The Impact of Knowledge Misrepresentation on Organization Performance Dynamics

Terrill L. Frantz
Carnegie Mellon University
terrill@cs.cmu.edu

Kathleen M. Carley
Carnegie Mellon University

Abstract

A computer-based simulation of a multi-agent networked, virtual organization is experimented with to explore the impact of human-agents misrepresenting information to others when performing an organization-level decision task. It is a common-place phenomenon that humans make misrepresentations (or errors) in judgment (accidental or otherwise) when evaluating information and then passing their resulting conclusion on to others, which can affect the accuracy of organization-level decisions. This can have great relevance in organizations under stress or undergoing large-scale change, e.g., as in the case of a merger or acquisition, as humans can make even more misrepresentations. For this study, repeated simulations are executed with different levels of actor knowledge-misrepresentation to assess the impact on the dynamics of organization performance. Organization performance is quantified by the accuracy of a knowledge-based task, repeated over multiple time periods. Each task is a single binary-decision, which is derived by actor majority vote. Actors base their vote on the partial, factual knowledge that each holds, but their best-effort vote can be reversed according to the actor's tendency to misrepresent the totality of the accurate information they truly hold. The underlying computational model is based on the Construct multi-agent network model, which is extended by incorporating a Monte-Carlo process on the actors' tendency to misrepresent their true knowledge of the actual facts. The impact on the dynamics of the organization performance over time are reported and discussed. The findings indicate that differing levels of misrepresentation by actors in an organization do indeed have an impact the organization performance, but with differing temporal dynamics at different misrepresentation levels.

Contact:
Terrill L. Frantz
School of Computer Science
Carnegie Mellon University
Pittsburgh, PA 15213

Email: terrill@cs.cmu.edu

Key Words: knowledge misrepresentation, misjudgment, decision-making, majority voting, Construct, organization behavior, multi-agent simulation

Support: This paper is part of the Dynamics Networks project in CASOS (Center for Computational Analysis of Social and Organizational Systems, http://www.casos.cs.cmu.edu) at Carnegie Mellon University. This work was supported in part by the Office of Naval Research

(ONR), United States Navy Grant No. 9620.1.1140071 on Dynamic Network Analysis under the direction of Rebecca Goolsby and by the Air Force Office of Sponsored Research (MURI: Cultural Modeling of the Adversary, 600322). Additional support on measures was provided by the DOD and the NSF. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the National Science Foundation or the U.S. government.

The Impact of Knowledge Misrepresentation on Organization Performance Dynamics

Terrill L. Frantz & Kathleen M. Carley

Introduction

Humans have mistakes embedded in their judgment when making even the simplest of decisions. They can make communication errors in expressing their view and they can even sometimes be blatantly deceit. Antidotally, their misrepresentation of true fact, or truth, in judgmental tasks must have influence on the performance of an organization--one way or another. During an organizational merger, the workforce is possibly under even greater emotional and performance stress and pressure, thus actors may be even more likely to make mistakes and misrepresent fact, thus the increasing level of misrepresentation--unintentionally or otherwise--when communicating factual and judgmental information to others. This line of questioning is not completely new as earlier studies have been conducted on information distortion in organizations under incomplete information (Carley & Lin, 1997). This paper aims to build on such earlier work.

It has been shown that many people are dishonest when given the chance (Hartshorne & May, 1928) and even are so in the workplace. A recent survey indicates that 1 in 5 workers admit to having told lies at work (Gurchiek, 2006) and some researchers even cite managers saying it is okay to lie to customers, but not to the boss (Management Issues News, 2005). Lying has been called a fact of life in the workplace (Penson, 1997). Given enough incentive, the uncomfortable tension and cognitive dissonance sometimes created by external incentives, e.g., monetary rewards, people will be led towards misrepresentation in their behavior, even by convincing themselves of the falsehood (Festinger & Carlsmith, 1959).

Incorrect information has a hefty impact on organization performance (Lin, 1996) as it has been shown that there is indeed a link between individuals' knowledge and group performance (Ashworth & Carley, 2006). While, even during a crisis a high performing organization will continue to perform well (Lin, Zhao, Ismail, Carley, 2006), events such as organizational mergers can cause stress in people and to the organizations involved (Cartwright & Cooper, 1990, 1993, 1994). Stress affects the group decision making and thus impacts the ensuing organization performance (Driskell & Sala, 1991). Deceit is a potential distress-reliving mechanism that people can fall into (Gover, 1993).

