Beyond Being in the Lab: Using Multi-Agent Modeling to Isolate Competing Hypotheses

First Author Name (Blank if Blind Review)
Affiliation (Blank if Blind Review)
Address (Blank if Blind Review)
e-mail address (Blank if Blind Review)
Optional phone number (Blank if Blind Review)

Second Author Name (Blank if Blind Review)
Affiliation (Blank if Blind Review)
Address (Blank if Blind Review)
e-mail address (Blank if Blind Review)
Optional phone number (Blank if Blind Review)

ABSTRACT

In studies of virtual teams, it is difficult to determine pure effects of geographic isolation and uneven communication technology. We developed a multi-agent computer model in NetLogo to complement laboratory-based organizational simulations [3]. In the lab, favoritism among collocated team members (collocators) appeared to increase their performance. However, in the computer simulation, when controlled for communication delay, in-group favoritism had a detrimental effect on the performance of collocators. This suggested that the advantage of collocators shown in the lab was due to synchronous communication, not favoritism. The canceling-out effects of in-group bias and communication delay explained why many studies did not see performance difference between collocated and remote team members. The multi-agent modeling in this case proved its value by both clarifying previous laboratory findings and guiding design of future experiments.

Author Keywords

Virtual team, computer supported cooperative work, multi-agent modeling, in-group favoritism, computer-mediated communication

ACM Classification Keywords

H.5.m [Information Interfaces and Presentation (e.g., HCI)]: Miscellaneous;
I.6. [Computing Methodologies]: Simulation and Modeling -- model development, simulation output analysis.

INTRODUCTION

Virtual teams involve both geographically distributed coworkers and technology-mediated communications [7]. As a result, for related fieldwork or laboratory experiments, it has been difficult to separate human factors from technology effects.

This paper describes our work to isolate competing hypotheses about virtual teams by using multi-agent computer simulations. In multi-agent models, "fundamental social structures and group behaviors emerge from the interaction of individuals operating in artificial environments under rules that place only bounded demands on each agent's information and computational capacity." (p4) [5] The "artificial" nature of agent-based modeling enabled us to eliminate inter-related factors and assess their pure effects. We found agent based modeling to be a surprisingly useful method to complement our fieldwork and laboratory experiments.

In our lab experiments about virtual teams, we found no significant differences in performance between collocated and remote players [3]. Remote players formed in-groups despite being deprived of face-to-face (FTF) social cues. Since both in-groups and computer-mediated communication (CMC) were present in our experiments, we could not conclude which variables led to the results. Such "mysteries" are faced by many other studies [6].

We developed the multi-agent model to focus on two competing hypotheses: in-group favoritism, and delay introduced by CMC. In-group favoritism is a major factor in the distribution of scarce resources in virtual teams [4]. On the other hand, time delay of CMC has substantial and pervasive effects on the performance of virtual teams [1]. Our work addresses three questions: (1) what are the independent effects of in-group favoritism and communication delay? (2) how do we account for the similar performance of remote and collocated team members, observed in previous studies? and (3) what lead to the formation of in-groups among people using CMC?

METHOD

Our overall program of research proceeded in four steps. First, we conducted a lab experiment to obtain behavioral patterns and benchmarks in a simulated virtual team setting. Second, we designed and tested a multi-agent model (in NetLogo) using observations from the lab as our benchmark. Third, we experimented in the controlled settings of the multi-agent model to examine pure effects of in-group favoritism and CMC delay. Finally, we explored the interesting model predictions by comparing them with more detailed findings in the lab data.

Organizational Simulation

We developed an online organizational simulation, Shape Factory as the apparatus of our lab experiment [3]. In Shape Factory, participants played the part of specialty producers of one of five different shapes. At the same time, each player had the task to collect different shapes from other players to fill "customer" orders. Players earned points by selling their specialty shapes and buying other shapes to fill orders. We imposed production limits to players so that there was a scarcity of shapes relative to customer demand. Players had varied communication channels. To be successful in Shape Factory, players had to use the available channels to contact, negotiate and cooperate with others.

We conducted 13 experimental sessions. Each session involved ten participants. Five "collocated players" were arranged in the same room, and five "remote players" were located in separate rooms, alone. Collocated players could communicate face-to-face or use text messages through the online game interface. Remote players could only interact with other players via text messages in the game. All players worked on laptop PC's with an Internet connection. Game scores measured performance of participants. Unequal trading volume among players was an indicator of in-group formation.

