


Information Flow and Decision Making
in Teams under Threat

Garold Stasser and Steve Hinkle

Department of Psychology
Miami University
Oxford, OH 45056

Final Report
Report Number 96-1
Office of Naval Research
Grant N00014-90-J-1790



19960503 084

REPORT DOCUMENTATION PAGE			Form Approved OMB No. 0704-0188	
Public report documentation page of information developed by the Department of Defense, Department of Energy, Department of Health, Education and Welfare, Department of Housing and Urban Development, Department of Justice, Department of Labor, Department of State, Department of Transportation, Department of the Interior, Department of Agriculture, and Department of Veterans Affairs. This report is the property of the Government and is loaned to your agency; it and its contents are not to be distributed outside your agency without the approval of the originating agency. This report is to be maintained in your agency's files and is to be made available to the public upon request. This report is to be maintained in your agency's files and is to be made available to the public upon request. This report is to be maintained in your agency's files and is to be made available to the public upon request.				
1. AGENCY USE ONLY (Leave blank)		2. REPORT DATE April, 1996	3. REPORT TYPE AND DATES COVERED Final Report: 15 March 1990-31 Dec 1994	
4. TITLE AND SUBTITLE Information Flow and Decision Making in Teams under Threat			5. FUNDING NUMBERS G N00014-90-J-1790	
6. AUTHOR(S) Garold Stasser, Steve Hinkle				
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Department of Psychology Miami University Oxford, OH 45056			8. PERFORMING ORGANIZATION REPORT NUMBER 96-1	
9. SPONSORING MONITORING AGENCY NAME(S) AND ADDRESS(ES) Cognitive and Neural Science Division Office of Naval Research 800 North Quincy Street Arlington, Virginia 22217-5000			10. SPONSORING MONITORING AGENCY REPORT NUMBER	
11. SUPPLEMENTARY NOTES				
12a. DISTRIBUTION AVAILABILITY STATEMENT Approved for public release; distribution unlimited			12b. DISTRIBUTION CODE	
13. ABSTRACT (Maximum 200 words) Information flow during team decision-making was examined using a Situation Assessment Simulation (SAS). SAS presents multiple trials of a task that requires seven-person teams to access and communicate 20 items of diagnostic information and to decide which of two courses of action is appropriate under severe time constraints. Correct decisions are rewarded and incorrect decisions penalized. Four studies tracked communication patterns in hierarchically-organized teams and examined emergent patterns across repeated episodes under varying levels of task complexity, information distribution knowledge, communication feedback, and decision centrality. Of primary interest was the development of efficient communication strategies when access to information was distributed unevenly across team members. Results informed the development of a theory of tacit coordination: members actions are guided by expectations of other actions and the identification of pivotal actions given what others' are expected to do and what is required by the collective task. Expectations regarding others' actions can be based on knowledge of others' skills, interests, role prescriptions, resource access and past behavior. This report also contrasts information flow in hierarchically-organized teams with collective information sampling in face-to-face discussions of decision-making committees.				
14. SUBJECT TERMS Team decision-making; team communication; group processes;			15. NUMBER OF PAGES 74	
			16. PRICE CODE	
17. SECURITY CLASSIFICATION OF REPORT Unclassified	18. SECURITY CLASSIFICATION OF THIS PAGE Unclassified	19. SECURITY CLASSIFICATION OF ABSTRACT Unclassified	20. LIMITATION OF ABSTRACT SAR	

Information Flow During Team Decision Making

Decision-making teams are often located in environments where information is unevenly distributed among members. Whereas some information may be available to most or all members, only one or two members may have access to other information. Moreover, information load is often high, making an exhaustive review of information infeasible given pressures to reach a timely decision. In such contexts, effective decision making often depends on effective communication and integration of information.

Background

Earlier work has examined the flow of information during the face-to-face discussions of decision-making teams (e.g., Stasser & Titus, 1985, 1987; Stasser, Taylor, & Hanna, 1989). This work demonstrated that the contents of teams' discussions is not ordinarily a representative sample of the available information. On the one hand, teams are more likely to discuss information that all members have before discussion (shared information) than to discuss information that is held by single members (unshared information). On the other hand, teams tend to discuss information that supports an the most popular decision alternative in the group.

Several implications of these earlier findings are important. First, data that all group members share at the onset of discussion has more impact on their final decision than data that is initially available to only one or a few (Gigone & Hastie, 1993). A second and related point is that communication in the group tends to restrict rather than amplify the range of data actively considered in making a decision. Third, information exchanged during discussion not only guides the emergence of a consensus but also much of discussion is devoted to bolstering a decision that is already functionally made (Stasser & Davis, 1981). Finally, these findings underscore the importance of considering how information is distributed among group members prior to their interaction. This consideration is particularly important in decision environments where individuals have differential access to various types of information (distributed information systems).

Communication in Hierarchically Organized Teams

The current work differed from earlier work in several ways. First, it examined information flow in hierarchically organized teams rather than in face-to-face discussion. Second, it examined the emergence of information processing strategies as a teams gained experience working together. Earlier studies typically used ad hoc teams that made one decision. Thus, there was little opportunity for teams in these studies to develop strategies. Third, earlier work often used judgmental tasks rather than tasks that have right and wrong answers (Laughlin &

Ellis, 1986).

A Situation Assessment Simulation (SAS) task was developed to be implemented on a local area computer network. SAS simulates a situation that requires a team to choose between two courses of action based on an array of 20 diagnostic items of information. A detailed description of the task is contained in Stasser, Hinkle, Fox-Cardamone, and Ely (1992, see Appendix A). SAS accommodates teams of seven persons and simulates as many as 32 episodes of the situation assessment tasks. Among other things, the experimenter can control who among team members can access each of the 20 items of information (distribution of information access), who can communicate with whom (communication organization), the time allotted for a decision during each episode, and the level of payoffs for correct and penalties for incorrect decisions. SAS continuously monitors and records members actions including what information is accessed by each member (information assessment), sent from a member to another (communication), and requested by a member from another (queries). Four studies were completed using the SAS simulation software.

Study 1: Communication and Decision Sensitivity under Threat

Stasser, Hinkle, Fox-Cardamone, and Ely (1992) examined team communication patterns and decision-making effectiveness under two organizational schemes. All teams were hierarchically organized in three vertical levels. Four members at the bottom of the organization were responsible for monitoring and communicating information to two members at the second level. The two middle members were responsible for selecting information to relay to the decision-maker at the top. In one variant of this organizational structure, members could only communicate with others immediately above or below them (vertical communication). In the other variant, communication within a level was also permitted (lateral communication). Task complexity was also manipulated. In a simple version, all 20 items of information were equally diagnostic. In a more complex version, nine of the 20 items were irrelevant, and teams not only had to gain proficiency in communicating information but also try to learn with experience what information was important.

The objectives of this study were to: a) examine the interactive effects of communication structure and task complexity on performance and adaptability (improvement in performance); b) examine the effects of stress (time limits and risk of gain and loss) on communication patterns, amount of information transferred, and performance; and c) demonstrate measures of team decision, communication, and member vote sensitivity. A fourth objective was to obtain verbal protocols describing emergent member and team communication strategies for dealing with time pressures. This fourth objective required

content coding of open-ended messages sent between trial blocks and open-ended responses to questions posed at the end of the study. These results have not been analyzed, but the messages have been coded and are ready to be analyzed.

Forty-four, seven-person teams completed the study; 11 teams in each of the four conditions defined by communication structure and task complexity. A complete report of results is contained in Appendix A. Briefly, in all conditions, team communication performance improved over time as evidenced by increases in the amount of important information reaching the centralized decision-maker. However, improvements in communication performance resulted in improved team decision performance only for the simple version of the task. For the complex task, there was evidence that team members learned to distinguish important from unimportant information only in the simpler, vertical communication structure. Nonetheless, the modest gains in member sensitivity to information importance in this condition did not result in improved team performance. Additionally, the results suggested that teams were quite tolerant of, and susceptible to, communication redundancy.

Study 2: Coordination of Communication

Stasser, Hinkle, Ely and Fox-Cardamone (1995) examined the coordination of communication over repeated episodes of the situation assessment task when some items of diagnostic information were available to only single members (unshared information) whereas other items were available to two members (partially shared information). Knowledge of information distribution among team members and availability of feedback about what information was reaching the Commander were manipulated. This study extends the aforementioned program of research on information sampling during face-to-face discussion in decision-making groups. This earlier line of research suggests that groups often fail to discuss information that is uniquely-held by one member. However, the kind of "committee" decision making simulated in this earlier work (e.g., Stasser & Titus, 1985, 1987; Stasser et al, 1989; and Stasser & Stewart, 1992) is different in several important ways from the team decision-making examined in this project. First, the domain of relevant information is more clearly delineated in the SAS team-decision making task and assigned roles in the team make potentially more salient the types of information that various members may uniquely access. Second, the SAS task provides for repeated episodes permitting the development of communication strategies with experience. Third, because of their distributed nature, teams in the SAS task had less feedback about what others were communicating than is the case for face-to-face interaction.

To examine how these differences may impact the

communication of unshared and partially-shared information, access to information was structured so that four items of shared information were available to all members of a seven-person team; eight items of partially-shared information were available to only two members; and eight items of unshared information were available to only one member. The partially-shared and unshared items were distributed among the four members at the bottom of the hierarchically-organized team. For this information to be considered by the "Commander" at the top of the hierarchy, it had to be communicated through the communication network before a final decision was required.

Three major predictions were examined. First, when knowledge of how information is distributed and communication feedback is unavailable, teams would evidence the bias observed in earlier work on committees -- namely, partially shared information would be communicated more frequently and reach the Commander more often than unshared information. This prediction follows from an extension of a collective information sampling model first proposed by Stasser & Titus (1987). Second, knowledge of how access to information was distributed among team members would reverse the bias observed in face-to-face committees so that unshared information would be more likely to be communicated than partially-shared information. This prediction is based partly on an extension of the reasoning expressed in Wegner's (1986) theory of transactive memory, but, less formally, follows if one assumes that team members accept responsibility for tasks which only they can accomplish. Third, with communication feedback, groups will adjust their communication strategies over time to correct or, at least, diminish these predicted biases.

The basic design of the experiment was 2 X 2 X 4 factorial of knowledge of information distribution (limited versus elaborated), inter-trial communication feedback (available versus not available) and trial blocks. Fifty-six seven-person teams participated in this study for two hours each. All teams completed 32 trials divided into four blocks of eight trials.

A more detailed report of results is contained in Appendix B. Primarily reports of results were also presented by Stasser and Wittenbaum (1995) and Stasser (1995a). The results support the predictions regarding the effects of information distribution knowledge on communication of partially-shared and unshared information. When the information distribution was not known, partially-shared information was more likely than unshared information to reach the Commander. This pattern became more pronounced over trial blocks, particularly when there was also no inter-trial feedback to team members regarding the items reaching the Commander. When the information distribution was known, the pattern was reversed in that unshared items were more likely to

reach the Commander than were partially-shared items. However, the availability of feedback virtually eliminated this bias in favor of unshared information by the fourth trial block. Thus, feedback was efficacious when members also knew who could access what information. Without this knowledge, feedback had little systematic effect.

