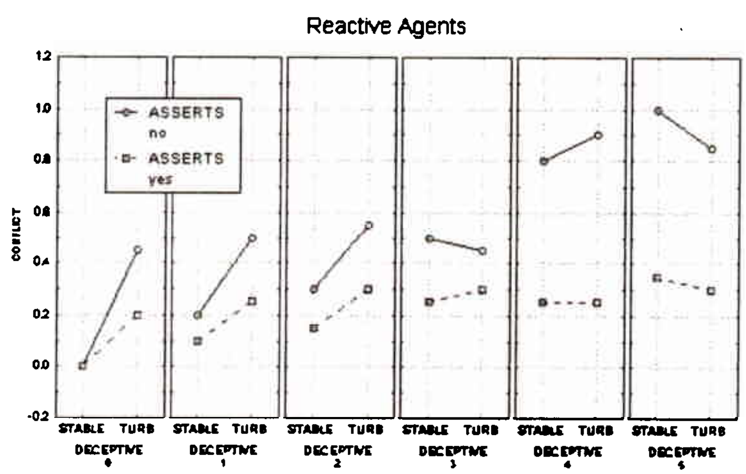

Figure 1.11 Reactive Agents and Conflict

larly, Reactive Agents (Figure 11) are quick to break up coalitions and increase Conflict. Again, rumor inhibits the direct experience and, by our definition, reduces conflict in the coalition.

How this works can be explained by examining the interaction between trust judgements adjustments and rumor. When Agents experience good or bad advice, they alter their trust judgements accordingly, though slightly differently as their trust models differ. Across Deceptive conditions, the general amount of good advice an Agent receives understandably decreases: more Deceptive Agents are generating bad advice and fewer Agents are generating good advice. Rumor, however, does not impact good advice as much as bad, as rumors are spread based on bad advice, not good (reflecting the general nature of rumor). Therefore, the effect is that rumor attenuates bad judgements, which necessarily depend on experience, and this impacts the Reactive Agents' ability to alter trust judgements down within Rumor conditions (Figure 12). This effect is also found in Turbulent environments, though there is an overall effect of less judgements up.

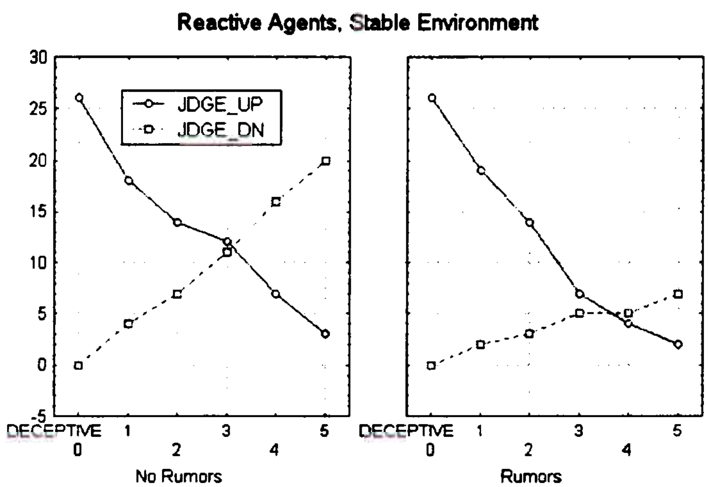

Figure 1.12 Reactive Agents and Judgements

6. CONCLUSION

The purpose of this paper was to craft an initial simulation in order to explore how different emotional-trust models of individual boundedly rational and emotional agents would impact group performance that depended, in part, on cooperation in the form of information. We defined four Trust models and explored how their decisions unfolded under a variety of conditions. The general nature of the Trust models were designed, in part, on empirical studies on how humans trust information from a series of different sources: peers, experts, and intelligent (computer-based) agents. We also described how two models of emotion could be woven together to generate emotion-based responses, though we did not implement this in this study.

Does trust matter? Overall findings suggest that different trust models have different organizational effects (recall Table 1) and we focused our results on two variations, Forgiving and Reactive Agents, and explored how groups would differ under varying conditions of Rumor, number of imposed Deceptive Agents, and whether the task environment was Stable or not. The most interesting results seem to occur with the Rumor construct and its interactions with the other manipulations.

Do rumors matter? Rumors indeed matter, but not the same way to all groups (and we are talking about groups of Agents).

- For Forgiving and Reactive Agents, as the number of Deceptive Agents increases in a group, rumor can reduce their general impact on increasing organizational effort;
- For Forgiving Agents, rumors increase the amount of information withheld (cooperation via advice) as the number of Deceptive Agents increases, but this effect does not carry over to Reactive Agents;
- For Forgiving Agents, rumors reduce the average size of an information coalition, but for Reactive Agents rumor increases the average size of the information coalition;
- For Forgiving Agents, rumors can increase the amount of group conflict, but for Reactive Agents rumor can reduce the group conflict.