To investigate knowledge misrepresentation in organizations, techniques available in the field of computational organization science will be used. Computational organization science, which has been called "a new frontier" (Carley, 2002b), provides a perspective and a rich and developing set of techniques to study organization behavior. Collections of intelligent, adaptive agents are studied in a virtual, i.e. computer-based, environment. The utlizaiton of underlying computational models can help us explore organizations (Cyert & March, 1963) and understand basic human information processing behavior in organizations (Simon, 1973). The computational modeling of organizations is influential to the development or organization theory (Carley, 1995b) and helps to bring alive models in the active systems realm (Burkov & Enaleev, 1994).

The Construct model is a well published (Carley, 1990, 1991a, 1991b, 1995a, 1999, 2002a; Caley & Hill, 2001; Carley & Krackhardt, 1996; Carley & Schreiber, 2002; Kaufer & Carley, 1993; Schneider & Graham, n.d.; Schreiber & Carley, 2003) computational model of organization behavior. Construct embodies social theories such as Structuration (Giddens, 1979, 1984) and Adaptive Structuration Theory (AST) (DeSanctis & Poole, 1994) in which it is put forward that: (a) social interaction is not merely the sum of actors' micro-activity, (b) it is actually the actions of individual actors that produce the structure of a group, and (c) such micro-level interaction meaningfully involves today's new technologies as a factor in the individual agents' behavior and interactions leading to the greater group structure (as per AST).

The Construct model incorporates a relational network to represent the social network of the agents' possible interactions as well and to represent the transactive knowledge held by each agent. As conceptualized by Festinger (1950) and later, Granovetter (1973), the knowledge held by an agent is gained through social interactions with other agents. . This network structure of agents and knowledge implements the organizational meta-matrix concept (Carley, 2002a; Krackhardt & Carley, 1998) which has been used in other virtual experiments, e.g., Carley and Hill (2001). The meta-matrix is a multiple-network social network construct that can be readily analyzed using Dynamic Network Analysis (Carley, 2003) techniques.

It has been shown that transactive memory play an important role and has a positive impact on group performance (Ren, Carley & Argote, 2006) and combined with social interaction it results in organizational cognition (Hutchins, 1991, 1995). The decision making process in Construct includes an equal-weighting scheme of actors' judgment on the binary decision, which is expected to be less accurate if the individual actors hold any decision bias (Einhorn, Hogarth, & Klempner, 1977), than if no actor bias is present. There has been much research into the decision making process and the voting method in particular has received focused attention in the political field, e.g., Riker & Ordeshook (1968) and Cranor (1996).

Method

The Construct model is used as the foundation and software host for this simulation. Construct incorporates a multi-agent network model for examining information diffusion and organization change. An organization's social structure evolves and changes according to the underlying knowledge network in the organization. Construct uses the meta-matrix approach to representing agents, knowledge and their relationships. Agents hold various bits of information and use the information to render an opinion on a binary-decision task. The organization of agents then each submits their vote in order for the organization to classify the binary-decision. Majority voting rules are utilized in this herein. Specifically for the experiment Construct software and extended to include a process by which the actors can have a probability to misrepresent their true vote during the process of polling all actors to determine the organization's final classification in the binary task.

Table 1. Input Parameters for Virtual Experiment

Parameter	Value	Meaning
FIXED INPUTS:		
Number of agents	101	Set to an odd number so group voting cannot result in a tie.
Number of tasks	1000	
Size of each task	101	Number of binary bits for each task.
Number of time periods	250	For each run
Number of runs	5	Number of times the simulation is run for each unique set of parameter settings.
VARIABLE INPUTS:		
Misrepresentation rate	Fixed or random uniform distribution assignment of rate to individual agents, using variety a of settings between 0 and 1	A total of 75 combinations of the setting for the misrepresentation rate we run for this experiment.