Team performance as measured by decision sensitivity to the expected values of decisions options (namely, covariation of decisions with the expected value of decision options, a measure akin to d' in signal detection theory) was better when information distribution was known; however, only the combination of known information distribution and feedback improved performance across trial blocks.

Somewhat contrary to predictions, communication feedback by itself did not reduce the bias favoring partially-shared over unshared items, nor did it improve performance. Team members were seemingly unable to implement a strategy that allowed them to address the disproportionate omission of unshared items. However, there is indirect evidence that they detected a problem. In the feedback without distribution knowledge condition, requests from one team member to another for specific items of information increased by threefold over the trial blocks and the number of these requests was 60% higher in trial blocks 3 and 4 than in the other conditions.

Several applied implications follow. First, when information is not uniformly available to all members, increasing the number of members who have access to particular items will increase the likelihood that they will be successfully transmitted through a communication network. Thus, consistent with intuition, information that is thought to be relatively more important should be immediately accessible to more team members. Second, however, if team members are aware of who, and how many, can access particular types of information, this recommendation may change. Under these conditions, each member is more likely to communicate information that is uniquely available to him or her. The implication is that the more people who are thought to have access to an item, the less responsibility any one of them will accept for accessing and communicating the item. Finally, whereas feedback (debriefing) regarding the information actually reaching the decision-maker is not sufficient by itself to ensure that teams can address the communication biases introduced by the patterns of information access across team members, feedback coupled with awareness of who has access to what information does permit teams to correct communication biases with experience.

Studies 3: Recommendations as Proxies for Information

Preliminary results of a third study were presented by

Stasser (1995b). This study examined the effects of a "voting" option on the communication of information in the SAS task. The interest in the voting option stems partly from another obvious difference between the earlier work on decision-making in face-to-face groups and the present work on decision-making in hierarchically-organized teams. Discussion of information in committees is typically accompanied by statements of preference for decision alternatives. However, the version of the SAS task used in the first two studies did not permit team members to recommend a course of action. We expected that the opportunity to make recommendations would suppress communication of information for three reasons. First, members may view sending recommendations as an efficient way of communicating the gist of their information without having to send every item, and this would be a particularly attractive strategy under severe time pressures. Second, given the presumed efficiency of sending recommendations rather than "raw" data, members would spend more time accessing and integrating information and, thus, have less time to send the raw data to supplement their recommendations. Third, having made a recommendation, members may be more selective in sending information and tend to send information that is consistent, and to suppress information that is inconsistent, with their recommended course of action.

We also examined the feedback manipulation used in the previous study although all groups were apprised of how information access was distributed across members. The basic design of the experiment was 2 X 2 X 4 factorial of voting (allowed versus disallowed), inter-trial communication feedback (available versus not available) and trial blocks. Sixty-two seven-person teams participated for two hours each in this study.

Preliminary analyses revealed that being able to send recommendations reduced the amount of unshared information communicated by 44% and the amount of partially-shared information by 28%. Both are highly significant reductions and the difference in reduction for unshared and partially-shared information is also significant. Thus, as predicted, voting reduced the amount of information communicated and the impact was greater for partially-shared than for unshared information. Replicating one finding of the previous study, unshared information was more likely to reach the Commander when the information distribution was known and voting was not allowed. Permitting voting exaggerated this pattern resulting in the Commander receiving only about half as much partially-shared as unshared information. In summary, allowing voting reduced the amount of information sent to the Commander, and this reduction was greater for partially-shared than for unshared information.

Whereas members communicated more information when voting was allowed than when it was not, the opposite pattern was

obtained for assessing (measuring) information. Team members accessed more of the information available to them when they were allowed to vote than when they were not. Thus, they knew more and communicated less when they were allowed to vote. The analyses examining the bias hypotheses (namely, voting would promote selective sending of information in order to support a recommendation) has not been completed.

Study 4: Voting and Knowledge of Information Distribution

A partial replication of both studies 2 and 3 crossed the distribution knowledge with the voting manipulation. Of particular interest was the combined effects of not knowing how information was distributed and having the option to send recommendations. Based on results of the two previous studies, we expected that the communication of unshared information would be quite low under these conditions. Additionally, in this study, the importance of information was changed during the last trial block. In the first three trial blocks, all information was equally predictive of the correct decision (as was the case for all blocks in studies 3 & 4). However, in the last trial block, the correct decision was perfectly predicted by the unshared information and all other information was superfluous. We expected that conditions that promoted the use and communication of unshared information would facilitate the detection of the change and improve performance in the last trial block. Conversely, we expected particularly poor performance in the last trial block when members did not know how information was distributed and were able to send recommendations.

The basic design of the experiment was 2 X 2 X 4 factorial of voting (allowed versus disallowed), knowledge of information distribution (limited versus elaborated), and trial blocks. Forty-eight seven-person teams participated for two hours each.

The communication effects of voting and information distribution knowledge that were observed in the previous two studies were clearly replicated in this study. The combination of voting and not knowing the information distribution suppressed dramatically the communication of unshared information; in this condition, only 12% of unshared items reached the Commander during a trial, on average. In contrast, when both the distribution of information was known and voting was disallowed, an average of 43% of unshared items reached the Commander per trial. There were also strong block effects on communication resulting in more unshared items being communicated during the last trial block. Nonetheless, the aforementioned effects of voting and limited knowledge of information distribution on the communication of unshared information were stronger in later than in earlier blocks.

The change in information importance implemented in the last trial block resulted in an overall decrease in decision sensitivity from the third to fourth trial block. The only exception to the pattern occurred when voting was disallowed and information distribution was known. Indeed, somewhat unexpectedly, allowing voting even when the distribution of information was known interfered rather substantially with teams' abilities to detect that unshared information was strongly (actually, perfectly) predictive of the correct decision over the last eight trials.

These results are relevant to contexts where teams need to be sensitive to emerging patterns in widely distributed information. In these contexts, communication strategies that fail to communicate sufficient amounts of unshared information to a central decision-maker risk overlooking highly diagnostic patterns of unshared information. Only the central decision-maker can "get the big picture"; individuals less central in the communication network necessarily have limited access to unshared information and, thus, are not in a position to recognize the diagnostic patterns in widely-distributed information. This study suggests that both the practice of sending recommendations and ignorance of how access to information is distributed will interfere with the detection of diagnostic patterns in unshared information.

Face-to-face Communication Revisited

The foregoing studies with distributed decision-making teams provided the basis for extending work with face-to-face interactions. The finding that knowledge of information distribution in a team reverses the often demonstrated finding that unshared information is underrepresented in a group's discussion raises the possibility that team members may use a variety of cues to anticipate what others can and will do and adjust their own behavior accordingly. Wittenbaum, Stasser and Merry (in press) articulated a model of tacit coordination in which members form expectations about other actions, assess the demands of the collective task, and then allocate their resources (time, effort, attention, etc.) to facilitate the successful completion of the task. Stasser and Stewart (1992) examined communication of shared and unshared clues when a mystery task was framed as either judgmental (identify the suspect who most likely committed the crime) or intellective (identify the guilty suspect; i.e., a demonstrably correct answer exists, Laughlin & Ellis, 1986). The prospect of being right (or the risk of being wrong) when the task was framed as intellective promoted the communication and use of unshared clues. Stasser and Vaughan (in press) reviewed, critiqued and integrated formal models of communication during in face-to-face interaction. Each of these papers acknowledged support of the grant.

References

- Gigone, D., and Hastie, R. (1993). The common knowledge effect: Information sampling and group judgment. Journal of Personality and Social Psychology, 65, 959-974.
- Laughlin, P. R., & Ellis, A. L. (1986). Demonstrability and social combination processes on mathematical intellectual tasks. Journal of Experimental Social Psychology, 22, 177-189.
- Stasser, G. (Sept, 1995b). Information communication in distributed decision-making teams. Invited presentation at the Small Groups Pre-Conference, Society of Experimental Social Psychology, Washington, D.C.
- Stasser, G. (May, 1995a). Information exchange strategies in hierarchically organized groups. Invited symposium presentation at the annual meeting of the Society for Industrial and Organizational Psychology, Orlando, FL.
- Stasser, G., & Davis, J. H. (1981). Group decision making and social influence: A social interaction sequence model. Psychological Review, 88, 523-551.
- Stasser, G., Hinkle, S., Fox-Cardamone, L., & Ely, P. (1992). Communication and decision sensitivity in decision-making teams under threat. Paper presented at the annual meeting of the Society for Industrial and Organizational Psychology, Montreal, Canada.
- Stasser, G., & Stewart, D. (1992). Discovery of hidden profiles by decision-making groups: Solving a problem versus making a judgment. Journal of Personality and Social Psychology, 63, 426-434.
- Stasser, G., Taylor, L. A., & Hanna, C. (1989). Information sampling in structured and unstructured discussions of three- and six-person groups. Journal of Personality and Social Psychology, 57, 67-78.
- Stasser, G., & Titus, W. (1987). Effects of information load and percentage of shared information on the dissemination of unshared information during group discussion. Journal of Personality and Social Psychology, 53, 81-93.
- Stasser, G., & Titus, W. (1985). Pooling of unshared information in group decision making: Biased information sampling during discussion. Journal of Personality and Social Psychology, 48, 1467-1478.

Stasser, G., & Wittenbaum, G. M. (May, 1995). Tacit coordination: Using task, social and self knowledge. Invited symposium presentation at the 67th Annual Meeting of the Midwestern Psychological Association, Chicago, IL.

Stasser, G., & Vaughan, S. I. (in press). Models of participation during face-to-face unstructured discussions. To appear in J. H. Davis & E. Witte (Eds.), Understanding group behavior: Consensual action by small groups (Vol. I). Hillsdale, N. J.: Erlbaum.

Wittenbaum, G. M., Stasser, G., & Merry, C. J. (in press). Tacit coordination in anticipation of small group task completion. Journal of Experimental Social Psychology.

Appendix A

Communication and Decision Sensitivity in Decision-Making Teams under Threat

Communication and Decision Sensitivity in
Decision-Making Teams under Threat

Garold Stasser, Steve Hinkle, Lee Fox-Cardamone,
and Pamela Ely

Miami University

Team Decision-making Under Threat

Abstract

This research examines team communication processes and decision-making when time is limited and losses are risked (threat). Hierarchically-organized teams monitored incoming information and communicated items to a centralized decision-making member who had to decide quickly between two courses of action. Over a series of 32 trials, 44 teams were rewarded (monetary gain) for correct decisions and penalized (monetary loss) for incorrect ones. Two versions of the task were used: either all information was weighted equally (equal-weight task) or some information was irrelevant (unequal-weight task). Also, some teams could send and receive messages only to those immediately above or below them in the hierarchy (vertical communication structure) whereas others could also use lateral links (lateral communication structure).