Shape Factory mimics four aspects of a virtual team environment: (1) a need to build and maintain relationships with others, (2) flexibility to choose collaboration opportunities, (3) limited and unequal resources, and (4) uneven communication channels. Among the four aspects, the limited and unequal resources produced a necessary and sufficient condition to generate in-group favoritism [4]. The uneven communication channels could create time delay for those using CMC [1].

Of all the findings, three were of particular interest to the model design. First, as the major quantitative finding of the lab experiment, both collocated and remote players formed in-groups of their own [3]. This result provided a benchmark for testing the model design.

Second, we found that collocated players tended to ignore remote players by prioritizing other collocators when selling and requesting shapes. In contrast, remote players did not purposefully attend or ignore any other players.

Third, we saw asymmetric communication delay in the experiment. Remote players experienced time delay in every interaction with other players. Collocated players could request and negotiate with each other verbally. They had the delay only when contacting others via text messages. The qualitative observations of favoritism and communication delay helped us design the behavioral rules for the multi-agent model.

Design and Test of the Multi-Agent Model

As the first step in model development, we started with a simple abstraction of the Shape Factory experiment. The model has 10 agents (simulated players): five "collocated" and five "remote". In addition to attending requests on a first-come-first-served basis, agent behaviors are driven only by in-group favoritism (for those who are collocated) or communication delay (for those who are remote). We do not implement cost, price, or negotiation strategies. Agents do not have cognition, emotions, or memory.

Our model is built in NetLogo [9]. NetLogo is a multi-agent programmable modeling environment. It permits the visual demonstration of agent behaviors. Figure 1 shows the interface of our model. Using the buttons in the upper left corner, we can set environmental conditions for the model. The charts in the lower left corner plot the major outputs of individual and cumulative results. The window on the right is the "lab." Ten agents are displayed in a circle. The five inside the shaded box are collocated agents, while the five outside are remote agents. The lines connecting pairs of agents are visualizations of the frequency of their interactions. Both the thickness of the lines and the numbers by the lines reflect the number of shapes exchanged between a pair of agents.

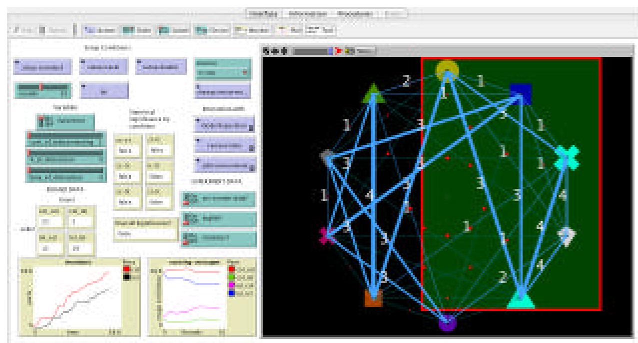

Figure 1: The model interface

We designed two behavioral rules to implement in-group bias and communication delay. The *favoritism rule* is defined as follows. When a collocated agent has a choice to request a shape from either another agent in the room or outside, it will choose the one in the room. Only when the other collocated agent has reached its production limit, will the agent place the request outside the room. Similarly, when presented with two requests to fill, a collocated agent will first fulfill the order of the agent in the room. After orders have been prioritized based on favoritism, they are treated on a first come first serve basis.

For remote agents, favoritism is absent from their behaviors because they are isolated, never in the same room as any other agents. For remote agents, both placing and fulfilling requests occur on a first come first serve basis.

The *communication delay rule* was operationalized by imposing a one-step time delay on communications of all remote agents. For collocated agents, they do not have delay when contacting each other. Only when collocated agents communicate with remote agents, will they have the one-step time delay.

To test whether the model design can generate meaningful behavioral patterns, we ran 50 rounds of simulations with both the favoritism and communication delay rules. As mentioned above, the formation of in-groups observed in the lab was the benchmark of the test. The model produced similar in-group patterns as the lab experiment. In the model output, collocated agents on average sold more to other collocators than to remote agents (24.4 vs. 5.6, p<.01), whereas remote agents on average sold more to other remote agents than to collocators (19.3 vs. 10.7, p<.01). Thus, our model is able to simulate a major collaboration pattern—formation of in-groups within both collocated and remote agents.