As shown in Table 1, for this study, each organization--each simulation run--consists of 101 agents, 101 facts for each binary-decision, and 1000 independent decision tasks. For this experiment, the Construct feature for forgetting is turned off, so as to keep the results understandable. Each simulation run is independent of all others. For each specific parameter setting, 5 runs were executed for 250 time periods. The accuracy of the binary decision task for each time period was averaged across the 5 same-parameter runs, resulting in a single vector with 250 accuracy values between 0 and 1 for each specific parameter setting. The output parameter for the experiment is shown in Table 2. This single output, one value for each time period, is the accuracy of the 1000 decision tasks in each period. This accuracy is considered the performance metric for the organization.

Table 2. Output Parameters for Virtual Experiment

Parameter	Meaning
Decision Accuracy	The performance measure of the organization collective; a percentage value between 0 and 100.

The knowledge in Construct is represented as a binary string 101 bits long. Each agent knows the true value for parts of this string. Thus, their individual knowledge can be thought as having a network consisting of a "knows" relationship between the agent and the knowledge bit(s) available to the agent, which in the aggregate, forms an agent x knowledge matrix for the entire organization.

The organization is faced with specific binary classification tasks (for this experiment there are 1,000 independent decision-tasks), each consisting of a true value based on the count of the odd or even number of the 101 bits in the corresponding binary string. To accomplish each task, the agents vote on their own classification of the same task given the knowledge that they have according to the agent x knowledge matrix. The agents each have an equal vote for each task and the organization-level decision is based on majority-rules. The classification task is scored as correct or incorrect depending on the organization-level decision being the same as, or different that the true value, respectively. The accuracy for all 1,000 tasks is summed giving a percentage accuracy score for the organization for each time period. This accuracy score is considered the performance of the organization for each time period and reported in the Results section of this paper.

Results

Of the 75 parameter-specific simulation runs, only a representative few will be reported in detail within this section as all 75 instances are each designed in an incremental fashion that can be fully represented by a representative subset of results. Logistically, the entire suite of simulations took two days to execute on a Linux-based, 4-CPU computer and it took about a day to develop the reporting capabilities and process.

Figure 1 shows the outcome dependent variable, the organization performance, for the full range of possible misrepresentation levels. The chart shows increments of 10 percent between 0% and 100% for the input independent misrepresentation-level. The misrepresentation level indicates the parameter setting for each and every actor's probability of misrepresenting their vote during the organization's majority-rule decision-making process. The resulting lines shown on the chart are aligned, in line, with the levels listed in the legend of the chart, that is, the top line is for the 0% level and the very bottom lines (there are several transposed over one another) are indicative of the higher levels, specifically 60% through 100%.

The Figure 1 graph shows that an organization with every actor representing their true vote in every task begins with an approximate 70% accuracy rate that ultimately reaches 100% accuracy in about 75 time periods. The 10% misrepresentation rate, which means that each and every actor will misrepresent their judgment in 10%, on average, of the tasks, has little impact on the organization's overall performance. The 20% and 30% misrepresentation rate levels begin to noticeably affect the performance, with more of an impact early on, but by approximately time period 100, the 100% accuracy level is reached.

However, once the misrepresentation rate reaches and is above 40%, the organization performance is impacted greatly. At 40%, the time period 1 performance is a meager 25%--less than random chance--which is 1/3 the accuracy performance of a 0% organization, i.e., an organization without any knowledge misrepresentation. Further, the performance level is now asymptotic to 100%, reaching a high of only approximately 94%. When the misrepresentation rate is in the neighborhood of 50%, the organization flat-lines at about 12% accuracy through the entire 250 time periods. Below the 50% misrepresentation rate, the organization essentially becomes utterly ineffective in its decision-making proficiency.