Team communication performance improved over time as evidenced by increases in the amount of important information reaching the centralized decision-maker and the sensitivity of this communicated information to changes in the available information. However, for the unequal-weight task, this improved communication did not result in a corresponding increase in team decision performance. Indeed, there was little evidence that team members learned to distinguish important from unimportant information. Moreover, members did not frequently use lateral communication links when they were available. These results are discussed in terms of the task demands when information load is high, information is distributed, performance feedback is uncertain, and time is limited. Implications for team learning in distributed information environments are considered.

Many decision environments require that teams react quickly in the face of impending danger. Both action and no action often risk costly outcomes, and the team does not have the luxury of carefully analyzing the situation before deciding what to do. A prototypical example is a military unit on a routine maneuver (e.g., a naval ship) under potential attack (e.g., by an approaching fighter plane) when the intention of the threatening entity is unclear. Such situations characteristically present a dilemma. On the one hand, launching a preemptive strike is desirable if an attack is imminent; however, an unnecessary military engagement may be costly, compromising diplomatic relations and foreign policy, destroying an unstable but peaceful coexistence, and precipitating other hostile action. On the other hand, taking no action is a regrettable decision in the event of an attack. There may be much information that allows one to judge the likelihood of attack, but the information is typically complex, uncertain and insufficient to make a definitive judgment. Moreover, there is too little time to consider all of the available information.

Risk of loss and time pressure are fundamental elements of threat (Gladstein & Reilly, 1985; Staw, Sandelands, & Dutton, 1981). The aforementioned military situation is one of many examples of teams having to choose a course of action under threat. Others include a disaster team monitoring unstable weather conditions and deciding whether to mobilize, a company's management considering costly plans to offset devastating changes that are threatening to occur in the market place, or a medical team deciding whether a radical and risky operation is necessary to stabilize a patient's condition during emergency treatment. Although the time scale is different in these examples, the critical element is that there is not enough time to obtain and analyze the information necessary to make a definitive decision. Additionally, the available courses of action risk negative consequences.

It is often the case that decisions under threat are made by a relatively small team. Such teams need not have any detectable organization; members may be able to communicate freely and share equally the responsibility of collecting, interpreting and integrating information to yield a decision. But often there is a formal or informal hierarchical organization. A hierarchical structure permits and often promotes some division of labor. The responsibility for the final decision often falls on one or two members at the top and the responsibility for monitoring and communicating incoming information is concentrated at the bottom of the organization. Team members in the middle may have diverse functions including collating information received from below, deciding what to relay to the decision makers at the top, and, perhaps, offering recommendations and advice to the top-level members.

Hierarchically-organized teams making decisions under threat have been rarely studied in the experimental literature. They fall at the interface between three research traditions in the behavioral sciences: small group performance, organizational structure and effectiveness, and individual performance under stress. Staw, Sandelands, and Dutton (1981) provided a rare attempt to integrate findings from these three traditions; they identified one common theme suggesting that individuals, small groups, and organizations "show restriction of information processing and constriction of control under threat" leading to rigid responding. At a minimum, their conclusions alert us to the possibility that dysfunctional reactions to

stress that are typical of individuals under psychological stress (e.g., narrowing of attention, incomplete consideration of relevant information, and dogged repetitions of maladaptive responses; Eysenck, 1976; Janis & Mann, 1977; Lazarus, 1966; Zajonc, 1965) may be compounded when they are functioning in a social unit. However, as Staw et al. (1981) noted, their threat-rigidity hypothesis needs more careful examination. For one thing, it is not clear that all types of groups and organizations are susceptible to restriction of information flow and concentration of control in the hands of a few under threat. For another, it is possible that these reactions to threat are functional in some contexts. For example, if a task is well understood and unchanging and a team is highly coordinated in their effort, their response to stress may be effective and efficient.

The approach in the present study is most closely allied with the small group performance literature (Davis, 1969; Steiner, 1972; McGrath, 1984) in that we examined communication and decision-making effectiveness within the context of the interacting group. We were particularly interested in information flow among team members en route to a decision (Stasser, 1988), the communication strategies that drive information exchange, and the sensitivity of the team to patterns of available information.

Unfortunately, much of the small group performance and decision-making literature has used groups that meet face-to-face in minimally structured discussions. One exception is the study of imposed communication networks in problem-solving groups (Shaw, 1964) which offers some insight into how small groups perform when communication is highly constrained. Additionally, we borrowed some concepts from the organizational literature (Duncan & Weiss, 1979; Katz & Kahn, 1978; O'Reilly, 1983) to understand how information exchange and decision-making processes might operate in hierarchically structured teams.

A shortcoming of traditional small group research for our purposes is that the decision tasks are typically one-shot and static and there is rarely threat (in the sense of limited time coupled with high risk). There is also a tendency in this tradition to focus either on the group product (i.e., the final decision or solution) or some limited aspect of interaction. We used a computer-simulated game as a research tool. Simulation offers several advantages including the ability to emulate a fast-paced and changing environment while maintaining experimental control over many task parameters (Leik & Gifford, 1986). Just as importantly, computer simulation permits one to monitor and record many measures of performance and process as they unfold over time.

Information Sampling in Decision-Making Groups

Recent theoretical statements have promoted the idea that the information considered during discussion shapes the final decision of a group (Burnstein & Vinokur, 1977; Kaplan & Miller, 1983; Stasser, 1988). Hoffman (1979) summarized a series of studies demonstrating that the support for a group's decision accumulates throughout discussion and that the final decision can usually be forecast well before it is announced by simply tabulating the number of statements made in support of each decision alternative. On the face of it, this work paints a reassuring picture of decision making processes; in spite of their inefficiencies and vagaries, groups are apparently guided by the available information. Moreover, conventional

wisdom (Tillman, 1960) and managerial practice suggest that groups can make more informed decisions because they have the opportunity to pool the diverse knowledge of their members. The implicit assumption seems to be that group discussion is composed of a representative sampling of the collective pool of information that members bring to discussion. Therefore, if particular members have critical facts that others do not have, there is the potential for an emergent understanding that could not occur if individuals acted alone.

However, several studies that have controlled what members know before discussion and either recorded the contents of discussion (Stasser, Taylor, & Hanna, 1989; Stasser & Stewart, in press) or the recall of information after discussion (Stasser & Titus, 1985, 1987) have shown that collective information sampling is incomplete and biased. The bias arises from several sources. First, groups are more likely to consider information that all members knew beforehand than information that only one or two knew. A second source of bias stems from the tendency for groups to communicate information that supports their initial predispositions even when it is clear that their initial opinions are based on partial information.

A third source of bias actually compounds the effects of the first two. Stasser et al. (1989) found that groups tend to repeat information throughout the course of discussion; thus, information introduced early is likely to be brought up again later. What is most interesting is the type of information that was repeated. One might suppose that novel information, once mentioned, would be likely to be reconsidered. To the contrary, information that was shared before discussion was not only more likely than unshared information to be mentioned, it was also more likely to be repeated once it was mentioned. Thus, it seems that groups are insensitive to redundancy in their communications. In fact, repetition may serve to focus their attention on a manageable subset of the available information and, thereby, help them form a coherent and shared portrayal of the decision options that would lead to a clear-cut decision.

Several implications of these findings are important. First, data that all group members share at the onset of discussion has more impact on their final decision than data that is initially available to only one or a few. A second and related point is that communication in the group tends to restrict rather than amplify the range of data actively considered in making a decision. Third, information exchanged during discussion not only guides the emergence of a consensus but also much of discussion is devoted to bolstering a decision that is already functionally made (Hoffman, 1979; Stasser & Davis, 1981). Finally, these findings underscore the importance of considering how information is distributed among group members prior to their interaction. This consideration is particularly important in decision environments where individuals have differential access to various types of information (distributed information systems).

This earlier work on information flow differs from the present focus on hierarchically structured teams in two ways. First, the earlier work has predominately used groups that met face-to-face with no imposed communication structure. Often the contexts that give rise to distributed information systems also result in members being physically separated and connected via relatively inflexible communication links. Second, and perhaps more

importantly, group members in this earlier work were aware that they may not have identical sets of information but had no way of discerning at the onset of their interaction what types of information were unshared nor who among them may possess critical unshared information. Organized teams often have some knowledge of each others' domains of specialized knowledge; thus, they may be able to enact more effective and efficient information exchange strategies than when members' unique areas of expertise are obscured (Stasser, 1992, in press).

Communication Structures in Hierarchical Decision Making Teams

The traditional organizational literature (Bass, 1983; Katz & Kahn, 1978) distinguishes between vertical and horizontal structures. Vertical structures typically involve specialized roles depending on a member's location in the hierarchy. For example, decision-making is often concentrated at the top of the organization and enactment of decisions delegated to members at the bottom. Information flow is primarily vertical. Horizontal structures permit less specialization with decision-making, information assessment, and production function distributed throughout the organization. This study focuses on a decision-making team that is a specialized organizational unit whose primary function is to monitor information and make discrete, binary (go/no go) decisions. For our purposes, enactment of the decision can be considered the responsibility of other units within a larger organization. The primary concerns of such a decision-making team are to maintain accurate and comprehensive accounts of relevant information and to integrate this information to yield a timely and informed decision.


Whereas this structure is similar to the traditional notion of vertical structure in that the decision making function is concentrated at the top, it differs in that the primary flow of information is upward. That is, the lower levels have responsibility for monitoring and conveying information whereas the upper levels must integrate the information to make a go/no-go decision.

The present study employs a seven-person team organized in three vertical levels. The four lowest members are responsible for monitoring information and selecting items to convey to the two middle team members. The middle members are primarily responsible for selecting the most critical information to relay to the top member. The top member must make the final decision based on what is conveyed from the lower levels. In the simplest structure, we allowed only a vertical flow of information as denoted by the arrows in the top panel of Figure 1; we refer to this structure as a vertical communication structure.

As depicted in the lower panel of Figure 1, we also used an extension of this basic hierarchical structure which permitted communication between members at a given level. For simplicity, we refer to this as a lateral structure although it is still predominantly hierarchical in nature. The lateral links permit cross-flow of information but seemingly add little advantage in getting information to the top of the system. However, they may add a distinct advantage when the importance of information is unclear and members must learn through experience to distinguish important from unimportant information. Quite simply, the lateral links permit more dissemination of information throughout the network which, in turn, gives individual members a broader basis for determining the relative importance of

Figure 1

Schematic of Vertical and Lateral Communication Structures



information. Moreover, cross-flows may build redundancy into the system; information that may be overlooked when it arrives from one channel has a second chance to be considered. However, such redundancy of communication opportunity may create actual redundancy in communication. If members do not carefully monitor and coordinate their communications, some information may be repeatedly communicated, particularly to members at the top and middle of the hierarchy. Communicating some information more than once will likely exclude the communication of other important information when time is limited.