How do rumors matter? Overall, rumors act by disseminating Agents' negative experiences (i.e., reactions to receiving bad advice) which, in turn, reduces the likelihood of a "deceptive" Agent deceiving other Agents (as rumors induce them to avoid advice from the offending Agent). On the other hand, by not taking the advice of an Agent, direct experiences cannot ensue that may alter the trust judgements of the Agent (up or down). Thus, one can think of adherence to rumors as a group trait, so to speak (i.e., do not supply information to an agent when there exists a rumor), while specific evaluations about an Agent's trustworthiness, based on direct experiences with that Agent, as an individual trait.

From the perspectives of the Agent types, Forgiving Agents have a "maladaptive under-reaction" to perform in extremely uncertain information environments (Turbulence and high Deception), so a group trait (rumor) facilitates the reduction of the information coalitions, and reduces the overall effort of the group (see Figure 6). These Forgiving Agents are prevented for forgiving other Agents (depending on general timing issues) based on their individual preferences for adjusting trust. Rumor facilitates the destruction of the coalition.

Reactive Agents possess individual traits to quickly terminate information coalition membership, but under certain conditions this might be an "maladaptive over-reaction" and rumor serves as a group behavior that mitigates that trait. Acting in the opposite manner, rumor inhibits Reactive Agents from direct experiences that would result in coalition termination. Rumor facilitates the preservation of the coalition. Overall, rumor has a dampening effect on the underlying trust judgements.

Does turbulence matter? Turbulence generally makes things worse, as expected, with minor exceptions. Furthermore, the effects of Turbulence generally decline as the number of Deceptive Agents increases. One interesting observation is the similarities between turbulence and deception as a source of uncertainty; that is, at what point does a Deceptive group become turbulent-equivalent? For example, in Figure 6 the amount of effect required for a group of Forgiving Agents with no Deceptive Agents (without Rumor) in a Turbulent environment exceeds that of a group with two Deceptive Agents in a Stable environment (without Rumor). If Rumor is added to the Stable case, the effort in the Turbulent case would exceed the effort exerted by an entire group of Deceptive Agents.

In conclusion, this work is embryonic and constrained. The next steps seek to incorporate the other specific components (e.g, Affect Control mechanisms), scale up (and down) the size of the group, alter the forms of agents, introduce stochasticity, and proceed with human-agent studies in order to determine the results that are more artifacts of the simulations than of reality.

7. ACKNOWLEDGEMENTS

We would like to thank Richard Burton (Duke), Benoit Morel (Carnegie Mellon), Andrew Ortony (Illinois), and Ray Levitt (Stanford) for their input and suggestions.

References

Albin, P. (1998). Barriers and Bounds to Rationality. Princeton, NJ: Princeton University Press.

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

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

Axelrod, R. (1985). The convergence and stability of culturs: Local convergence and global polarization. Working Paper 95-03-028, Sante Fe Institute, Sante Fe, NM.

Bates, J. (1994). The role of emotion in believable agents. Communication of the ACM, 37, 122-125.

Bennett, G. (1985). What's modern about the modern legend. Fabula, 26 (3/4), 219-229.

Canamero, D. (1997). Modeling motivations and emotions as a basis for intelligent behavior. In W. Johnson (ed.), Proceedings of the First International Conference on Autonomous Agents. New York, NY: ACM Press.

Carley, K. (1986). An approach for relating social structure to cognitive structure. Journal of Mathematical Sociology, 12, 108-116.

Carley, K. & A. Newell (1994). The Nature of the Social Agent, Journal of Mathematical Sociology, 19(4), 221-262.

CLIPS. (1993). CLIPS Reference Manual, Version 6.0. Lyundon B. Johnson Space Center, Software Technology Branch, Houston TX.

Cyert, R. and J. March (1963). A Behavioral Theory of the Firm. Englewood Cliffs, NJ: Prentice Hall.

DeAngelis, D ¿ & L. Gross (eds.) (1992). Individual-based Models and Approaches in Ecology. New York, NY: Chapman & Hall.

Degh, L. & A. Vazsonyi. (1976). Legend and belief. In D. Ben-Amos (Ed.), Folklore Genres. Austin, TX: University of Texas Press.

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

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

Epstein, J. & R. Axtell. (1996). Growing Artificial Societies: Social Science from the Bottom Up. Cambridge, MA: MIT Press.

Gaylord, R. & L. D'Andria. (1998). Simulating Society. New York, NY: Springer-Verlag/Telos.