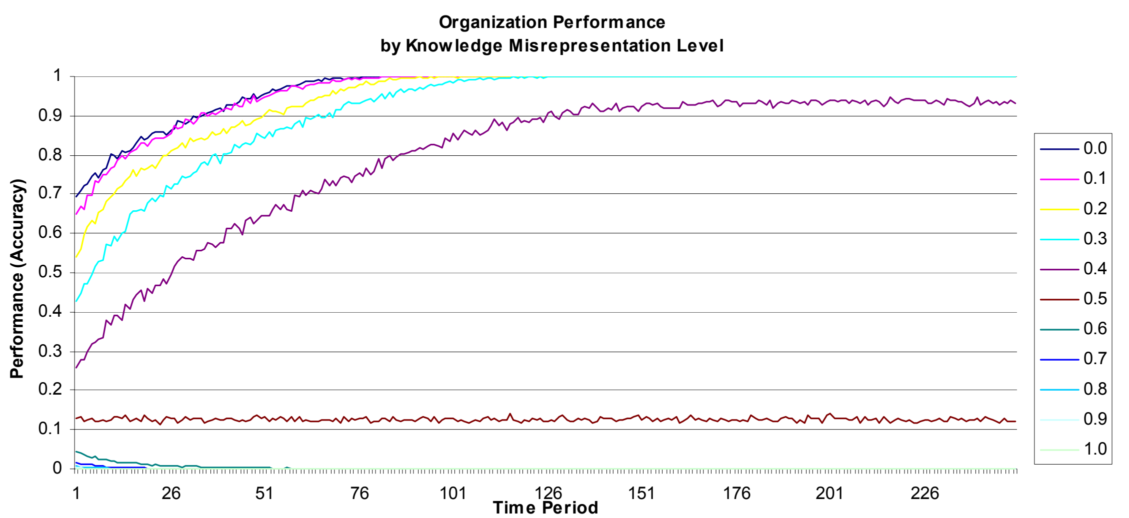

Figure 1. Performance at different uniform levels of knowledge misrepresentation by actors.

To explore in greater detail the tipping–point nature of the performance outcome occurring between the 40% and 50% misrepresentation-rate levels, a series of more detailed rates in this range are presented in Figure 2. This chart shows the rather similar nature of the temporal dynamics for the various misrepresentation-rate levels, but at significantly different level of performance outcome when misrepresentation rates are between 41% and 49%. For example, the 41% rate is asymptotic to the 91% accuracy; the 42% level is asymptotic to 85% accuracy; 43% asymptotic to 80% accuracy; and so on down to 49% rate, which is asymptotic to a relatively poor 20% accuracy level.

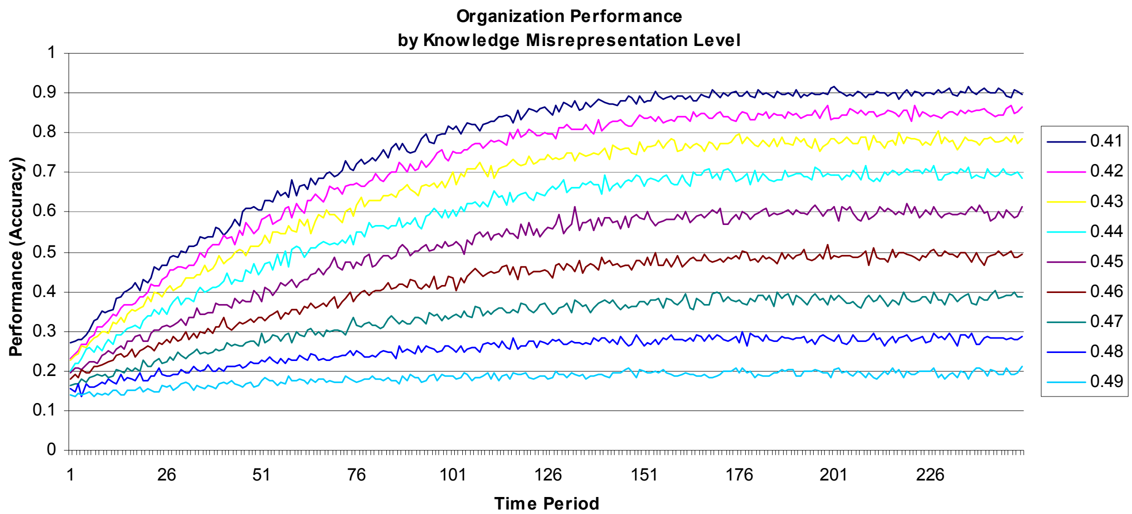

Figure 2. Performance at different uniform levels between 41% and 49% of knowledge misrepresentation by actors.

