Agents, Trust, and Organizational Behavior

Michael Prietula

Department of Commerce & Technology

Johns Hopkins University

201 N. Charles Street, Suite 200

Baltimore MD 21201

prietula@jhu.edu

Kathleen M. Carley

Department of Social and Decision Sciences

Carnegie Mellon University

Pittsburgh PA 15213

Kathleen.Carley@.cs.cmu.edu

An abbreviated version of this paper was published in the AAAI Technical report as "Agents, Trust, and Organizational Behavior", Michael Prietula and Kathleen M. Carley, pp. [[[[MIKE PLEASE FILL THIS IN]]], in Socially Intelligent Agents, AAAI Press Technical Report FS-97-02, 1997, American Association for Artificial Intelligence. Support for this paper was provided in part by the NSF under

Appearing in Dautenhahn, K.(Ed.) "Socially Intelligent Agents: Papers from the 1997 AAAI Fall Symposium," AAAI Press Technical Report FS 97-02, Menlo Park, CA: AAAI Press.

Abstract

Executives often argue that trust is critical if an organization is to perform well. Although articulated as a (functional) property of individuals, the collective effect of events influencing (and being influenced by) trust judgments can impact organizational behavior. In theory, without trust, work needs to be re-checked, decisions are continually in need of re-evaluation, personnel become less motivated, and performance takes a nose-dive. Trust, however, is a computational construct and is therefor capable of being explicitly modeled and studied. In this paper we summarize our work on agent trust and explore the joint impact of the trust, cooperativeness, and benevolence of agents on organizational behavior.

Agents, Trust, and Organizational Behavior

Introduction

To what extent does trust really influence organizational performance? Executives, and the literature on management, point to the importance of trust. In fact, trust is often argued to be a prerequisite for good organizational performance (Nicholas, 1993). In a recent survey on job performance respondents rated trust and getting along with others at work as the two most important skills (Frazee, 1996). But is trust really necessary for good organizational performance? To what extent is untrustworthiness even noticeable in an organizational environment?

Organizations are complex dynamic systems faced with environments which may be rapidly changing. For organizations, a large number of factors interact in complex and often non-linear fashions to determine performance (Carley and Svoboda, 1996). One important source of these non-linearities is the adaptiveness, not just of the organization as a unit; but also of the set of agents within the organization. In other words, change in what individual's know and in how the organization operates affects performance. However, change and ambiguity go hand in hand; i.e., what was so yesterday may not be so tomorrow. Lies can be thought of as a type of ambiguity. From an organizational perspective, then, the question becomes to what extent are lies even noticeable in a dynamic environment. Is lying disruptive to organizational performance if it is not possible to determine that the source is lying?

Most researchers on organizations have been concerned with one of two logics of performance, which are often at odds, the logic of the task and the logic of interaction. In the logic of task, the set of agents in the organization work collectively to solve some problem or achieve some goal. Such a logic assumes a set of cooperative and trustworthy agents. In such a system, ambiguity, should it occur comes from information errors external to the individual. In the logic of interaction, the set of agents communicate and exchange information in order to create and maintain social norms and to achieve individual goals. This logic, assumes neither cooperation nor trustworthiness but it does assume that task is not an overly influential basis for interaction. In this case, ambiguity, may arise either from the individual or from transmission errors between individuals. In this paper, we pit these two logics against each other when we ask, how does lying influence organizational performance.

Typically, studies of organizational behavior assume that the logic of task dominates and so task based collective behavior becomes interpreted and re-enforced by social norms of cooperation and expectations centered around trust. Empirical evidence does suggest, at the organizational level, that trust and performance go hand-in-hand. Further, at the inter-organizational level, Zaheer, McEvily, and Perrone (forthcoming) demonstrated that exchange performance was higher for those organizations where boundary spanners reported higher levels of trust in the partner organizations. At the individual level, however, the relation between trust, cooperativeness, and performance is less clear. Cooperativeness is not necessarily in the individual's best interest; lying can produce an advantage; and levels of trust can change (Klein, 1997).

Herein our concern is with trust in the organizational context. Nonetheless, our understanding of trust draws from the social and behavioral sciences more generally. Issues of affect, and trust in particular, have long been of interest to social scientists. Research in psychology and sociology has examined the impact of trust in individuals (Barber, 1983; Rotter, 1971; Shapiro, 1990), in close relationships (Johnson-George and Swap, 1982; Rempel, Holmes and Zanna, 1985; Rempel and Holmes, 1986), and in groups (Lewis andWeigert, 1985a, 1985b; Luhman, 1979; Zucker, 1986; Shapiro, 1987). Most of these studies have not been particularly interested in organizations as such. Recently, a few organizational researchers have begun to look at issues of trust. Much of this work is in the area of strategy and so often focuses on inter-organizational relations (Bradach and Eccles, 1989, Ring and Van de Ven, 1982; Zaheer and Venkatraman, 1995; Zaheer, McEvily and Perrone, forthcoming).

In the social and behavioral sciences much of the research is not on trust per se', but on related concepts such as the emergence of conventions (such as truth telling) and the evolution of cooperation. Often computational models are used to explore the impact of collective adaptation on the formation of group level behavior (Axelrod 1987; Skvoretz and Fararo 1995; Turner 1993). Computational models of cooperation abound (e.g., see Macy 1991b). Many of these models demonstrate that environmental and institutional factors such as payoffs, population dynamics, and population structure influence the evolution of cooperation in a discontinuous fashion (Axelrod and Dion 1988); that simple interactions can lead to the development of and change in beliefs and the underlying culture (Ballim 1991; Carley 1991a; Kaufer and Carley, 1993). Such studies also demonstrate that when collections of agents interact and each agent has a simple goal, such as trying to attain the highest cumulative award, interesting and non-trivial social dynamics and norms often arise (Shoham and Tennenholtz 1994). Macy (1991a) demonstrates the ability of social learning to result in cooperation in social groups. Kephart et al (1992) has argued that social chaos is reduced when intelligent adaptive agents determine their actions using strategies based on observations and beliefs about others. But would this be the case if agents lie? Following in this tradition we develop a computational model of organizational behavior in which the organization is modeled as a collection of intelligent adaptive agents. We examine whether or not truthfulness, as a pre-requisite for continued cooperation, is a necessary the reduction of social chaos and the consequent improved performance of the organization.