In the terminology of the communication network literature (Shaw, 1964, 1978), lateral links should increase information dispersion (the number of nodes in the network receiving critical items), channel saturation (or load: the number of communication links passing through a node), and information saturation (or load: the amount of information passing through a node). An often-cited conclusion from this literature is that decentralization of communication enhances performance on complex tasks, but hinders performance on simple tasks (McGrath, 1984; Forsyth, 1983). Although task complexity is ill-defined, it typically revolves around the degree to which effective communication depends on either the integration of disparate information or feedback regarding the relative importance of information.

From the perspective of Duncan and Weiss's (1979) analysis of organizational learning, our vertical communication structure is a simple functional organization whereas the addition of lateral links creates a mixed functional organization. Their mixed functional organization is a hybrid of the simple functional and a decentralized organization. They concluded that mixed functional, as compared to simple functional, organizations are more adaptable, enhancing their ability to respond to novel and changing environments.

In our study, we defined two levels of task complexity. The task required that a team monitor 20 items of information and decide between two courses of action (binary choice). The value of information determined the likelihood that a particular course was appropriate. In a simple version of a task, all information was equally diagnostic (i.e., weighted equally) whereas, in a more complex version, nine of the items were irrelevant (i.e., zero weight). Thus, in the simple task, it did not matter what information was communicated (except to avoid redundancy) and the more information that reached the decision maker at the top the better. In contrast, efficient communication in the complex task required that members learn to distinguish important from unimportant information. Based on the foregoing ideas from the communication network and organizational learning literatures, we entertained a flexibility-adaptability hypothesis: the increased communication flexibility afforded by adding lateral links would permit more learning by individual members and thus permit a team to adapt more readily to the complex task. Related to this first hypothesis is a flexibility-complexity hypothesis: communication flexibility would lead to better ultimate performance on the complex task whereas the purely vertical communication structure would result in better performance initially on both tasks and better ultimate performance on the simple task.

Conceptual Overview of the Team Task

Our experimental task is a computer simulation of binary choice under uncertainty. An analogy to conventional statistical hypothesis testing provides a useful way of thinking about the contingencies in the binary choice task. There are two possible states of the world, S_1 and S_2 , and two courses of action, A_1 and A_2 . The consequences of either course of action depend on the actual state of the world which is unknowable within the time available to make a choice. Let V_{ij} be the value (either loss or gain) associated with making the decision A_i when the current state is S_j . Then, in conventional decision theory notation, we can express the expected value of either decision as a function of the probability, P_j , that each state, S_j , exists and the values, V_{ij} . That is, the expected value, EV_1 , of choosing A_1 is

$$EV_1 = P_1 V_{11} + P_2 V_{12}. \quad (1)$$

The expected value of choosing A_2 can be expressed similarly,

$$EV_2 = P_1 V_{21} + P_2 V_{22}. \quad (2)$$

A rational model of action dictates that the decision maker should compare EV_1 and EV_2 , and choose A_1 if $(EV_1 - EV_2) > 0$, and A_2 if $(EV_1 - EV_2) < 0$. By using the fact that $P_2 = 1 - P_1$ (because S_1 and S_2 are mutually exclusive and exhaustive states), this critical comparison can be expressed as

$$EV_1 - EV_2 = P_1 (V_{11} + V_{22} - V_{12} - V_{21}) + V_{12} - V_{22}. \quad (3)$$

That is, the difference between the expected values of the two actions depends on the probability that state S_1 exists and the values of the gains realized by the correct decisions (V_{11} and V_{22}) and the losses incurred by the incorrect decisions (V_{12} and V_{21}). In fact, $EV_1 - EV_2$ will exceed 0 (indicating that A_1 is the rational choice) whenever

$$P_1 > \frac{1}{1 + [(V_{11} - V_{21}) / (V_{22} - V_{12})]}. \quad (4)$$

Applying such formulations to decisions in naturally-occurring situations is notoriously difficult. First, not only is the state of the world unknown, but the objective probability P_1 is also unknown (and perhaps unknowable). Second, the values, V_{ij} , of the various correct and incorrect decisions may be difficult to specify. Therefore, the objective probability is often replaced with a subjective probability (the degree of belief that S_1 is true), and the values, V_{ij} , are replaced with derived scales of psychological worth (utilities). Whereas these are concerns when trying to describe what decision makers actually do, our intent is to use these relationships to evaluate the effectiveness of a decision making process when we can set the probabilities and values.

How P_1 relates to the diagnostic information (integrative complexity) and the amount of diagnostic information (information load) are two ways of operationalizing task complexity. In the present version of the task items of information can assume one of two values: one value increases the probability of A_1 and the other value decreases the probability from the specified baseline. For example, the bearing of the incoming plane may be an intercept course (raising the likelihood of attack) but the altitude may be slightly off (lowering the likelihood of attack).

The value of P_1 is a function of the values of the diagnostic information that the team receives and a baserate, P_0 :

$$P_1 = \underline{f}(P_0, \mathbf{x}), \quad (5)$$

where $\mathbf{x}$ is a vector containing weights denoting the direction and amount of adjustment from P_0 associated with the diagnostic information given to the team. Although the function $\underline{f}$ can be defined by any rule, it is a linear function for the current study. (We also note that even a linear function can be extended to include interaction terms to capture disjunctive and conjunctive relationships -- e.g., a given type of plane at a given approach altitude increases the likelihood of attack but another type of plane at the same approach altitude decreases the likelihood of attack).

Framing the task in this way emphasizes the fact that the number of correct decisions may be an insensitive way to measure team performance and improvement in team performance over time. Outcomes are probabilistic and optimal performance does not ensure correct decisions. That is, a team may follow perfectly the rational model of decision making (correctly assessing the expected values of the alternate courses of action and choosing the one with the higher expected value) and still make a number of incorrect decision over a series of trials. Moreover, a team that is completely unresponsive to the expected values (i.e., unresponsive to the information and the relative payoffs) will be right 50% of the time by simply guessing in a binary choice task. As a result, measures of performance that reflect the covariation of team action with expected values will be more sensitive than simply noting how frequently they make correct decisions.

Parameter Settings for Current Study. We used the pattern of weights in $\mathbf{x}$ (equation 5) to operationalize task complexity. In the simple task, all of the weights are positive and equal whereas in the complex task approximately half of the weights are zero and the rest are positive and equal. Because our primary interest for this study is to examine information flow, we equated the values of the two kinds of correct decisions (i.e., $V_{11} = V_{22} = V_G$) and the two kinds of errors (i.e., $V_{12} = V_{21} = V_L$). Moreover, we made gains and losses symmetrical by setting V_L equal to $-V_G$. Additionally, the baserate P_0 in equation 5 was maintained at .5. As a result, changes in the differential expected values for the two courses of action across trials were driven solely by changes in the patterns of information given to the team. Note, however, that the task would also permit one to systematically vary payoffs or change the baserate over time. If payoffs or baserates were manipulated, the measures that are discussed in the following section could also be used to assess the sensitivity of the team's actions to these changes.

Performance Measures

The computer program controlling the simulated decision making task can monitor several aspects of performance including what information is transmitted, in what temporal order and at what rate, along each link in the communication structure. Additionally, the program computes from the current configurations of information and valuations of correct and incorrect decisions, the expected values of the two possible courses of action. The program keeps a cumulative record of decisions and outcomes across trials.

Based on these basic data, several indices of performance can be obtained. Obvious ones are the proportion of correct decisions and the cumulative net value of outcomes across a set of trials: the gains realized by correct decisions minus the losses incurred by incorrect ones. While these are obvious measures, they lack sensitivity because the probabilistic nature of the task can result in poor outcomes over a series of trials even when a team is using an effective strategy (or, vice versa, an inappropriate strategy can result in a series of correct decisions). Therefore, we used three kinds of performance measures that can potentially provide more reliable and sensitive indicants of team's effectiveness. These measures are based on the notion of sensitivity - - the degree to which various aspects of team performance covary with changes in the differential expected values expressed in equation 3.

Team Decision Sensitivity. A team's decisions should be sensitive to changes in the relative expected values associated with the decision alternatives. Thus, we used a sensitivity index which measures the degree that a team's decisions covary with the difference in the expected values of the two courses of action (equation 3). More formally, let the responses of the team be coded as a variable, R , that is defined as follows:

$$R = \begin{cases} 1, & \text{if the decision is } A_1 \\ -1, & \text{if the decision is } A_2 \end{cases}$$

Also, let D be the difference between the expected values of the two actions: $D = EV_1 - EV_2$. Then, the sensitivity index is the point-biserial correlation, r_{RD} , between R and D . While the point-biserial can in principle vary between -1 and +1, substantially negative values are unlikely in this application because they would suggest that the behavior of the team was sensitive to changes in the expected values across trials but were reacting in the wrong way; this should happen only if there were a fundamental misunderstanding of the task. Thus, the sensitivity index will ordinarily vary between 0 (indicating no sensitivity) and +1 (indicating that the team's decisions varied perfectly with the changes in the expected values).

Another potential problem with this measure is that the maximum value obtainable for a point-biserial correlation can be impacted substantially by the underlying distributions of the variables (R and D in this case; see Crocker & Algina, 1986). For example, a team could be perfectly sensitive to changes in the expected values but biased in their responding such that they made too many or too few A_1 responses. This bias could reduce the maximum obtainable value of the point-biserial to something less than +1. Thus, the sensitivity index can be considered in the context of a bias measure, $B = rf_1^* - rf_1$, which is the difference between the relative frequency of trials for which A_1 was the appropriate response (rf_1^*) and the relative frequency of A_1 responses (rf_1). Moreover, the program will control the distribution of D which can also depress the value of r_{RD} . If these distributional problems arise, an alternate sensitivity indexed, S , may be preferable:

$$S = r_{RD} / \max(r_{RD}),$$

where $\max(r_{RD})$ is the maximum obtainable value of the point-biserial correlation given the underlying distributions of R and D .

Perfect performance over a series of trials would be indicated by a bias of zero ($\underline{B} = 0$) and a sensitivity of one ($\underline{S} = +1$). A team with good sensitivity but erring in the direction of too many A_1 decisions (preemptive attacks in our task) would obtain a sensitivity near one and a positive bias value. Random responding could result in negative, zero, or positive bias (depending on the baserates of appropriate A_1 responses) but sensitivity would be near zero.

Information Transfer and Communication Sensitivity. Given the nature of the task, good performance depends on getting information from the bottom of the communication hierarchy to the top. Moreover, when information is differentially weighted, transmitting the most important information is critical. However, the information transfer may be optimal given the time constraints of the task but performance could still suffer because the information is ignored or used inappropriately once it arrives at the top. Thus, it is useful to evaluate the communication processes of a team as well as its decision performance. One way of indexing communication effectiveness is to note simply the proportion of the available information that makes it to the top. However, when the importance of information varies, the information transfer index should incorporate the relative importance of the information.

Let $\underline{w}_k$ reflect the absolute amount that item k increases (or decreases) P_1 , and $\underline{c}_k$ be 1 if item k reaches the top, 0 otherwise. Then, define $\underline{T}$ as follows:

$$\underline{T} = \sum \underline{c}_k \underline{w}_k / \sum \underline{w}_k,$$

where the summation is across all items of available information. $\underline{T}$ not only reflects the amount of information but also the quality of the information that passes through the communication network.