The experiments run with an input-range for the misrepresentation rate, rather than the fixed rate across actors provides a more realistic scenario that is closer to more atypical real-world organization realities; that being that all actors likely do not have the exact same misrepresentation rate, instead the actors likely each lie at a point within a range of possible misrepresentation rates. Figure 3 shows five broadly representative and illustrative settings for the misrepresentation rate ranges run in this experiment. In each of these simulations, the specific misrepresentation rate for each agent is determined by a uniformly random Monte Carlo draw, bounded by the range, where the lower bound is the minimum value and the upper bound is the maximum value of the range.

As with prior figures (shown above), the lines in the chart and the order of the legend are associated exactly, that is the top line is the top label in the legend, namely 0% to 10%, and, bottom line is the bottom in the legend, namely the 0% to 100% misrepresentation-rate range. The various start and end levels of the performance outcome values and the associated dynamics of the performance over time for each individual performance lie closely to one another, when compared to the fixed misrepresentation rate experiments (see Figure 1). When the mean value of the range for each experiment is matched with the corresponding fixed value the performance outcomes results are similar—this correlation is due to the nature of the probability assignment process for the group or actors as a whole. Specifically, the uniform random Monte Carlo draw between a number range, in large numbers, should be equivalent to a fixed misrepresentation rate equaling the midpoint of the range.

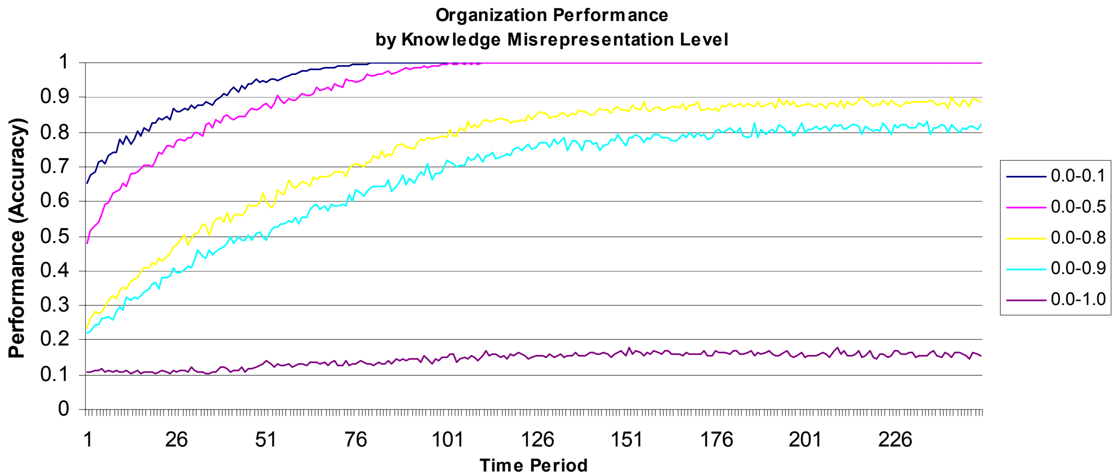

Figure 3. Performance at different average level (via different ranges) of knowledge misrepresentation by actors.

However, contrary to a priori expectations, there does appear to be some variation when the different width ranges of the same mean misrepresentation rates are used. Figure 4 shows five different ranges that all have the same mean midpoint value, i.e., 45%. The performance for each at time 1 is consistent, however over time, the different ranges appear to differ in the level of accuracy performance they are asymptotic with. Note: in this--Figure 4--unlike the others, the performance outcome value lines in the chart do not match up with the legend. In this case, the 20%-70% range has the poorest accuracy performance (it is the bottom line). This chart shows that, for ranges with the same midpoint at the 45% misrepresentation level, the wider the spread or the greater the range, the better the performance.