Importantly, from an agent perspective, the information gained during an interaction can also alter the likelihood of future interactions. In this paper, we take a constructural perspective (Carley, 1991a; Kaufer and Carley, 1993) on this interaction process. According to the constructural perspective, as individuals interact they gain information which changes the way they look at the world, their actions, and attitudes, and beliefs and this change in information, this knowledge level change, results in changes in who interacts with whom and the extent of those interactions. This dynamics driving this change is an evaluation of the relative similarity between agents. That is, two agents are more likely to interact if they both believe that they are more similar to each other than they are to others in the group. We extend this perspective by basing these similarity judgments on the perceived truthfulness of the other. In our model, agents always see themselves as truthful, whether or not they are. Thus, as they encounter others, if they come to believe that those others are less like themselves, i.e., they are liars, then they are less likely to engage that other in a future interaction.

Within organization science, researchers have examined the impact of uncertainty on organizational performance. Some of this work focuses on the relationship between individual uncertainty and overall organizational performance. Importantly for our purposes, some of this work led to the development of organizational models that begin to address issues of uncertainty in a formal fashion. For example, Carley, Prietula and Lin (in press) examine how the performance of the organization is affected by the agents being given incorrect information. Not surprisingly, they find that performance decreases under conditions of uncertainty. More to the point, they also demonstrate that there is an interaction between uncertainty, the organizational structure (team versus hierarchy), and the division of labor. Thus, some organizations, just because they have different structures may be less affected by uncertainty (Carley 1991b, Carley and Lin 1995). Differences at the individual level can have profound effects on organizational performance; however, the strength of those effects depends on the structure of the organization (Lin and Carley, 1993).

Recently, researchers have begun to explore issues of trust using computational agents. In particular, interest in the role of affective, social, and personality based agent attributes are of increasing interest to researchers in the areas of computational social and organizational theory (see for example, Castelfranci, de Rosis and Falcone, this volume; Canemero and Van de Velde, this volume) and in distributed artificial intelligence (Rosenschein and Zlotkin, 1994; Sandholm and Lesser, 1996).

One such study revolves around whether or not the crafting information and advice through a particular medium -- the computer -- has human-like interactional characteristics. In a communication, it is not simply the medium, it is the context in which the message is presented that affects the receiver. One point of impact of a system is the effect it can have in delivering advice to a human, who then subsequently makes decisions based, in part, on that advice. It is a form of exploring trust. This exploration can be carried out using a "good-bad" model in which the advice is either right or wrong or a more elaborate message model in which the "right or wrong" advice appears and is linked to a richer context of whom the agents are who are delivering this advice.

In Lerch, Prietula and Kulik (1997), a series of human-computer interaction experiments were conducted, comparing how trust levels varied with how the source of advice was characterized over the computer. That is, all advice was received as a form of email, but the source of the advice was differentially described (human experts, expert system, and peers). Results revealed a "Turing effect" in which the characterization of advice as coming from an expert system had significant effects on trust (e.g., they trusted the expert systems less than the human experts). In addition, the results provided insight into how these subjects "viewed" as an internally defined model (e.g., expert systems could not exert "effort," but humans could). Expert systems are seen, as characterized in these studies, as significantly different types of creatures. Furthermore, evidence reveals possible sources for the Turing effect and demonstrates that manipulation of how advice is characterized can influence trust judgments (Lerch, Prietula, Kim and Buzas 1997). Trust, in this experiment, was a function not of the message content, but of expectations about the nature of the agent delivering the information.

Herein, we describe a computational model of organizational performance under conditions of agent uncertainty. One source of agent uncertainty is inaccurate information. Since agents interact both with each other and the environment, this uncertainty may derive from agents not telling the truth or from agents facing a changing environment. We describe a series of virtual experiments that we ran, using this model, that address issues of trust in organizations. Results from these virtual experiments are discussed and the implications of these findings for organizations of humans and WebBots are discussed.

The Computational Approach to Trust

We can understand organizations, and develop our theories of them, by crafting and analyzing computational models of organizations as collections of intelligent adaptive agents. Computational theorizing is an invaluable asset to the organizational theorist as it helps to lay bare the implications of the agent, adaptation, task, and organizational characteristics on organizational performance. Computational theorizing helps us to reason through, systematically, the consequences of the multiple interacting factors present within organizations. Both social interaction and the capabilities of a social agent can be defined computationally (Carley and Newell 1994, Carley and Prietula 1994). As such, it is meaningful to conduct computational experiments which address social interaction as well as (associated) cognitive properties in order to explicate properties of organizations of such organizational and cognitive agents (Carley and Prietula in press, Carley, Kjaer-Hansen, Newell and Prietula 1992, Prietula and Carley 1994).

Computational theorizing is clearly useful when theorizing about adaptive or dynamic systems. In such systems, the level of complexity and the existence of important non-linearities dictate ultimate performance and make the system intractable from a mathematical point of view. In this situation, computational theorizing acts as an aid to the theorist by enabling the theorist to systematically explore the space of possibilities.

One issue that computational theorizing helps us to reason about in this paper is the impact of human characteristics (like trust) and human-like interaction (communication based exchanges) on organizational performance under different task and structural constraints. For example, we might ask, if the organization is facing a dynamic environment, where the lessons learned today are not necessarily applicable tomorrow, are there organizational consequence if agents lie?

To operationalize trust in this way implies that it is a computational entity. Our argument is that trust is primarily cognitive, and only secondarily affective. Since trust is primarily cognitive it is fundamentally computational. Further, we view trust as resulting from the long term cognitive processing of information about the actions and attitudes of others. Our view of trust as first and foremost a cognitive construct, and secondarily as an affective construct, is consistent with empirical findings on the nature of trust. For example, in a study of 194 managers and professional, McAllister (1995) found that levels of cognition-based trust were higher than levels of affect-based trust. This finding is consistent with the argument that trust develops as a cognitive construct and only then as an affective construct.