Of course, this information transfer index can be averaged to track performance across a series or block of trials, but, unlike the decision sensitivity index (which must be computed across trials) its value can also be tracked trial by trial. Additionally, note that the transfer index can be computed between adjacent levels in a communication hierarchy. Thus, the overall transmission effectiveness of the hierarchy can be decomposed into lower-to-middle and middle-to-top communication effectiveness.

Whereas the information transfer index has intuitive appeal, it may be unduly impacted by the volume of information and relatively insensitive to other aspects of communication effectiveness. It does not measure how well the communicated information represents the available information, and, even more importantly, it does not measure how well communicated information reflects changes in the patterns of available information. As a result, some types of potentially effective communication strategies may fair poorly when assessed by the information transfer index. For example, consider a team that develops the implicit strategy of sending all of the information that supports a given course of action (e.g., send only information that suggests attack). Such a strategy, if followed consistently across team members and across time, may be quite effective particularly when limited time prevents sending all information. Nonetheless, the value of the information transfer index would be low for a team using this strategy compared to a team that haphazardly sent as much information as possible.

Thus, we propose a communication sensitivity index that is similar to the aforementioned decision sensitivity index. Let $\underline{C}$ be the difference in the expected values of the two actions based on the information that is received by the decision maker at the top of the hierarchy and $\underline{D}$ be the difference in the expected values based on the entire set of available information (as defined above). Then, the correlation between $\underline{C}$ and $\underline{D}$ across a series of trials would assess the degree to which communicated information systematically reflects changes in pattern of available information.

Member Vote Sensitivity. The rationale for the flexibility-adaptability and the flexibility-complexity hypotheses presumes that team performance depends partly on how well members adapt to the task. Effective communication is promoted by members learning the importance or relevance of information. Thus, we expected that when a centralized decision-maker must rely on communicated information, improvements in team performance would lag behind member performance. In order to evaluate more directly the performance of team members, we solicited their individual decisions after the team decision had been made but before they were informed of either the team decision or the correct decision for each trial. As with team decisions, we computed that proportion of member votes that were correct but also computed an index that is analogous to the decision sensitivity index. More specifically, we correlated across trials the number of team votes for A_1 with the difference in expected values of the two courses of actions. Not only does this measure provide an alternate way of describing members' performance, it also gives an indirect assessment of how well a team might have performed if decision-making were decentralized (e.g., a majority voting process).

Experimental Objectives

The objectives of this study were to: a) examine the interactive effects of communication structure and task complexity on performance and adaptability (improvement in performance); b) examine the effects of stress (time limits and risk of gain and loss) on communication patterns, amount of information transferred, and performance; and c) demonstrate the proposed measures of team decision, communication, and member vote sensitivity. A fourth objective was to obtain verbal protocols describing emergent member and team communication strategies for dealing with time pressures. The results pertaining to this fourth objective are not presented in this report.

Method

Design. The design can be represented as a $2 \times 2 \times 2 \times 2 \times 4$ factorial of communication structure (vertical versus lateral: see Figure 1), task complexity (linear integration with equal or unequal weights), payoff level (gains or losses of either \$.05 or \$.50 per trial), time stress (90- versus 60-sec trials) and trial blocks. A team participated in one of four conditions as defined by structure and task complexity; the payoff, time stress and trial blocks factors were repeated measures factors. Each trial block contained 8 trials: 5 of these were "normal" trials with time limits of 90 sec and payoffs set at the \$.05 level, 1 was a high payoff trial with a 90-sec time limit, 1 was a low payoff trial with a 60-sec time

Team Decision-making Under Threat

limit and 1 was a high payoff trial with a 60-sec time limit. The occurrence of high payoff and short trials within a block was randomly assigned with the restriction that the first trial was always a normal (low payoff/90-sec) trial.

Subjects. University students were recruited via announcements in the student paper, on campus bulletin boards, and during meetings of campus service organizations to participate for up to two hours. The announcements described the general nature of the research and the minimum amount of pay for participation (\$10.00). Three hundred and fifty student participated in 50 7-person teams. Data from six of these teams were unusable because of experimenter errors (4) or equipment failures (2). The remaining 44 teams were distributed evenly across the four conditions created by the communication structure and task complexity factors.

Overview of Experimental Task. The experimental task simulates a situation which requires the decision-making team to choose between two courses of action. Each choice risks loss if it is inappropriate given the current state of affairs and avoids loss if it is appropriate. From the participant's point of view, the state of affairs is not knowable with certainty but the probability that a particular state exists can be inferred from diagnostic information. The task was implemented on a computer network with team members seated at separate work stations.

The simulation program emulated the military decision task that we described earlier. That is, participants were told that they were to play the roles of a ship's crew that is monitoring the approach of a potentially hostile plane which could be either "friend or foe". If the plane intends to bomb the ship, there is a limited time for the crew to launch a preemptive strike. On the one hand, given the political climate and current military policy, shooting down the plane is costly and to be avoided unless absolutely necessary. On the other hand, failing to launch a preemptive strike if the plane intends to attack is also costly. The costs were represented as a loss of money that team members were paid for their participation.

It was explained that, as is typical of such situations, they would not be able to ascertain with certainty the actual intent of the approaching plane until it was too late to react. However, they would have information that allowed them to estimate the likelihood of an actual attack. For example, the plane's velocity, altitude, and heading could provide useful diagnostic information that either raised or lowered the judged likelihood of attack. Moreover, the ship's location (contested versus uncontested international waters), current weather conditions, recent intelligence reports, and the like, may provide additional indicators of the likelihood of hostile action by the approaching plane. The types and meaning of the diagnostic information were reviewed thoroughly. However, it was further noted that the various types of information may not be equally informative and that success on the task would depend in part on their ability to distinguish important from unimportant information. This instruction was given even when items were equally-weighted. Thus, participants in the simple and complex (equal- and unequal-weight) task conditions began with the same instructional set.

These instructions were supplemented with a diagram (see Appendix A) that summarized the types of available information, who had direct access to the information, and who could communicate with whom in the network. This diagram was available to each member throughout the experiment.

There were 20 items of information that were divided into two sets: situational information consisted of 4 items (weather conditions, time of day, location of ship, and the ship's current defense capability) that were always accessible by all team members; specialized information consisted of 16 items, each of which was directly accessible by only one of the four team members at the bottom of the communication hierarchy. Each of the four lower members was assigned four of the 16 specialized items (see Appendix A). In the unequal weight version of the task, one situational and eight of the specialized (two for each of the four lower nodes) items did not affect the probability of attack ("zero-weight"). For each team, these zero-weight items were randomly chosen.

The patterns of information values across trials were constructed so that each block contained two trials in each of the four categories defined by: a) probability of attack greater than .5 versus less than .5; and b) the pattern of situational information did or did not reflect the balance of "safe" and "danger" items in the overall pattern of information. (Note: In half of the cases where situational information did not accurately reflect the overall pattern, it was ambiguous: 2 "safe" and 2 "danger" items.) In this way, the expected base rate was maintained at .5 within each block, and performance based solely on situational information would have resulted in relatively poor performance. The location of these types of trials was randomly determined within each trial block. Appendix B gives the distributions of values for probability of attack across the 32 trials for the equal- and unequal-weight versions of the task. Whereas the distributions are symmetrical around .5, values near .5 did not occur.

Procedures. Participants were scheduled in seven-person teams. Upon arriving, they were assigned randomly to computer workstations; each workstation had a designated role (e.g., Ship Commander, Radar Captain, Air Position Radar, etc.) that was associated with a particular node in the communication network (see Appendix A). After team members had completed the introduction to the task, they completed 4 practice trials and then had the opportunity to discuss with the experimenter any questions or concerns about the procedures, the task, or their participation in the study.

Then the team completed 32 trials that were divided into 4 8-trial blocks. A trial was terminated either by the Commander making a decision or by expiration of trial time. If time expired before the Commander made a decision, the default decision was not to fire. (Team members were aware of this default.) After each trial, each team member reported his/her individual decision and then received feedback about the Commander's decision and the correct decision. Accumulated earnings were also displayed. Between trial blocks, team members were also given the opportunity to send a message to any other team member. The recipients of these inter-block messages were not constrained by the communication links that were operative during the trials.

Each trial began with a display of the available time for that trial; the shorter 60-sec trials were also signaled by a bell. The trial clock was updated continuously throughout the

Team Decision-making Under Threat

trial. Additionally, high payoff trials were accompanied by a flashing dollar sign that was displayed under the trial clock in the computer monitor.

During each trial, team members could take the following actions: assess the values of items to which they had direct access; send information that they had either directly assessed or received from another member; receive information sent to them by other members; review items already assessed or received; and request specific items from another specified team member. Members could send, receive, and request information only through the operative links in the network. When a message was sent, the recipient was notified that his/her mail file contained a message and it remained in the mail queue until it was received. Messages sent to a node had to be received in the order that they were sent. A request for information had to be acknowledged in that the request remained on the screen and other actions were suspended until the "space bar" on the keyboard was pressed. However, the target of the request did not have to comply with the request (and could not comply if he or she had not already assessed or received the desired information).

The Commander (central decision maker) had the additional option of making a decision. A decision could be made at any time and once it was made the trial was terminated. Each station was notified of this action when it occurred and all other actions were suspended at that point. However, they were not told what the decision was until after they had registered their individual decisions.

At the end of the session, team members responded to two open-ended questions that were designed to obtain verbal reports of strategies that they had employed and strategies that they would recommend to others. The first question asked them to describe the strategy (or strategies) that they used to select information to send to other members. The second question asked them what advice they would give to members of a new team.

Results

Volume of Communication Activity

Information Transfer. The information transfer index measures the quality of information held by the decision maker at the top of the hierarchy when a decision is made. In the present application, it is equivalent to the proportion of the available important information that is either communicated to the Commander or, in the case of situational information, directly measured by the Commander. As is evident in Table 1, the value of the information transfer index increased steadily across trial blocks, $F(3, 120) = 122.28$, $p < .0001$. Because situational information could be directly accessed by the Commander, the transfer index was recomputed for only specialized information which had to be communicated and, thus, is a cleaner measure of communication activity. A similar block effect emerged; the proportion of specialized information reaching the Commander increased from .19 in the first block to .51 in the final block, $F(3, 120) = 108.46$, $p < .0001$.

Communication structure and task complexity did not significantly impact the transfer index. However, the Commander acquired less of the important situational and specialized

Team Decision-making Under Threat

Table 1. Communication and Information Processing Activity across Trial Blocks.

Measure	Block			
	1	2	3	4
Information Transfer				
All Information	.34	.45	.53	.60
Specialized Only	.19	.32	.41	.51
Send				
All Information	17.4	22.9	25.9	30.1
Situational Only	2.1	2.5	2.7	3.0
Receive	12.7	18.7	22.7	27.2
Loss	4.8	4.2	3.2	2.8
Assess				
All Information	31.0	34.8	36.2	37.2
Situational Only	15.1	18.4	19.9	20.9
Requests	5.8	6.7	7.3	7.1
Review	5.5	5.9	5.0	4.8

information during 60-sec trials ($M = .41$) than during 90-sec trials ($M = .51$), $F(1, 120) = 149.12$, $p < .0001$. This drop is primarily due to the communication of important specialized information (for 60-sec trials, $M = .28$ and, for 90-sec trials $M = .40$), $F(1, 120) = 154.13$, $p < .0001$.