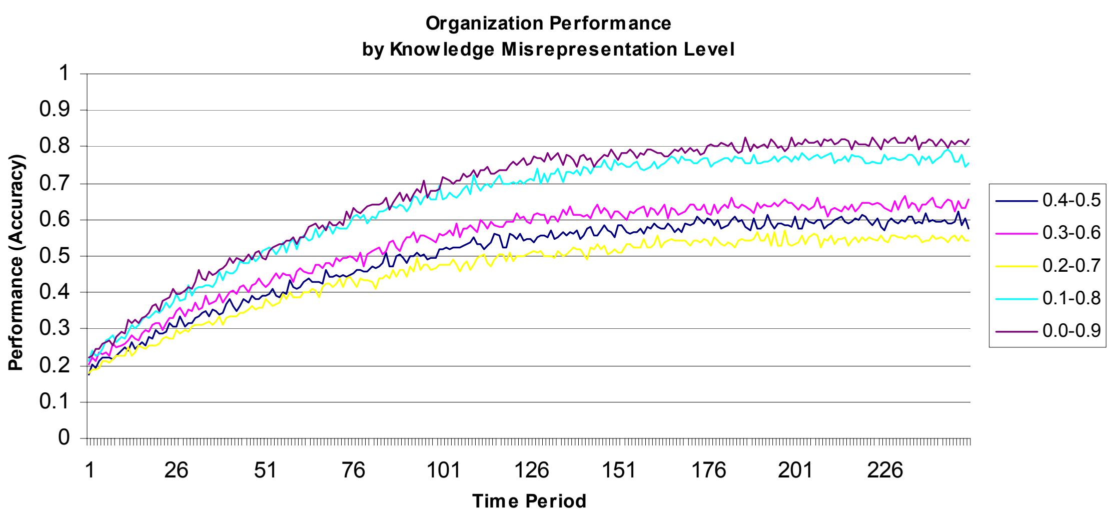

Figure 4. Performance at same mean/midpoint value (45%), but with different ranges of knowledge misrepresentation by actors.

Discussion

Knowledge misrepresentation, unintended or otherwise, is a reality in human organizations and warrants more attention in the realm of organization-behavior research. The experiments conducted and the discussion in this paper explores the impact of individuals' knowledge misrepresentation on the overall organization performance.

The 75 independent simulations present three major points that are meaningful to our understanding of organizational behavior. First, the results provided in detail in the Results section indicate that a mean level of actor misrepresentation below 40% does indeed impact the organization performance as one may anecdotally expect, however the impact over time ultimately disappears completely. That is, initially the increase misrepresentation rate (up to the 40% level) impacts the organization's accuracy performance only in the first 100 time periods. After the initial 100 time periods, the performance accuracy still reaches 100%, regardless of the starting point, or the underlying misrepresentation rate. The organization appears to reach peak performance regardless to the level of knowledge misrepresentation.

Second, between a 40% and 50% misrepresentation rate, the organization's performance dynamics change quite visibly. While the performance outcomes for misrepresentation rates in this range de novo are practically the same, the performance levels over time reach substantially different asymptotic levels of performance accuracy; that is to say that the organization's peak accuracy performance rate differs greatly when the misrepresentation rate is within this 40-50% range. Each additional percentage point of the misrepresentation rate impacts the peak performance level for the organization in an appreciably increasing negative manner.

Third, a misrepresentation rate greater than 50% in an organization could be considered disastrous to the organization's accuracy performance. A knowledge misrepresentation rate at 50% or greater, severely limits the maximum level of performance to no more than 18% accuracy, at best. Misrepresentation rates slightly higher than the 50% drop the outcome accuracy performance to a tragic level of less that 1% in very few time periods. Anecdotally, an organization in which the members, for whatever reason, misrepresent their true knowledge in a decision-making task, will perform poorly; so poorly that the very raison d'être of the organization construct is in jeopardy as a single accurate actor may perform better than the collective organization.

The results of this experiment demonstrate the repercussions and impact of different levels of knowledge misrepresentation in an organization. The findings allow managers to begin to quantify the impact of this phenomenon on their organization's performance and possibly provide a practical tolerance level that can be perhaps actively managed. The results of 75 different combinations of probabilistic distributions of actors' knowledge misrepresentation levels in a decision-making organization herein are surely valuable, but such information may merely be a step closer to understanding the phenomenon and may serve as a seed to greater leaps forward, rather than being functional in real-world organizations. More work undoubtedly needs to be carried out in this area or research.

Limitations

However valuable the findings of this paper are, before reaching specific conclusions from this research it is imperative that some caveats be presented. The most critical limitation of this study to be understood, is that the method of voting in the Construct model, and thus used in this simulation, is a simplistic majority-rule method (see Ragenwetter, Marley & Grofman, 2002) and may be an oversimplification for strategic decision-making (Schwenk, 1984). While this method may in-fact be used in some real-world situations, this pure decision-making method may be considered as being too simplistic for real-world application of the findings discussed herein. More realistic approaches such as collective decision making in a hierarchy (Berg & Paroush, 1998) need to be investigated. Certainly, experiments employing a more real-world decision making process are necessary to fully understand the impact of knowledge misrepresentation on organization performance.