Secondly, as argued by Corazzini (1977) we view trust as a multi-dimensional construct. Trust encompasses a number of attributes including belief in the predictability of others and an expectation that the others will act with goodwill. Predictability of others implies that each agent thinks that others will behave as they did in the past. Hence, it is reasonable for each agent to predicate their actions on both their own knowledge and on the past actions of others. The expectation of goodwill means that each agent begins by assuming that others are honest. Truthfulness, if you will, is the default. Further, since goodwill is expected, even if agents think that others have lied they are going to discount the event, somewhat. Agents may learn that others lie and so respond to others as liars rather than truth tellers but it will take a number of instances of lying for that change to come about fully.

Task and Agent Models

In this paper we present results from two different models of team behavior. In both models there are multiple agents who interact and communicate as they perform a task. Some of this communication takes the form of requests for information and advice. Agents act on the advice they receive. Agents begin by assuming that others tell the truth. Agents, however, may or may not be truthful. Over time, each agent will develop a perception of whether or not they think each other agent is a liar. This perception, as it changes, will change the likelihood that the agent seeks out the other for information.

In both models the task is a generic search task which can be viewed isomorphically as search on the internet (e.g., Carley and Prietula in press), or searching for items in a warehouse (Prietula and Carley

1994). These two alternative versions are shown in Figure 1. Viewed from the latter perspective, each agent has task knowledge about how to get orders and how to tell when an order is filled. Each agent moves about in the warehouse locating items to fill the order. Agents can and do interact. One form of interaction is that agents can ask each other where the items they need are located. They are seeking simple advice. In response to this query, the other agents can either tell the truth (report where they last saw the needed item) or lie (purposely report that the needed item is somewhere other than where they last saw it).

*** Place Figure 1 About Here ***

Viewing this task as a warehouse, there are 10 item locations (where items are stored), an order queue (where requests are taken), and a delivery queue (where items are taken). A total of twenty unique items were in the order queue (non-redundant) and each agent takes one order at a time. The purpose of this preliminary study is simple. We wish to explore the extent that misinformation can impact organizational performance and behavior. Our model should generate the fundamental collective behaviors it purports to explain.

The major difference between the two models is the level of cognitive sophistication in the agents. In the first study we use very sophisticated cognitive agents -- Soar agents. In the second study, we abstract the main components relative to this task and use cognitively simple algorithmic agents. The second difference is that in the first study the items move about (because agents move them); whereas, in the second study the items remain where they are.

Computational Study 1

Demonstrating that differential trust judgments exist and can be manipulated on an individual level, we move to computational agent research. We use a multi-agent model of teamwork in which a collection of Soar (Laird, Newell and Rosenbloom 1987) agents seek to fill orders by collecting items from various stacks in a warehouse. As noted agents can communicate, request and provide information. In providing information some agents always tell the truth and others lie. Lying increases uncertainty as it decreases the likelihood that the information seeker will find the item it needs in the expected location. The objects in the warehouse, however, are not stationary. Specifically, as agents fill orders they may move the items they do not need to other locations so that they can access the items they do need. The fact that agents can move items makes the environment volatile and increases the level of uncertainty. For the searching agent, when an item is not where it is expected to be, the agent does not know for certain if the information it received was false (the responding agent lied) or if the information was correct but the environment changed (some third party moved the item). The number of agents in the organization were varied from 1 to 5 and all agents were either honest or liars.

The size of the organization and the trustworthiness of the agents interact. Thus, as the size of the organization increases organizations of Honest agents ask and answer more questions; whereas, organizations of Liar agents ask and answer fewer questions. This is simply the effect of the value of information. More accurate information increases the willingness to ask questions. More inaccurate information decreases the willingness to ask questions. However, the effects are not linear. There is also an impact of environmental volatility. Specifically, as the size of the organization increases the effects just mentioned continue to occur but at decreasing rates. As organizations increase in size the environment becomes more volatile. Thus, because items are moving, honest agents appear dishonest and dishonest agents appear honest. This results in, over time, organizations of Honest agents appearing less cooperative than organizations of Liar agents because the ration of answered to asked questions for Honest agents actually drops below that for Liar agents (see Figure 2). We would speculate that eventually there may be a point where volatility is sufficiently high that both curves "turn" and no difference is discernible due to trust. In the meantime, many different conditions may prevail and the uncertainty of the environment may counteract or augment the uncertainty due to lying.

*** Place Figure 2 About Here ***

Computational Study 2

Our second study weaves trust models from empirical human studies into simple computational models of groups. Specifically, we reimplemented the Plural-Soar model in a less-cognitive, agent-based form. As the prior research demonstrated the general nature of trust judgment deliberation, the particular deliberation model could be extracted from the data and crafted directly in algorithmic form.

There are three components to the agent's social behavior: honesty (will they lie or not), benevolence (how forgiving are they), and cooperation (under what circumstances do they give advice). In this study we hold benevolence and cooperation constant -- they are quick forgivers and they cooperate when they have the first opportunity. We vary the first social component: whether they lie or not about the location of an item.

Each agent has an item-memory (it recalls where it has encountered all items), a social memory (it recalls the value of the advice from other agents), and a communication capability (it can ask and receive advice). All three work to define the social interaction among agents. In this study, the nature of the interaction is as follows. Agents are either Liars or they are Honest. Essentially, this defines how an agent responds to a general request for advice. An Honest agent will respond directly to the questioning agent if it knows the location of the item in question. A Liar agent will respond to any request for advice, supplying incorrect location information. An agent, however, recalls the advice provided by an agent, and engages a simple social judgment model of advice acceptance, based on three judgment states: trustworthy, risky, and untrustworthy. Good/bad advice moves the judgments up/down. Advice from risky/untrustworthy agents is not accepted and questions from them are not answered (unless the answering agent is a Liar). If an agent is unforgiving, the untrustworthy state is absorbing an agent deemed untrustworthy remains so judged. As noted, all agents in this study are benevolent. It is a forgiving group.