Glance, N. & B. Huberman. (1994). The dynamics of social dilemmas. Scientific American, March, 76-81.

Heider, F. (1946). Attitudes and cognitive organizations. Journal of Psychology, 21, 107-112.

Heise, D. (1992). Heise, D. R. 1992. Affect control theory and impression formation. Vol. I, pp. 12-17. E. Borgatta and M. Borgatta (Eds.), Encyclopedia of Sociology. New York: Macmillan.

Heise, D. (1987). Affect control theory: Concepts and model. Journal of Mathematical Sociology, 13(1-2), 1-33.

Heise, D. (1979). Understanding Events: Affect and the Construction of Social Action. Cambridge, England: Cambridge University Press.

Heise, D. & N. MacKinnon. (1987). Event likelihoods. Journal of Mathematical Sociology, 13(1-2), 133-151.

Heise, D. & L. Smith-Lovin. (1981). Impressions of goodness, powerfulness, and liveliness from discerned social events. Social Psychology Quarterly, 44, 93-106.

Huberman, B. & T. Hogg. (1995). Communities of practice: Performance and evolution. Journal of Computational and Mathematical Organization Theory, 1(1), 73-92.

Hyder, E., M. Prietula & L. Weingart. (to appear). Getting to best: Efficiency and optimality in negotiations. Cognitive Science.

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

Kauffman, S. (1993). The Origins of Order. New York, NY: Oxford University Press.

Langton, C. (ed.) (1989). Artificial Life. Redwood City, CA: Addison-Wesley.

Lave, L. & J. March. (1993). An Introduction to Models in the Social Sciences. New York, NY: University Press of America.

Lerch, J., M. Prietula & C. Kulik. (1997). The Turing Effect: The nature of trust in expert systems advice. In P. Feltovich, K. Ford & R. Hoffman (Eds.), Expertise in Context. Menlo Park, CA: AAAI Press/MIT Press.

Lerch, J., M. Prietula & J. Kim. (1998). Unraveling the Turing Effect: Measuring trust in computer advice. Unpublished manuscript.

March, J. (1978). Bounded rationality, ambiguity, and the engineering of choice. The Bell Journal of Economics, 9(2), 587-608.

Morignot, P. & B. Hayes-Roth. (1996). Motivated Agents. Knowledge Systems Laboratory, Report KSL 96-22, Department of Computer Science, Stanford University, Stanford, CA.

Newell, A. (1990). Unified Theories of Cognition. Cambridge, MA: Harvard.

Ortony, A., G. Clore & A. Collins (1988). The Cognitive Structure of Emotions. Cambridge, England: Cambridge University Press.

Picard, R. (1997). Affective Computing. Cambridge, MA: MIT Press.

Prietula, M. & K. Carley (1999). Exploring the effects of agent trust and benevolence in a simulated organizational task. Applied Artificial Intelligence, 13, 321-338.

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

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

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

Rowe, N. (1989). Rules and Institutions. New York, NY: Philip Allan.

Shibutani, T. (1966). Improvised News: A Sociological Study of Rumor. Indianapolis, IN: Bobs Merrill.

Simon, H. (1993). Altruism and economics. Papers and Proceedings of the 105th Annual Meeting of the American Economic Association, 83(2), 156-161.

Simon, H. (1991). Organizations and markets. Journal of Economic Perspectives, 5, 25-44.

Simon, H. (1990). A mechanism for social selection and successful altruism. Science, 250, 1665-1668.

Simon, H. (1976a). Administrative Behavior (3rd Edition). New York, NY: Free Press.

Simon, H. (1976b). From substantive to procedural rationality. In S. Latsis (ed.) Method and Appraisal in Economics, Cambridge, England: Cambridge University Press.

Simon, H. (1967). Motivational and emotional controls of cognition. Reprinted in H. Simon (ed.), Models of Thought. New Haven, CT: Yale University Press, 1979.

Simon, H. (1956). A behavioral model of rational choice. Quarterly Journal of Economics, 69, 99-118.

Simon, H. (1955). Rational choice and the structure of the environment. Psychological Review, 63, 129-138.

Smith-Lovin, L. (1987). Impression from events. Journal of Mathematical Sociology. 13(1-2), 35-70.

Sober, E. & D. Wilson. (1998). Unto Others: The Evolution and Psychology of Unselfish Behavior. Cambridge, MA: Harvard University Press.

Williamson, O. (1975). Markets and Hierarchies. New York, NY: Free Press.

Wilson, D. & E. Sober. (1994). Reintroducing group selection to the human behavioral sciences. Behavioral and Brain Sciences, 17, 585-654.

Wittek, R. & R. Wielers (1998). Gossip in organizations. Computational and Mathematical Organization Theory, 4(2), 189-204.