Further, this research does not address the complex issue of trust in organizations. Inevitably as any actor increasingly misrepresents the truth, over time, the other actors will discount the deceitful opinion. Thus others will ultimately ignore the untrustworthy vote, and possibly expel the deceitful out of the organization; which additionally impacts the organization performance via the phenomenon of turnover. This is an example of this author's call for more realistic, thus more complicated, decision-making models to be utilized in this line of research.

Lastly, the various aspects and complexities of organizational ecology (see Hannan & Freeman, 1989) including organization turnover and forgetting (Douglas, 1986; Reason, 1990) have been disabled in the Construct model for this study, in order to keep this initial analysis simple and somewhat straight-forward. It should be clear that such complex dynamics such as forgetting, et cetera, in an organization are natural and ubiquitous and therefore should be included in future more advanced studies on the impact of knowledge misrepresentation.

References

Ashworth, M. J. & Carley, K. M. (2006). Who you know vs. what you know: The impact of social position and knowledge on team performance. *Journal of Mathematical Sociology, 30,* 43-75.

Berg, S., & Paroush, J. (1998). Collective decision making in hierarchies. *Mathematical Social Sciences, 35,* 233-244.

Burkov, V.N., & Enaleev, A.K. (1994). Stimulation and decision-making in the active system theory. *Mathematical Social Sciences, 27*, 271-291.

Carley, K. M. (1986a). An approach for related social structure to cognitive structure. *Journal of Mathematical Sociology, 12*, 137-189.

Carley, K. M. (1986b). Knowledge acquisition as a social phenomenon. *Instructional Science, 14*, 381-438.

Carley, K. M. (1990). Group stability: A socio-cognitive approach. In E. Lawler, B. Markovsky, et al (Eds.), Advances *in group processes: Theory & research, VVI*, pp.1-44. Greenwich, CN: JAI Press.

Carley, K. M. (1991a). Growing up: The development and acquisition of social knowledge. In J. Howard & P. Callero (Eds.), *The Self-Society Dynamic: Cognition, Emotion, and Action,* (p. 72-105). Cambridge, England: Cambridge University Press.

Carley, K. M. (1991b). A theory of group stability. *American Sociological Review, 56*, 331-354.

Carley, K. M (1995a). Communication technologies & their effect on cultural homogeneity, consensus, & the diffusion of new ideas. *Sociological Perspectives, 38*, 547-571.

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

Carley, K. M. (1996). Communicating new ideas: The potential impact of information and telecommunication technology. *Technology in Society, 18*, 219-230.

Carley, K. M. (1991). On the evolution of social and organizational networks. In S. B. Andrews and D. Knoke (Eds.), special issue of *Research in the Sociology of Organizations on Networks In and Around Organizations, 16,* 3-30.

Carley, K. M. (2002a). Smart agents and organizations of the future. In: L. Lievrouw and S. Livingstone (Eds.). *The Handbook of New Media*, (pp. 206-220). Thousand Oaks, CA: Sage.

Carley, K. M. (2002b). Computational organization science: A new frontier. *Proceedings of the National Academy of Sciences of the United States of America, 99,* 7257-7262.

Carley, K. M. (2003). Dynamic Network Analysis. In R. Breiger, K. Carley, and P. Pattison, (Eds.), *Dynamic Social Network Modeling and Analysis: Workshop Summary and Papers*, pp. 133-145. Washington, DC: Committee on Human Factors, National Research Council, National Research Council.

Carley, K. M. & Hill, V. (2001). Structural change and learning within organizations. In A. Lomi and. E. R. Larsen (Ed.). *Dynamics of Organizations: Computational Modeling and Organization Theories*. Menlo Park, CA: MIT Press/AAAI.

Carley, K. M. & Krackhardt D. (1996). Cognitive inconsistencies and non-symmetric friendship. *Social Networks, 18*, 1-27.