The first questions we address is simply the general impact lying would have on group behavior. This serves as a baseline for subsequent studies. Figure 3 summarizes the Total Moves taken to complete the 20 item task for organizations of either all Liar agents or all Honest agents. The size of the organization ranges includes 1, 2, 4, 6, 8, or 10 agents. As can be seen in Figure 3, an organization of Liar agents quickly incur penalties for their incorrect advice and so collectively expend more effort. The primary reason for this resides in the nature of the advice being generated by each group. In

*** Place Figure 3 About Here ***

In Figure 4, the relationship between advice and the size of the organization is shown. Here we another interaction between organizational size and truth-telling. As organizational size increases BAD advice (occurring in the Liar group) and GOOD advice (occurring in the Honest group) increase. However, the bad gets worse and good gets attenuated.

*** Place Figure 4 About Here ***

In general, in stable environments, it is expected that organizations of Liar agents will do worse than organizations of Honest agents. What is perhaps less obvious is the impact Liar agents would have in a Benevolent group. This was explored as follows. The organization's size was held constant at 10, while the number of Liar agents in the group was varied according to the schedule: 0, 1, 3, 5, 7, 9, 10. Thus, the extremes considered are the prior reported conditions: all Liar agents, all Honest agents. It is the mix that is of interest. Figure 5 summarizes the Total Moves for the seven groups simulated. What this shows is that with only 30% Liars the group as a whole behaves as though the entire organization was composed of liars.

*** Place Figure 5 About Here ***

The underlying events can be traced to the good and bad advice generated under the different conditions. Figure 6 summarizes these results. Note that in Figure 6, the amount of Bad advice generated quickly dominates the organizational communication. Thus, at low ratio levels a small number of Liar agents can be extremely disruptive to organizational performance. How this occurs is two fold. First, in order to provide Good advice, an agent must have good information. That is, the agent must know where a given item is located. On the other hand, giving Bad advice requires no information, unless an agent wishes to ensure that an item is not in a known location. These agents, however, simply generate the same incorrect response. They do not search their memories for the correct one, nor do they require it. It is the preference for these agents to trust communication. Without a given advised location, they will engage in a systematic search. However, information provided will cause them to do directly to the suggested location (without investigating interim locations). This has the effect, for Liar agents, of sending agents to false locations, causing Location faults going to a location without success. Figure 7 summarizes this result.

*** Place Figures 6 and 7 About Here ***

Discussion

We have presented an exploratory analysis of the relationship between an affective response, trust, and organizational performance. Further research in this area needs to explore the impact of other affective responses. Individual agents develop cognitive coping mechanisms for dealing with uncertainty, and their affective response to this uncertainty. These individualized cognitive coping mechanisms may be, and are often assumed to be, detrimental to the organization as these cognitive mechanisms may lead to individuals acting in more rigid, less flexible, less efficient fashions. However, as even our very simple models illustrate, there are interactions between agent based sources of uncertainty (such as lies) and environmental uncertainty (movement of items) which can alter the effectiveness of the various coping mechanisms. Further, in a volatile and uncertain environment a certain amount of rigidity and inflexibility at the individual level may be beneficial at the group or team level. These kinds of tradeoffs need more exploration.

One of the major limitations of the two models described herein is the lack of even a primitive model of motivation. In the simple models we have described, the agents either are or are not liars. If they are liars then they do not change whom they lie to, what they lie about, or how often they lie. Rather, given the opportunity to provide information, they lie. Other than severe pathological cases, humans do not behave in this fashion. Nonetheless, even with this very simplified view, we find that it is not always possible for other agents to recognize a lie when they hear it. Other types of uncertainty can mask lies and make truth-tellers appear to be liars and vice-versa. Organizational theorists have been quick to point out that structures, processes, and routines can stabilize the task environment to the extent that it is possible for trust to develop and persist (Heide and Miner, 1992). Our argument, is truth may persist even in the absence of a stable environment if the various sources of uncertainty serve to mitigate each other.

Incorporating motives, although increasing model veridicality, will also increase the complexity of the results. Imagine, for example, what might happen if agents were motivated to lie only if that lie was likely to increase their own relative performance. In this cases, Liar agents would reserve their lies for others who were high performers. If the Liar happened to be the highest performer the agent might never lie. In group settings, we might find it behooves a group to have some group members who are liars so that the group as a whole can outperform the other group.

Finally, Barnes (1991) argued that excessive reliance on trust can result in exclusive and dysfunctional reliance on soft data rather than hard data in making decisions. The upshot of which would be for the organization to make decisions that appear whim based and are often wrong. One such types of soft data is the "opinion" of a local expert. In our study, if the environment is volatile the opinion of any one agent is unlikely to match the underlying reality. Thus trust in a single agent may result in major errors even when that agent is inherently honest. An important extension of the work we have provided would be to examine how much worse or better organizational decisions would be if they were made on the basis of expert opinion or data when the environment is uncertain and one source of that uncertainty is the proportion of agents in the organization that are lying.

Conclusion

Our focus has been on truth-telling. We have asked, what would the outcome be for organizations if agents were not truthful? Turner (1993) argues that one of the basic multi-agent problems is that of resource consumption. He goes on to argue that this problem exists for both artificial societies and human societies. Thus, computational studies about how to avoid or mitigate this problem will have value both to the design of computer systems and to societies. We note that issues of uncertainty, and trust, are also a basic multi-agent problems, and that such issues will affect issues of resource consumption. To the extent to which Turner's argument holds, then understanding the impact of truthfulness on organizational behavior is also important for organizations of both humans and artificial agents.

Our argument is simple. Trust is first and foremost a cognitive construct. Hence, trust is a computational construct. We can therefore explore in a meaningful way the relation between trust and organizational behavior using multi-agent models. The models and results in this paper are a small step in this direction. Although limited, these models do show that at the organizational level, lies are just one source of uncertainty and thus mechanisms for reducing uncertainty, whether or not they are aimed at liars may be performance enhancing. Further since the impact of lying on organizational performance is a function of the size of the organization, then the cost of tolerating liars will vary with organizational size.