Level of payoff also had a statistically significant, although modest, impact on the values of the information transfer index. For example, slightly less of the important specialized information was communicated during the low payoff ($M = .32$) than during the high payoff ($M = .35$) trials, $F(1, 120) = 8.52$, $p < .01$. This difference translates to less than one item of important information per trial.

Information Sent, Received, and Lost. The mean number of items sent and received, summed across all team members, are given in Figure 2 for the vertical and lateral communication structures across the four trial blocks. (The aggregate block means are also given in Table 1.) There is a systematic increase across blocks for both numbers of items sent, $F(3, 117) = 141.37$, $p < .0001$, and received, $F(3, 117) = 188.67$, $p < .0001$. (One team's data from the lateral/equal-weight condition are omitted from these results because communication data were lost for one of the lower-level workstations; thus the error degrees of freedom for block effects is reduced by 3.) Moreover, there is a significant interaction of communication structure with trial blocks for items sent, $F(3, 117) = 6.10$, $p < .001$, and received, $F(3, 117) = 3.33$, $p < .02$. Teams with a lateral communication structure increased their communication activity more across trial blocks than did teams with a vertical structure.

The difference in the solid and shaded bars in Figure 2 represents items that were lost in the mail queues and never received. Information loss is sizable (over 5 items, on average, per trial) during the first block but diminishes somewhat over blocks, $F(3, 117) = 21.56$, $p < .0001$. There is also a communication structure by block interaction, $F(3, 117) = 3.58$, $p < .02$, due to loss decreasing more over blocks in vertical than in lateral communication structures.


During 60-sec trials, teams sent ($M = 20.8$) fewer items than during 90-sec trials ($M = 25.8$), $F(1, 117) = 412.29$, $p < .0001$. Likewise, they received fewer items during 60-sec ($M = 16.0$) than during 90-sec trials ($M = 22.2$), $F(1, 117) = 558.86$, $p < .0001$. While team members are clearly doing less during the shorter trials, it is interesting to note that the proportional reduction in number of messages sent and received is not as large as the proportional reduction in time. That is, teams apparently responded to having less time by working faster.

Teams also sent slightly more messages during high payoff ($M = 23.4$) than during low payoff ($M = 22.8$) trials, $F(1, 117) = 5.95$, $p < .05$. There was a corresponding tendency for them to receive more messages during high payoff, as compared to low payoff, trials; however, the difference was statistically nonsignificant, $F(1, 117) = 3.78$, $p < .1$.

Other Activities. Table 1 presents the mean number of times that a team (summed across members) assessed, requested, and reviewed information during a trial as a function of trial blocks. Across blocks, there was an increase in the assessments, $F(3, 117) = 55.62$, $p < .0001$, and requests, $F(3, 117) = 4.97$, $p < .01$. However, the number of

Figure 2

Items Sent and Received: Communication Structure by Block



Loss = Sent-Received

times that team members reviewed information tended to decrease across blocks, $F(3, 117) = 31.53, p < .0001$.

In addition to these block effects, teams in the vertical communication structure assessed the value of items more often ($M = 36.2$ times per trial) than did lateral structure groups ($M = 33.4$), $F(1, 39) = 7.45, p < .01$.

Assessing and Sending Situational Information. All members of a team could directly assess situational information, but only the Commander needed to assess it for it to be considered in the team's decision. If only the Commander had assessed the four items of situational information on each trial, the mean number of assessments in Table 1 would have been 4. The means for number of assessments of situational information were considerably larger than this minimally necessary level, increasing from 15.1 during the first trial block to 20.9 during the last trial block, $F(3, 117) = 72.09, p < .0001$. The assessment of situational information was slightly more frequent in teams with vertical structures ($M = 19.7$) than in teams with lateral communication links ($M = 17.5$), $F(1, 39) = 5.19, p < .03$. This difference is consistent with the aforementioned overall tendency for vertical teams to assess the value of items more often than lateral teams.

Whereas team members, except for the Commander, did not need to assess situational information, they may have used it in their attempt to learn about the relative importance of information and, for teams with vertical communication structures, assessing situational items was the only feasible way for members at the bottom to broaden the informational basis for making judgments of relative importance. However, there was no apparent need to send situational information. Nonetheless, teams not only sent situational items but also the frequency of sending them increased over trial blocks from about 2 initially to about 3 finally, $F(3, 117) = 3.75, p < .02$. Thus, teams were processing situational information more than necessary for efficient performance.


Communication Sensitivity

We proposed a communication sensitivity index that is the covariation across trials of the differential expected values of the two courses of action (fire versus not fire; equation 3) computed first on the total available information and second on the information acquired by the Commander. In the present application, changes in expected values across trials are due solely to changes in the information. Therefore, the value of the sensitivity index would be +1 if the balance of important items supporting "fire" (or "not fire") that reached the Commander varied directly with the balance of available "fire" (or "not fire") items in the total set of information. Of course, this could occur if all of the information reached the Commander on every trial. As the results for the transfer index document, teams did not even approach this extreme of complete information transfer. However, the sensitivity could be perfect even when information transfer is less than complete.

The communication sensitivity index was computed across the eight trials in each trial block for each group. As illustrated in Figure 3, communication sensitivity increased over trial blocks; the overall linear trend is significantly greater than zero, $F(1, 40) = 41.86, p < .0001$. Moreover, there was a task complexity by linear trend interaction, $F(1, 40) = 4.73, p < .04$. This interaction arises because communication was less sensitive initially for

Figure 3

Informational Sensitivity of Communications: Equal- and Unequal-Weight Tasks



N=44

the unequal-weight task than for the equal-weight task, $F(1, 40) = 6.68, p < .02$, but, by trial blocks 3 and 4, this difference had virtually disappeared.

The results for the communication sensitivity index provide no support for the flexibility-adaptability hypothesis. Vertical and lateral communications structures gave rise to nearly identical trends across time for both types of task.

Team Decisions


The proportions of correct team decisions were computed within each trial block. As depicted in Figure 4, the proportion of correct decisions increased across trial blocks for the equal-weight task, $F(1, 40) = 7.51, p < .01$ for positive linear trend. In contrast, there was no systematic improvement for the unequal weight task, $F(1, 40) = 0.52$ for the linear trend. As a result, teams with the equal-weight task were correct more often during the final block than teams with the unequal-weight task, $F(1, 40) = 12.10, p < .002$. As with communication sensitivity, communication structure had little effect on the overall proportion of correct decisions or on the improvement of performance over time. Moreover, it is noteworthy that the aforementioned improvement in communication sensitivity for the unequal-weight task (evident in Figure 3) did not result in improved team accuracy.

A similar pattern emerged when we examined the team decision sensitivity index. Recall that this index captures the degree to which team decisions change in response to changes in the difference between the expected values of the decision alternatives. For the present case, this is equivalent to the covariation of decisions with the relative numbers of "fire" and "no fire" items in the set of important information. The index was computed for each team within each trial block. Figure 5 gives these values for the two types of tasks across trial blocks. For the equal-weight task, there is a significant linear improvement across blocks, $F(1, 40) = 7.72, p < .01$; however, for the unequal-weight task, the linear trend is nonsignificant, $F(1, 40) < 1$. By block 4, the difference between the equal- and unequal-weight task is approaching statistical significance, $F(1, 40) = 3.67, p = .06$. Again, there is not strong evidence that communication structure impacted either the level of sensitivity or its improvement over time; however, there is a trend for the vertical communication structure to show more improvement and better ultimate sensitivity, particularly on the equal-weight task. Specifically, for the final block on the equal-weight task, teams with only vertical links had a sensitivity index of .67 as compared to .47 for teams with lateral links although this difference was not statistically significant, $F(1, 40) = 1.38, p < .25$.

Comparing the communication sensitivity values in Figure 3 and the decision sensitivity values in Figure 5 highlights two points. First, the decision sensitivity is noticeably lower than communication sensitivity. Second, the rather sizeable improvement in communication sensitivity for the unequal-weight task did not translate into a corresponding improvement in decision sensitivity. These differences underscore a limitation of centralized decision-making. The ultimate responsibility for integrating the information rests with the decision-maker who has little control over what information he or she receives. Particularly for the unequal-weight task, team performance is dependent on either the decision-maker learning what information to ignore or on the team sending only high quality information.

Figure 4


Proportion of Correct Team Decisions: Equal- and Unequal-Weight Tasks



N=44

Figure 5

Informational Sensitivity of Team Decisions: Equal- and Unequal-Weight Tasks



However, the Commanders who worked with the unequal-weight task were receiving unreliable patterns of information during the first and, to some degree, the second trial blocks. This undoubtedly interfered with the development of effective information integration of strategies.

Member Decisions (Votes)

Our method of collecting individual decisions at the end of each trial allowed us to assess indirectly how effective decentralized decision-making might have been. The proportion of members who were correct is graphed in Figure 6 by trial block for the equal-weight and unequal-weight tasks. Comparing these results with the results for team decisions reveals similar patterns: relatively poor performance on the unequal-weight task with no evidence of improvement over time and higher performance on the equal-weight task with improvement over time. The only notable difference is that members tended to do better than the team initially (blocks 1 and 2) on the equal-weight task although team and member performance was nearly identical by block 4.

We also covaried the number of members who decided "fire" with the changes in the relative expected values of the two actions. This member vote sensitivity index has a metric identical to the communication sensitivity and team decision sensitivity indexes that were presented earlier. As is apparent in Figure 7, the votes of members were more responsive to changes in the important information for the equal-weight task than for the unequal-weight task, $F(1, 40) = 32.88, p < .0001$. Additionally, the interaction of task, communication structure, and linear trend over blocks was significant, $F(1, 40) = 9.61, p < .005$. The interaction of communication structure and trend is not significant for the equal-weight task, $F(1, 40) = 2.57, p > .1$, but is significant for the unequal-weight task, $F(1, 40) = 7.73, p < .01$. For the unequal-weight task, the sensitivity of member votes in the vertical communication structure improved significantly over time, $F(1, 40) = 4.98, p < .05$. In contrast, member vote sensitivity in teams with lateral communication links seemingly deteriorated over time although the negative linear trend was not significant, $F(1, 40) = 2.89, p < .1$.


Discussion

Communication Processes

Communication activity increased substantially as teams became more experienced. They were not only sending and receiving more messages in the last trial block, but were also getting more of the important information to the Commander. Moreover, there was a decrease in the amount of information that was lost in the network.


We were particularly interested in the degree to which the Commander received patterns of information that changed systematically with the available information. One interpretation of our communication sensitivity index (Figure 3) is in terms of the degree to which variations in information received by the Commander are predictive of the input data base. For the equal-weight task, any process that generated items that were representative of the balance of "safe" and "danger" in the input set would produce a high value of the index.