Experiment in the Multi-Agent Model

Our experiment in the agent model was a 2 ? 2 factorial design. Presence and absence of favoritism among collocators were the levels of one factor, while CMC delay and no delay were the levels of the second factor. The baseline for the experiment was the no-favoritism, no-delay treatment.

RESULTS

Results were based on 50 simulated sessions in each factorial treatment. In the model, performance was evaluated by the number of shapes an agent bought. Table 1 summarizes the performance outcome. As expected, collocated and remote agents performed equally well in the baseline treatment (30.2 vs. 29.8, p=.48). This confirmed that we did not introduce unintended variables into the model. Formation of in-groups was assessed by the unequal trading volumes among agents in different locations. Table 2 shows the formation of in-groups in two treatments, when there was only favoritism and when there was only CMC delay.

	Favoritism among Col			**No Favoritism**		
Delay	Col	Rem	p	Col	Rem	p
	35.2	24.8	<.01	37.3	22.7	<.01
No delay	Col	Rem	p	Col	Rem	p
	29.4	30.6	.038	30.2	29.8	.48

Col: collocated agents; Rem: remote agents
Numbers are the average amount of shapes Col or Rem bought per round.

Table 1. Performance outcome of the multi-agent model

Only Favoritism in Col			**Only CMC Delay in Rem**		
CC	CR	p	CC	CR	p
21.6	8.4	<.01	20	10	<.01
RC	RR	p	RC	RR	p
7.8	22.2	<.01	17.2	12.8	<.01

CC: collocators sold to collocators; CR: collocators sold to the remote.
RC: remote agents sold to collocators; RR: the remote sold to the remote.
Numbers are the average amount of shapes sold per round.

Table 2. Formation of in-groups of the multi-agent model

The effect of favoritism on performance is apparent in the "No delay" row of Table 1. With favoritism, collocated agents performed worse than remote agents (29.4 vs. 30.6, p<.05). Without favoritism, collocated agents performed as well as remote agents (30.2 vs. 29.8, p=.48).

The effect of CMC delay on performance was strong across all treatments. When there was favoritism among collocators, CMC delay reversed the performance outcome between collocated and remote agents. Collocated agents did better than remote agents when there was delay (35.2 vs. 24.8, p<.01), worse when there was no delay (29.4 vs. 30.6, p<.05). In treatments without the favoritism effect, collocated agents did better than the remote when there was CMC delay (37.3 vs. 22.7, p<.05). Without the delay, there was no difference between their performances (30.2 vs. 29.8, p=.48).

As for the formation of in-groups, results in the left side of Table 2 indicate that, with favoritism among collocators, collocated agents sold more to other collocators than to the remote (21.6 vs. 8.4, p<.01), and remote agents sold more to other remote agents than to collocators (22.2 vs. 7.8, p<.01). When there was only CMC delay, collocators still sold more to other collocators than to remote agents (20 vs. 10, p<.01), but remote agents sold significantly more to collocators as well (17.2 vs. 12.8, p<.01). Thus, CMC delay only led to in-groups within collocated agents, but not within remote agents.

Verification and Exploration of the Model Results

The negative effect of favoritism on performance was surprising. To verify our model predictions, we consulted the lab data. By examining the observation notes from the 13 experimental sessions, we identified three categorical usages of the FTF channel: no-use (4 sessions), business-use (4 sessions) and distractive-use (5 sessions). In the no-use sessions, CMC was the only channel used by both collocated and remote players. Collocated players still showed favoritism towards each other but remote players were not at a disadvantage from the CMC delay. The no-use sessions offered us a situation comparable to the no-delay, with-favoritism treatment in the model (southwest cells of Table 1). Consistent with our model results, collocated players performed worse than the remote players. Meanwhile, in the business-use sessions, collocated players utilized the FTF channel heavily to obtain shapes. These sessions provided a situation comparable to the with-delay, with-favoritism condition in the model (northwest cells of Table 1). Similar to our model prediction, collocated players did significantly better than remote ones.

The distractive usage of FTF was "serendipity" in the exploration. Collocators did significantly worse than remote players when noise (talking unrelated to Shape Factory tasks) dominated the FTF channel. We speculate that FTF distraction also played out in the similar performance observed in the lab. Although our current model does not have a comparable treatment to test this hypothesis, we will implement FTF distraction in the model as part of future work. The new model will create a more realistic virtual team environment. It will not only replicate the similar performance but also provide finer-tuned predictions of task performance in virtual teams.