References

Axelrod, Robert M. 1987. "The Evolution of Strategies in the Iterated Prisoner's Dilemma." Pp. 32-41 in Genetic Algorithms and Simulated Annealing, edited by L. Davis. London: Pitman.

Axelrod, Robert M. and Doug Dion. Dec 9, 1988. "The Further Evolution of Cooperation." Science 242(4884) 1385-1390.

Ballim, Afzal and Yorick Wilks. 1991. Artificial Believers: The Ascription of Belief. Hillsdale, NJ: Lawrence Erlbaum Associates.

Barber, B. 1983. The Logic of Trust. New Brunswick, NJ: Rutgers University Press.

Barnes, L.B. 1981, "Managing the Paradox of Organizational Trust." Harvard Business Review, 59: 107-116.

Bradach, J.L. and R.G. Eccles, 1989. "Price, Authority, and Trust: From Ideal Types to Plural Forms." Annua Review of Sociology, 15: 97-118.

Bromiley, P. and L.L. Cummings, 1995. Transaction Costs in Organizations with Trust in R. Bies, B, Sheppar and R. Lewicki (Eds.) Research on Negotiation in Organizations, Greenwich, CT: JAI Press.

Carley, K. M. 1991a. "A Theory of Group Stability." American Sociological Review 56(3): 331-354.

Carley, K. 1991b. "Designing Organizational Structures to Cope with Communication Breakdowns: A Simulation Model." Industrial Crisis Quarterly , 5: 19-57.

Carley, K. and Lin, Z. 1995. "Organizational Designs Suited to High Performance Under Stress." IEEE - Systems Man and Cybernetics, 25(1): 221-230.

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

Carley, K. and Prietula, M. 1994. ACTS Theory: Extending the Model of Bounded Rationality. In K. Carley ar M. Prietula (Eds.), Computational Organization Theory, Hillsdale, NJ: Lawrence Erlbaum.

Carley, K. and Prietula, M. In press. WebBots, Trust and Organizational Science. In M. Prietula, K. Carley an L. Gasser (Eds.), Simulating Societies: Computational Models of Institutions and Groups. Cambridge, MA: AAAI/MIT Press.

Carley, K., Prietula, M. and Lin, J. In press. Design versus Cognition: The Interaction of Agent Cognition and Organizational Design on Organizational Performance. In R. Conte and E. Chattoe (Eds.), Evolving Societies: The Computer Simulation of Social Systems.

Carley, K., J. Kjaer-Hansen, J., Newell, A. and Prietula, M. 1992. Plural-Soar: A Prolegomenon to Artificial Agents and Organizational Behavior. In M. Masuch and M. Warglien (Eds.), Artificial Intelligence in Organization and Management Theory, Amsterdam: North-Holland.

Carley, K.M. and David M. Svoboda, 1996, "Modeling Organizational Adaptation as a Simulated Annealing Process." Sociological Methods and Research, 25(1): 138-168.

Corazzini, R. 1977. Trust as a Complex Multi-Dimensional Construct. Psychological Reports, 40: 75-80.

Cummings, L.L. and P. Bromiley, 1996, The Organizational Trust Inventory (OTI): Development and Validatio in R.N. Kramer and T.R. Tyler (Eds.) Trust in Organizations, Thousand Oaks, CA: Sage.

Frazee, Valerie (1996). Employees value workplace relationships. Personnel Journal, 75(6): 25.

Heide, J.B. and A.S. Miner, 1992, The Shadow of the Future: Effects of Anticipated Interaction Frequency of Contact on Buyer-Seller Cooperation, Academy of Management Journal, 35(2): 256-291.

Johnson-George and Swap 1982

Kaufer, David S. and Kathleen M. Carley. 1993. Communication at a Distance: The Effect of Print on Socio-Cultural Organization and Change. Hillsdale, NJ: Lawrence Erlbaum Associates.

Kephart, Jeff O., Huberman, Bernardo A. and Tad Hogg. 1992. "Can Predictive Agents Prevent Chaos?" Pp 41-55 in Economics and Cognitive Science, edited by P. Bourgine and B. Walliser. Oxford: Pergamon Press.

Klein, Daniel B. (Ed.) 1997. Reputation : studies in the voluntary elicitation of good conduct. Ann Arbor, MI : University of Michigan Press

Laird, J., Newell, A. and Rosenbloom, P. 1987. Soar: An Architecture for General Intelligence, Artificial Intelligence, 33, 1-64.

Lerch, J., Prietula, M. and Kulik, C. 1997. The Turing Effect: The Nature of Trust in Machine Advice. In P. Feltovich, K. Ford and R. Hoffman (Eds.), Expertise in Context: Human and Machine, Cambridge, MA AAAI/MIT Press.

Lerch, J., Prietula, M., Kim, J., and Buzas, T. 1997. Unraveling the Turing Effect: Measuring Trust in Machir Advice. Working Paper, Graduate School of Industrial Administration, Carnegie Mellon University, Pittsburgh PA.

Lewis, J.D. and A.J. Weigert, 1985a, Trust as a Social Reality, Social Forces. 63: 966-985.

Lewis, J.D. and A.J. Weigert, 1985b, Social Atomism, Holism, and Trust. Sociological Quarterly, 26: 455-471

Lin, Z. and K. Carley, 1993, "Proactive or Reactive: An Analysis of the Effect of Agent Style on Organization Decision Making Performance." International Journal of Intelligent Systems in Accounting, Finance an Management , 2(4): 271-288.

Luhmann, N., 1979, Trust and Power, New York: Wiley.

Macy, Michael W. 1991a. "Learning to Cooperate: Stochastic and Tacit Collusion in Social Exchange." Americ Journal of Sociology, 97(3): 808-43.

Macy, Michael W. 1991b. "Chains of Cooperation: Threshold Effects in Collective Action." American Sociological Review 56: 730-747.

McAllister, Daniel J. 1995. Affect- and cognition-based trust as foundations for interpersonal cooperation in organizations. Academy of Management Journal, 38(1): 24-59.