Figure 6



N=44

Figure 7



However, for the unequal-weight task, the index would be suppressed to the degree that teams communicated unimportant items to the exclusion of important ones. Apparently as a result of this added source of noise in the unequal-weight task, communications in the early trial blocks were more sensitive when all information was equally weighted. Nonetheless, teams working on the unequal-weight task caught up by the third trial block so that teams, regardless of task, were calibrating reasonably well their communications to the top of hierarchy with the input information.

The flexibility-adaptability and flexibility-complexity hypotheses asserted that such improvements over time on a complex task would be more evident in teams with flexible communication options. However, there was no evidence in our data that the presence of lateral communication links promoted such improvement. The rationale for more improvement on a complex task when communication flexibility exists rests on at least two premises. First, learning, or, more generally, adaptability, is promoted by being exposed to more information. This seems particularly critical in our unequal-weight task where effective performance depends on knowing the relative importance of different types of information. Second, the presence of lateral links would facilitate the dispersion of information through the network.

Teams with lateral links, as compared to those without, did increase more rapidly the volume of their communications and were sending and receiving more messages at the end. However, a follow-up analysis revealed that teams used the lateral links sparingly. During the first trial block, less than 1% of the messages sent were routed along lateral links, and the percentage had increased to a rather modest 9% by the last trial block. In sum, the use of lateral links was too little and too late to expect any substantial benefits from increased lateral dispersion of information.

There was also evidence of certain communication inefficiencies that persisted even as teams gained experience. Most notably, teams communicated situational information even though it should have been readily apparent that all members had direct access to this information. And, the rate of sending situational information tended to increase with experience. Additionally, more detailed analyses of the content of messages showed that about 10% of the messages sent were redundant, primarily due to team members sending the same message more than once. Whether these inefficiencies were due to inattention or to members trying to emphasize certain types of information is unclear, but their prevalence, given the design of our task, is noteworthy.

First, there was relatively little situational (commonly available) information in our task; only 4 of the 20 types of information were situational. Second, as already noted, it was obvious that such information was directly accessible by all members. It is not hard to imagine contexts where the amount of commonly-held information may be more extensive than the specialized information and the fact that it is commonly-held less salient. In these cases, the aforementioned tendencies to communicate unnecessarily such information may be compounded (Stasser, Taylor, & Hanna, 1989; Stasser, 1992).

Finally, communication redundancy should not have been a problem in our teams. Specialized information was given to one, and only one, of the team members, and this

specialized distribution was apparent to all. Thus, redundancy was created primarily by single members sending redundant messages. In many distributed information environments, there is not such a simple mapping of information into team member roles. Some information may be available to more than one member but not to all. In this event, the information would need to be communicated, but members would need to coordinate their communications to avoid redundancy.

Teams with hierarchal communication structures are vulnerable to two types of mistakes when information is redundantly input into the network. If members ignore the fact that more than one person can communicate the information (or such input redundancy is not apparent to them), then considerable redundancy could be introduced as messages are sent up in the hierarchy. Conversely, if all members take note of the input redundancy and focus their efforts on communicating information that only they individually can directly access, inputting information redundantly may actually result in its being lost. Such communication coordination problems are much more complex than the ones that we presented to teams in this study. Yet, their apparent lapses of attention to the simple coordination problems that they encountered suggests that teams may not be particularly good at effectively managing redundancy.

Decision Performance

We have suggested that measuring performance by the number of correct decisions may be an insensitive way to track performance when outcomes are probabilistically related to information. Optimal communication and information use will not yield perfect performance. Additionally, in a binary choice task, ineffective and degenerate processes should yield, on average, 50% correct decisions. For both our equal- and unequal-weight tasks, optimal use of information would have resulted in a little over 70% correct decisions (see Appendix B). Thus, by the final trial block, team decision accuracy was near this optimal level for the equal-weight task, but never approached it for the unequal-weight task. In fact, neither decision accuracy (Figure 4) nor decision sensitivity (Figure 5) improved for the unequal-weight task. That is, the systematic improvement in communication sensitivity for the unequal-weight task did not translate into improved decision performance.

There was, however, a sizeable improvement in member vote sensitivity on the unequal-weight task in teams with a vertical communication structure. Again, this improvement of member sensitivity to variations in values of important information was not accompanied by any apparent improvement in team decisions. Moreover, it is the only evidence that we obtained for an interaction of communication structure with task complexity. However, it is a pattern that is contrary to that predicted by the flexibility-complexity hypothesis that was derived from the traditional literature on communication networks in problem-solving groups (see Shaw, 1978, for a review of this research). Indeed, it seems that the added flexibility of communications allowed by the lateral links in our teams, if anything, interfered with the calibration between member decisions and available information when irrelevant information was present.

One implication of the pattern of results for the unequal-weight task is that improvements in team decision performance lag behind both improvements in communication

effectiveness and member sensitivity to information. There are several possible reasons for such a lag. First, and foremost, when decision-making is centralized but information is distributed, the decision-maker must rely on members to convey information. However, until members develop reliable communication strategies, the information that they convey, from trial to trial, is infused with additional noise that makes learning more difficult. Second, when information is differentially important or diagnostic, effective communication depends on members distinguishing, at least to some degree, important from irrelevant information. Third, the volume of information reaching the decision-maker is not the critical mediator of accurate decisions. This is obvious when information is not equally diagnostic, but sheer volume may also be disruptive when all information is important because it creates information overload for the centralized decision-maker. In other words, it may be better for a team to send only partial information (i.e., reduce communication volume) and concentrate on strategies that reliably reflect the implications or patterns of the available data.

More generally, there are two features of probabilistic cue learning tasks with distributed information that retard team learning. First, because information is distributed, any team member has easy and reliable access to only a subset of information until effective communications strategies emerge. However, as we have noted, effective communication, particularly under severe time limits and when a large part of the information is irrelevant, depends to some degree on members being able to judge the importance of information. Thus, these interdependent processes of cue learning and communication seemingly produce an impasse when teams face novel tasks and information is distributed. This impasse may be partly due to the constraints on interaction that our computer-mediated task imposed on members. For example, in many situations, teams have the opportunity to reflect on their prior performance and exchange information after the fact without the constraints of limited time and restricted communication channels ("debriefing"). Debriefing, and the enriched feedback that it permits, may be fundamentally important to a team's ability to adapt quickly (i.e., with relatively few "trials") to novel tasks.

A second feature of our task which makes learning difficult is its probabilistic nature. Just as we have argued that decision accuracy is a fairly insensitive measure of performance, it is also a degraded form of feedback. Nonetheless, it was the only feedback that team members received in our study. Unfortunately, this is not an unrealistic feature of our task. Few decision environments or tasks contain a circumscribed set of information that determines, without uncertainty, a correct response. In fact, we have avoided one source of uncertainty in that the value of information at any point in time was determined whereas, in more realistic settings, information may be invalid, or inaccurately reported.

One way to alleviate partially this problem is to train members to focus more on long-term success rates and less on the feedback of immediately preceding trials. Although only indirectly related to the current problem, individual decision-making research documents that people are not sensitive to base rate information (see, for example, Argote, Seabright, and Dyer, 1986; Nesbitt and Ross, 1980). Similarly, getting team members to focus on success rates of performance over several trials without overreacting to immediate success or failure may be difficult. One way might be to use measures that summarize success over a series of

trials or episodes as feedback, such as our vote sensitivity index, rather than trial-by-trial success or failure. From a training perspective, it may also be possible to produce more effective communication more quickly by providing feedback that captures communication effectiveness (e.g., feedback based on measures such as redundancy, communication sensitivity, and information transfer).

Centralized or Decentralized Decision Making

Our procedures forced teams to use centralized decision making. Notwithstanding this procedural constraint, we obtained some evidence that teams would have performed better if decision making had been decentralized. On the equal-weight task, members' decisions were correct more often on the early trials than were team decisions. For teams with a vertical communication structure and working on the unequal-weight task, the sensitivity of members' decisions to variations in the input information improved over time whereas the sensitivity of team decisions did not. Additionally, we compared actual team performance to the predicted performance of teams generated by using individual members' decisions in a majority decision scheme model (Davis, 1973). Overall, teams were correct on 60% of the trials. The majority scheme yielded a predicted 67% success rate.

On the one hand, this may be a conservative estimate of how well teams would have done if members had been able to vote or make decision recommendations. Our procedures did not reward team members for correct individual decisions and they were aware that their individual choices did not directly impact team performance. If they had felt that their individual decisions would potentially impact team performance, they may have expended more cognitive effort and processed information more carefully and, thereby, improved individual success rates beyond what we observed.

On the other hand, allowing individuals to make decision recommendations or basing the team decision on some form of voting may make team performance particularly susceptible to the way in which information is distributed and may not be well suited for decision tasks for which configurations of information (not simple linear integration) are crucial. For example, locally available sets of information in a network may not adequately represent, or may even misrepresent, the information that a team has collectively (Stasser, 1988, 1992). If configurations of information are important, as they frequently are in natural contexts, it may be necessary to collate items arriving from different nodes of the network before one can make an informed judgment. In these cases, many of the team members, particularly in a hierarchically-organized team working with a distributed information base, may rarely have sufficient information to determine the presence or absence of a critical profile. Finally, we note that allowing team members to make recommendations may bias their communications. For example, members concerned with justifying their own decisions, may selectively attend to information and edit their communications to this end (Tetlock and Kim, 1987). As a result, communicated information may be severely biased.

Future Research

The results of this study suggest that a) team performance does not improve substantially when they have to learn to distinguish important from unimportant information based on the information that they received during a trial; b) team member performance

tends to improve more rapidly than team performance when decision-making is centralized; and c) team communication is susceptible to redundancy even when the input pattern of information should minimize redundancy.

The Benefits of Debriefing. The software for the experimental task is currently being redesigned to permit more complete feedback between trials. In particular, the debriefing option will give team members the entire pattern of information as well as the correct decision for the preceding trial. If we are correct in assuming that learning is retarded because members have to work with incomplete information, this option should result in faster improvement.

Input Redundancy and Knowledge of Information Distribution. One objective is to examine team strategies to deal with input redundancy. We expect that when this input redundancy is not apparent to members, their communications will compound the redundancy. However, when the input redundancy is apparent, teams may tend to omit the information particularly when time is limited.

Decentralized Decision Making. The current revision of the software is also designed to include a vote option that will permit team members to forward recommendations as well as information. Our current results suggest that decentralized decision making may result in better performance and faster improvement when a task involves simple linear integration of information. However, when teams have to recognize configurations of information (critical profiles), centralized decision-making may be superior.

Heuristics for Team Performance. We are currently coding the content of the inter-block and final messages. From these we hope to extract commonly expressed strategies and conceptions of how the team members think the team should have approached its task. Ultimately, we want to examine the efficacy of heuristics to guide communications under time limits. For a heuristic to be useful, it must not only work (yield reliably good performance) but also be workable (a rule that team members can use under stress). Moreover, we expect that a useful heuristic will need to make sense to team members (i.e., be plausible).