Carley, K. M. & Lin, Z. (1997). A theoretical study of organizational performance under information distortion. *Management Science, 43*, 976-997.

Carley, K. M. & Schreiber, C.. Information technology and knowledge distribution in C3I teams. 2002 Command and Control Research and Technology Symposium. Monterey, CA, 2002. p.

Cartwright, S. & Cooper, C. L. (1990). The impact of merger and acquisitions on people at work: Existing research and issues. British Journal of Management, 1, 65-76.

Cartwright, S. & Cooper, C. L. (1993). The psychological impact of merger and acquisition on the individual: A study of building society managers. Human Relations, 46, 327-345.

Cartwright, S. & Cooper, C. L. (1994). The human effects of mergers and acquisitions. In C. L. Cooper & D. M. Rousseau (Eds.), Trends in organizational behavior, pp. 47-61. New York: Wiley.

Cranor, L. F. (1996). Declared-Strategy Voting: An Instrument for Group Decision-Making. Washington University Dissertation.

Cyert, R., & March. J. G. (1963). A behavioral theory of the firm. Englewood Cliffs, NJ: Prentice-Hall.

DeSanctis, G., & Poole, M. S. (1994). Capturing the complexity in advanced technology use: Adaptive structuration theory. Organization Science, 5, 121-147.

Douglas, M. (1986). How institutions think. Syracuse, NY: Syracuse University Press.

Driskell, J. E., & Sala, E. (1991). Group decision making under stress. Journal of Applied Psychology, 76, 473-478.

Einhorn, H. J., Hogarth, R. M., & Klempner, E. (1977). Quality of Group Judgment. Psychological Bulletin, 84, 158-172.

Festinger, L. (1950). Informal social communication. Psychology Review, 57, 271-282.

Giddens, A. (1979). Central problems in social theory. Berkeley, CA: University of California Press.

Giddens, A. (1984). The constitution of society: Outline of the theory of structuration. Berkeley, CA: University of California Press.

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

Gurchiek, K. (2006). 1 in 5 workers admit to telling lies at work. SHRM. www.shrm.org/hrnews_published/archives/CMS_016154.asp

Hartshorne, H. & May, M. A. (1928). Studies in the Nature of Character. New York: Macmillan.

Hutchins, E. (1991). The social cognition of distributed cognition. In L. B. Resnick, J. M. Levine and S. D. Teasley (Eds.), Perspectives on Socially Shared Cognition. Washington, D.C.: American Psychological Association.

Hutchins, E. (1995). Cognition in the wild. Cambridge, MA: The MIT Press.

Krackhardt, D. & Carley, K. (1998). A PCANS model of structure in organization. In Proceedings of the 1998 International Symposium on Command and Control Research and technology (pp. 113-119). Monterrey, CA.

Lin, Z. (1994). Organizational performance – Theory and reality. PhD dissertation, Carnegie Mellon University.

Lin, Z., Zhao, X., Ismail, K., & Carley, K. M. (2006). Organizational design and restructuring in response to crises: Lessons from computational modeling and real-world cases. Organization Science, 15, 598-618.

Management Issues News (25 Feb 2005). It's OK to lie-but not to the boss. www.management-issues.com

Penson, P. (1997). Lying a fact of life in today's workplace. Jacksonville Business Journal. March 7, 1997.

Reason, J. (1990). Human error. Cambridge, MA: Cambridge University Press.

Regenwetter, M., Marley, A.A.J., Grofman, B. (2002). A general concept of majority rule. Mathematical Social Sciences, 43, 405-428.

Ren, Y., Carley, K. M., & Argote, L. (2006). The contingent effects of transactive memory: When is it more beneficial to know what others know? *Management Science, 52*, 671-682.

Riker, W. H., & Ordeshook, P. C. (1968). A theory of the calculus of voting. *The American Political Science Review, 62*, 25-42.

Simon, H. A. (1973). Applying information technology to organizational design. *Public Administration Review, 33,* 268-278.

Schneider, M. & Graham, J. (n.d.) *Effect of external memory tools on task performance in a future command and control environment: A virtual experiment using Construct*. Unpublished paper, Carnegie Mellon University.

Schwenk, C. (1984). Cognitive simplification processes in strategic decision-making. *Strategic Management Journal, 5,* 111-128.