Nicholas, Ted. 1993, Secrets of entrepreneurial leadership : building top performance through trust and teamwork. Chicago, Ill. : Enterprise Dearborn

Prietula, M. and Carley, K. 1994. Computational Organization Theory: Autonomous Agents and Emergent Behavior. Journal of Organizational Computing, 4(1), 41-83.

Rempel, J.K. and J.G. Holmes, 1986, How do I trust thee? Psychology Today, Feb: 28-34

Rempel, J.K., J.G. Holmes and M.P. Zanna, 1985, Trust in Close Relationships, Journal of Personality and Social Psychology, 49: 95-112.

Ring, P.S. and A.H. Van de Ven, 1992, Structuring Cooperative Relationships Between Organizations, Strategi Management Journal, 13: 483-498.

Rotter, J.B. 1971, Generalized Expectancies of Interpersonal Trust. American Psychologist, 26: 443-452.

Shapiro 1990

Shapiro, S.P. 1987, The Social Control of Impersonal Trust. American Journal of Sociology. 93(3): 623-658.

Shoham, Yoav and M. Tennenholtz. 1994. Co-learning and the evolution of social activity. Stanford University Dept. of Computer Science; STAN-CS-TR-94-1511 Report , Stanford, CA.

Skvoretz, John and Thomas J. Fararo. 1995. "The Evolution of Systems of Social Interaction." Pp. 275-299. Current Perspectives in Social Theory vol. 15, edited by B. Agger . Greenwich, CT: JAI Press.

Turner, Roy M. 1993. "The Tragedy of the Commons and Distributed AI Systems." University of New Hampshire, Dept. of Computer Science, Technical report series 93-01. Durham, NH.

Zaheer, A. and N. Venkatraman, 1995, Relational Governance as an Interorganizational Strategy: An Empirical Test of the Role of Trust in Economic Exchange, Strategic Management Journal, 19(5): 373-392.

Zaheer, A., B. McEvily, and V. Perrone (1998). "Does Trust Matter? Exploring the Effects of Interorganization and Interpersonal Trust on Performance," Organization Science (forthcoming).

Zucker, L.G. 1986, Production of Trust: Institutional Sources of Economic Structure, 1840-1920 in L.L. Cummings and B.M. Staw (Eds.) Research in Organizational Behavior, 8, Greenwich, CT: JAI Press, Inc. 53-111.

Tables

None

Figure Captions

Figure 1. The Generic Search Task

Figure 2. Cooperativeness by Organizational Size

Figure 3. Total Moves by Organizational Size

Figure 4. Advice by Organizational Size

Figure 5. Impact of Group Proportion on Activity

Figures

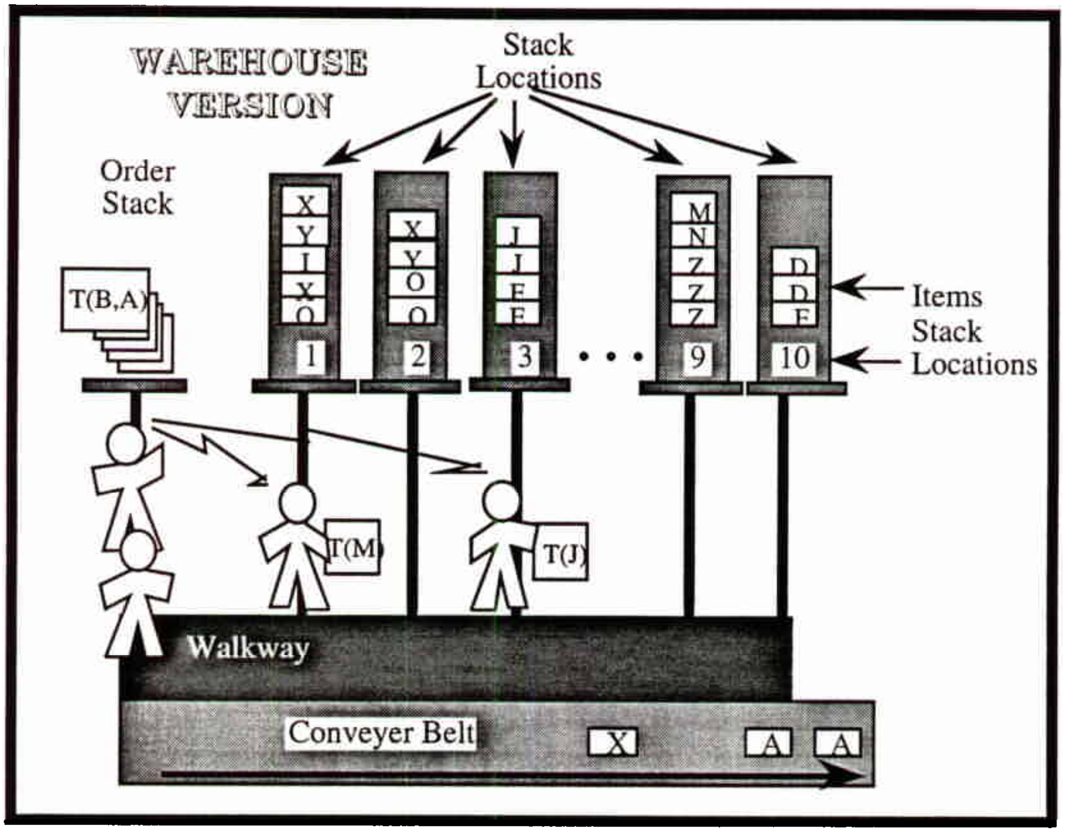

WAREHOUSE VERSION
Stack Locations
Order Stack
T(B,A)
X Y I X O
X Y O O
J J F F
M N Z Z Z
D D F
1
2
3
9
10
Items
Stack Locations
T(M)
T(J)
Walkway
Conveyer Belt
X
A
A