DISCUSSION AND IMPLICATIONS

With our organizational simulations and the findings from the current stage of our agent-based model, we can now answer the three research questions.

1: What are the independent effects of in-group favoritism and communication delay?

In-group favoritism has a detrimental effect on the performance of collocated team members. CMC delay impairs the performance of remote members.

The harmful effect of favoritism on in-group members may be counter-intuitive. Two theories can help explain this result. Social network theory says that larger networks render more access and better coordination of diverse resources [10]. In-group bias encourages the formation of smaller networks. Consequently, it lowers performance of in-group members by confining their available resources. Alternatively, microeconomic theory tells us that constrained optima are often inferior to unconstrained ones [8]. In-group bias constrains cooperative opportunities. As a result, in-group members often settle with sub-optimal strategies when trying to access and integrate different functional areas.

Our findings and the related theories suggest that the effect of in-group favoritism is context dependent. In workplaces requiring collaboration, favorable bias among a subset of coworkers can reduce rather than increase their allocation of scarce resources (e.g., individual expertise).

2: How do we account for the similar performance of remote and collocated team members observed in previous studies?

The answer to the previous question implies that the canceling-out effects between in-group favoritism and CMC delay could lead to the indistinguishable performance of remote and collocated team members. Observations of previous studies [6] often led to the conclusion that distance does not affect virtual teamwork. This paper shows that specific aspects of distance (i.e., in-group favoritism and CMC delay) have strong impacts on performance. However, the impacts of distance can often confound and cancel out each other.

3: What lead to the formation of in-groups among people using CMC?

Related studies proposed that shared communication delay was a base for remote team members to form in-groups [2]. Our results indicate that it is the exclusion from collocators, not the time delay, that drive remote people into subgroups.

Empirical Implications

To virtual team managers, this paper provides three guidelines. First of all, they should respect the power of distance. It still affects behaviors and performance of virtual teams. Secondly, they should prevent favoritism among subsets of virtual team members. In-group favoritism not only impairs the cohesion of overall teams, but also negatively affects the performance of the in-group members. Finally, to avoid in-groups among remote members, managers need to deploy organizational strategies rather than technological initiatives.

ACKNOWLEDGMENTS

This research is supported by xxx. We acknowledge the assistance and insights from a number of people not on the author list, including xxx, xxx, xxx, and xxx.

REFERENCES

1. Arrow, H., Berdahl, J.L., Bouas, K.S., Craig, K.M., Cummings, A., Lebie, L., McGrath, J.E., O'Connor, K.M., Rhoades, J.A. and Ann Time, S. Technology and groups: An integration. *Computer Supported Cooperative Work*, *4*, 1996, 253-261.
2. BeLanger, F. Communication patterns in distributed work groups: A network analysis. *IEEE Transactions on Professional Communication*, *42*, 4 (1999), 261-275.
3. Bos, N., Sadat Shami, N., Olson, J.S., Cheshin, A. and Nan, N. In-group/out-group effects in distributed teams: An experimental simulation. In *Proc. CSCW 2004*, ACM Press (2004), 429-436.
4. Dobbs, M., and Crano, W. D. Outgroup accountability in the minimal group paradigm: Implications for aversive discrimination and social identity theory. *Personality and Social Psychology Bulletin*, 27, 3 (2001), 355-364.
5. Epstein, J.M. and Axtell, R. *Growing Artificial Societies: Social Science from the Bottom up*. Brookings Institution Press, Washington, D. C., 1996.
6. Finn, K.E., Sellen, A.J., and Wilbur, S.B. *Video-Mediated Communication*. Lawrence Erlbaum Associates, Mahway, NJ, 1997.
7. Hinds, P.J. and Bailey, D.E. Out of sight, out of sync: Understanding conflict in distributed teams. *Organization Science*, *14*, 6 (2003), 615-632.
8. Mas-Colell, A, Whinston, M.D. and Green, J.R. *Microeconomic Theory*, Oxford University Press, Oxford, New York, 2002.
9. NetLogo website. http://ccl.northwestern.edu/netlogo/
10. Rosenthal, E. Social networks and team performance. *Team Performance Management*, *3*, 4 (1997). 288.