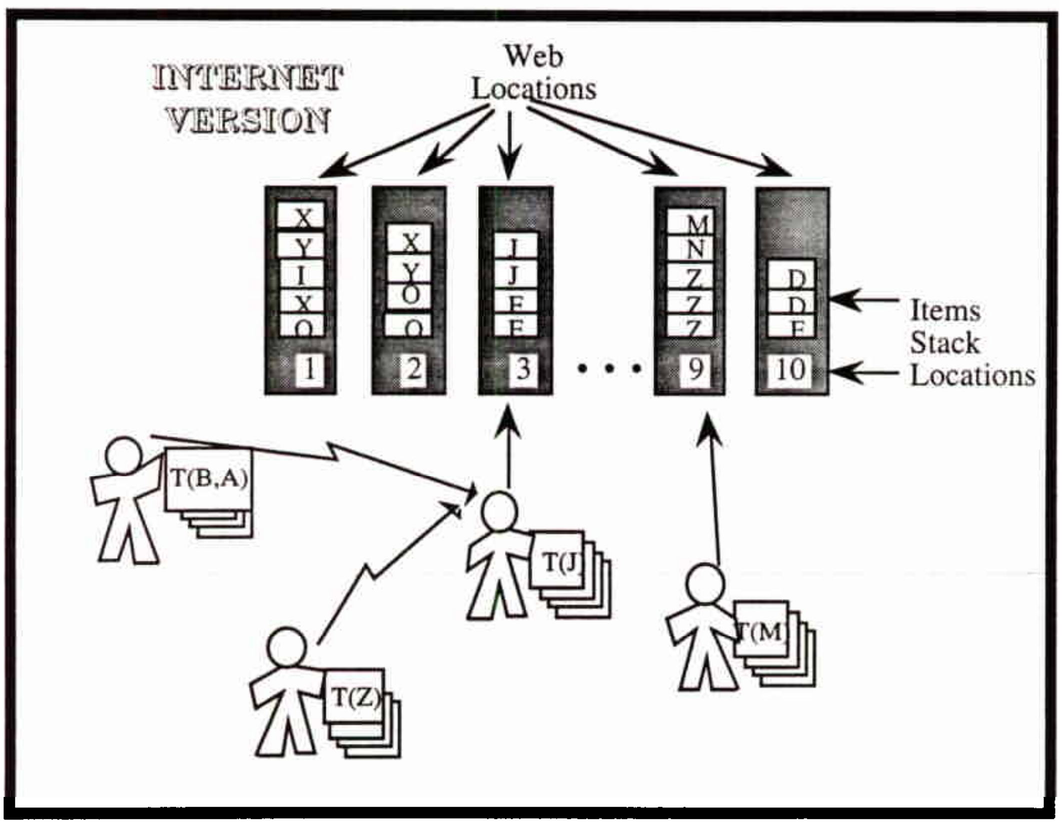

INTERNET VERSION
Web Locations
X Y I X O
X Y O O
J J F F
M N Z Z Z
D D F
1
2
3
9
10
Items
Stack Locations
T(B,A)
T(J)
T(M)
T(Z)

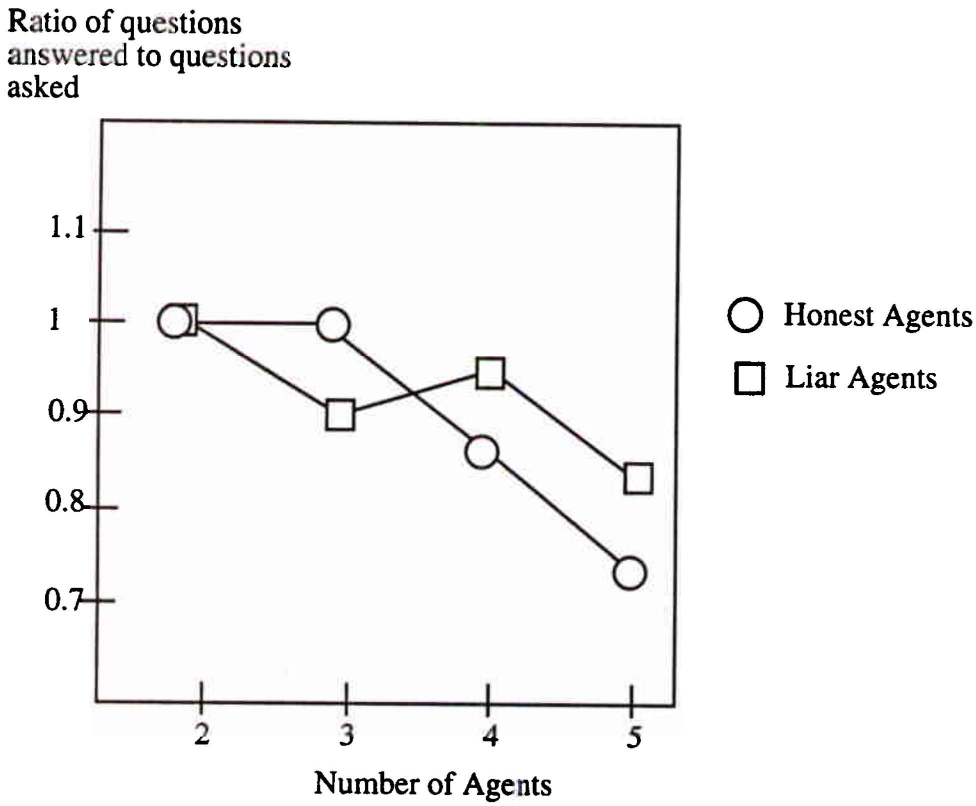

Ratio of questions answered to questions asked
1.1
1
0.9
0.8
0.7
2
3
4
5
Number of Agents
Honest Agents
Liar Agents

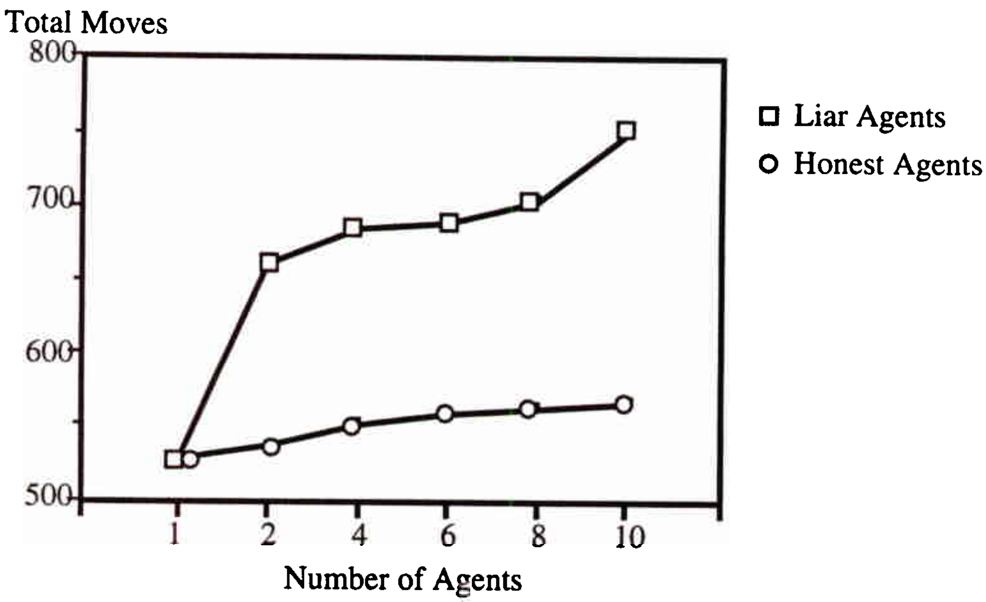

Total Moves
800
700
600
500
1
2
4
6
8
10
Number of Agents
Liar Agents
Honest Agents

Table 2

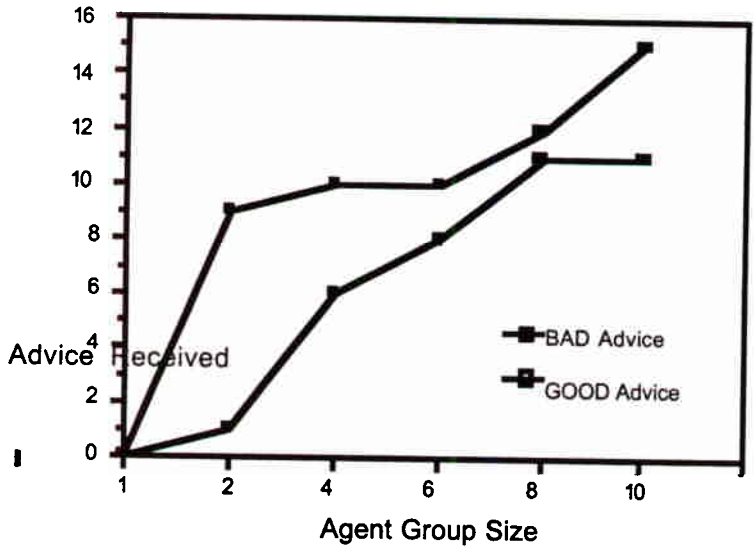

[[[Mike can you redraw Figure 4, Figure 2 is a guesstimate based on figures from old paper, if you have real data please double check the figure]]]

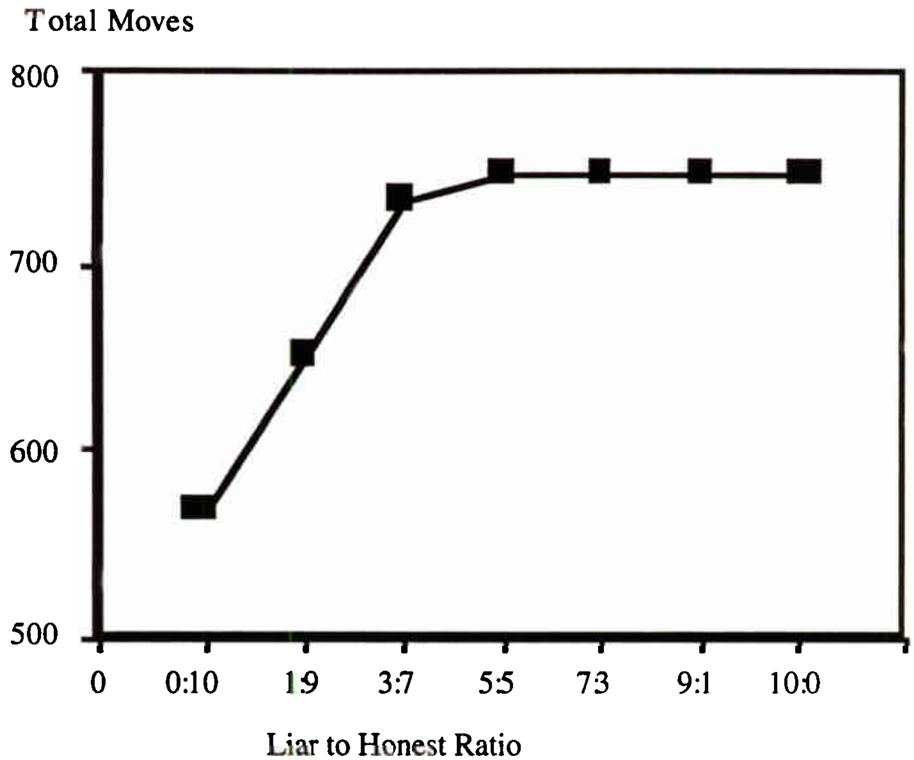

Total Moves
800
700
600
500
0
0:10
1:9
3:7
5:5
7:3
9:1
10:0
Liar to Honest Ratio

Table 4

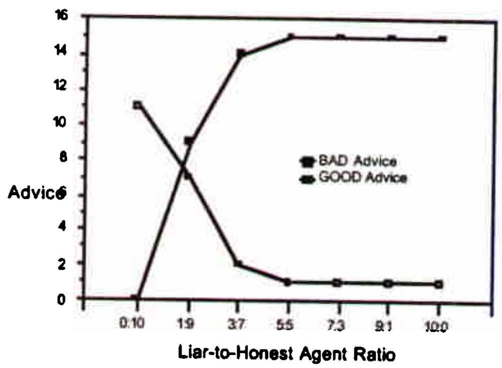

Table 5