References

- Argote, L., Seabright, & Dyer, L. (1986). Individual versus group use of base-rate information. Organizational Behavior and Human Decision Processes, 38, 65-75.
- Bass, B. M. (1983). Organizational decision making. Homewood, IL: Irwin.
- Burnstein, E., & Vinokur, A. (1977). Persuasive argumentation and social comparison as determinants of attitude polarization. Journal of Experimental Social Psychology, 13, 315-332.
- Crocker, L., & Algina, J. (1986). Introduction to classical and modern test theory. New York: Holt, Rinehart, & Winston.
- Davis, J. H. (1973). Group decision and social interaction: A theory of social decision schemes. Psychological Review, 80, 97-125.
- Davis, J. H. (1969). Group performance. Reading, MA: Addison-Wesley.
- Dublin, R. (1959). Stability in organizations. In M. Haire (Ed.). Modern organization theory (pp. 218-253). New York: Wiley.
- Duncan, R., and Weiss, A. (1979). Organizational learning: Implications for organizational design. In B. M. Staw (Ed.). Research in organizational behavior (Vol. 1, pp. 75-123). Greenwich, CT: JAI Press.
- Eysenck, M. W. (1976). Arousal, learning, and memory. Psychological Bulletin, 83, 389-404.
- Forsyth, D. R. (1983). An introduction to group dynamics. Monterey, CA: Brooks/Cole.
- Gladstein, D. L., & Reilly, N. P. (1985). Group decision making under threat: The TYCOON game. Academy of Management Journal, 28, 613-627.
- Hoffman, R. L. (Ed.). (1979). The group problem solving process: Studies of a valence model. New York: Praeger.
- Janis, I. L., & Mann, L. (1977). Decision making: A psychological analysis of conflict, choice, and commitment. New York: The Free Press.
- Kaplan, M. F., & Miller, C. E. (1983). Group discussion and judgment. In P. B. Paulus (Ed.). Basic group processes (pp. 65-94). New York: Springer-Verlag.

- Katz, D., and Kahn, R. L. (1978). The social psychology of organizations (2nd. Ed.). New York: Wiley.
- Lazarus, R. S. (1966). Psychological stress and the coping process. New York: McGraw-Hill.
- Leik, R. K., & Gifford, G. A. (1986). Decision making under threat of disaster: An experimental simulation. In A. R. Harris (Ed.). Rationality and collective belief. Norward, NJ: ABLEX.
- McGrath, J. E. (1984). Groups: Interaction and performance. Englewood Cliffs, NJ: Prentice-Hall.
- Nisbett, R. E., & Ross, L. (1980). Human inference: Strategies and shortcomings of social judgment. Englewood Cliffs: Prentice-Hall.
- O'Reilly, III, C. A. (1983). The use of information in organizational decision making: A model and some propositions. In L. L. Cummings and B. M. Staw (Eds.). Research in organizational behavior (Vol. 5, pp. 103-139). Greenwich, CT: JAI Press.
- Shaw, M. E. (1964). Communication networks. In L. Berkowitz (Ed.). Advances in experimental social psychology (Vol. 1). New York: Academic Press.
- Shaw, M. E. (1978). Communication networks fourteen years later. In L. Berkowitz (Ed.). Group processes. New York: Academic Press.
- Stasser, G. (in press). Information salience and the discovery of hidden profiles by decision-making groups: A "thought experiment". Organizational Behavior and Human Decision Processes.
- Stasser, G. (1992). Polling of unshared information during group discussion. In S. Worchel, W. Wood, & J. A. Simpson (Eds.). Group process and productivity (pp. 48-67). Newbury Park, CA: Sage.
- Stasser, G. (1988). Computer simulation as a research tool: The DISCUSS model of group decision making. Journal of Experimental Social Psychology, 24, 393-422.
- Stasser, G., & Davis, J. H. (1981). Group decision making and social influence: A social interaction sequence model. Psychological Review, 88, 523-551.

Team Decision-making Under Threat

- Stasser, G., Taylor, L. A., & Hanna, C. (1989). Information sampling in structured and unstructured discussions of three- and six-person groups. Journal of Personality and Social Psychology, 57, 67-78.
- Stasser, G., & Titus, W. (1985). Pooling of unshared information in group decision making: Biased information sampling during discussion. Journal of Personality and Social Psychology, 48, 1467-1478.
- Stasser, G., & Titus, W. (1987). Effects of information load and percentage of shared information on the dissemination of unshared information during group discussion. Journal of Personality and Social Psychology, 53, 81-93.
- Stasser, G., & Stewart, D. (in press). The discovery of hidden profiles by decision-making groups: Solving a problem versus making a judgment. Journal of Personality and Social Psychology.
- Staw, B. M., Sandelands, L. E., & Dutton, J. E. (1981). Threat-rigidity effects in organizational behavior: A multilevel analysis. Administrative Science Quarterly, 26, 501-524.
- Steiner, I. D. (1972). Group processes and productivity. New York: Academic Press.
- Tillman, R. (1960, May-June). Committees on trial. Harvard Business Review, pp. 4-5, 8, 11-12, 162-164, 166, 168, 171-173.
- Zajonc, R. B. (1965). Social facilitation. Science, 149, 269-274.

Team Decision-making Under Threat

Appendix A

Examples of Workstation Schematics for the Vertical and Lateral Communication Structures

You are here

Ship Commander

Radar Captain

Communication Capt.

Air position radar

Coordination Radar

Radio Operator

Intelligence Post

Safe item	Danger item
Aircraft heading	
Distancing range	Closing range
Cruising altitude	Attack altitude
Ascending	Decending
Cruising speed	Attack speed

Safe item	Danger item
Aircraft radar	
Navigation mode	Target search mode
Defense profile	Offensive profile
No jamming effort	Jamming effort
No support planes	Attack air support
	Support planes

Safe item	Danger item
Reconnaissance	Radio message type
Interceptor	Type of aircraft
Friendly response	Attack contact
Non-enemy profile	No radio response
	AWAC confirmation
	Enemy profile

Safe item	Danger item
No hostilities	Recent hostilities
Routine training	Probable mission
Neutral	Unknown
Low preparedness	Diplomatic status
	Enemy defenses
	Strained
	High preparedness

Situational information


Safe item	Danger item
Overcast sky	Weather conditions
Dark	Clear sky
Uncontested waters	Time of day
Heavily-armed	Daylight
	Location of ship
	Contested waters
	Ship defense state
	Vulnerable

Key:

Each of the four lower stations (i.e., Air position radar, Coordination radar, Radio operator, and Intelligence post) have access to the items listed beneath their respective positions.

All seven positions have access to the situational information (listed at left).

Areas in grey are not directly accessible to your position.



Situational Information (available to all)			
SAFE	DANGER		
<i>Weather Conditions</i>			
Overcast Sky	Clear Sky		
<i>Time of Day</i>			
Dark	Daylight		
<i>Location of Ship</i>			
Uncontested Waters	Contested Waters		
<i>Ship Defense State</i>			
Heavily-armed	Vulnerable		

KEY:

Each of the four lower stations have access to the items listed beneath their respective positions.

All seven positions have access to the situational information (listed at left).

Shaded positions are not directly accessible from your position.

Appendix B

The Distributions of Probabilities of Attack across
32 Trials for the Equal- and Unequal-Weight Tasks

Team Decision-making Under Threat

Equal-Weight Task

Probability of Attack	Number of Trials
1.00	1
0.83	4
0.67	7
0.58	4
0.42	4
0.33	7
0.17	4
0.00	1

Note: Optimal Expected Performance is 22.7
correct decisions out of 32 (0.71).

Unequal-Weight Task

Probability of Attack	Number of Trials
1.00	1
0.82	4
0.73	2
0.64	9
0.36	9
0.27	2
0.18	4
0.00	1

Note: Optimal Expected Performance is 23.0
correct decisions out of 32 (0.72).

@

Appendix B

Coordination of Communication in Hierarchically Organized Teams: Working Paper

**** Working Paper ****
Do not cite without permission

Coordination of Communication in
Hierarchically Organized Teams

Garold Stasser, Steve Hinkle, Pamela Ely
and Lee Fox-Cardamone

Department of Psychology
Miami University
Oxford, OH 45056

Abstract

Seven-person teams completed 32, limited time trials of a military, situation assessment game. Communication links were constrained to simulate a hierarchical organization with critical information available to four members at the bottom and decision-making centralized at the top. Thus, successful performance required information to flow quickly from the bottom of the hierarchy to the top. Eight items of information were available to two members at the bottom (partially shared) and eight items were available to only one member (unshared).

Two features of decision environment were manipulated: knowledge of information distribution and communication feedback. Some teams had limited knowledge of how information access was distributed among members; each member knew only what he/she could and could not access. Other teams had elaborated knowledge of information access; for every item in a member's domain of access, he/she also knew if others could access that item. Some teams received detailed feedback about what information reached the centralized decision-maker on each trial; other teams did not receive such feedback.

As predicted based on a collective information sampling model, teams that knew little about how information was distributed accessed and communicated more partially shared than unshared information. However, elaborated knowledge of information access reversed this pattern: more unshared than partially shared information reached the decision maker. It was expected that communication feedback would reduce over trials these biases by informing members of systematic omissions in the patterns of information reaching the decision-maker at the top. Feedback coupled with elaborated knowledge of information access did equalize the amount of unshared and partially shared information reaching the top of the organization during the last eight trials. However, when teams had limited knowledge of information distribution, feedback did not alter communication patterns.

Distributed information systems are ubiquitous (Hutchins, 1991; Radner, 1986). In teams and organizations, individuals often have access to different sets of information because they have different roles, interests, or responsibilities. In some cases, differential access to information may arise simply because individuals are located in different places in a social, organizational or physical space. At critical junctures, impending collective action evokes attempts to assess and integrate information that is distributed across a social system. Broadly speaking, pooling distributed information entails a mix of two processes: communication and information integration. Social systems can be distinguished based on where information integration occurs. At one extreme, each member of the system may integrate the locally available information and communicate a decision recommendation (distributed decision-making). At the other extreme, members may simply communicate raw information to a centralized decision maker who must integrate the information (centralized decision-making).

Each approach to pooling information has advantages. For example, distributed decision making potentially eases the load on communication channels and decision makers, especially when time is limited. Each member can integrate a subset of the pertinent information and communicate only a recommendation. Moreover, information integration can occur in parallel, theoretically reducing the information processing time for the team. However, effective distributed decision making requires distributed expertise. That is, individual members need to have the skills, knowledge and experience to integrate effectively information. Moreover, distributed decision making risks collective action that is not supported by the entire profile of information (Stasser, 1988). That is, each set of localized information may support one action whereas the total information supports another action.

Centralized decision-making allows peripheral team members to concentrate on information search and communication and often permits the centralized decision-maker to base a decision on a more comprehensive set of information than is available to any of the peripheral members. Thus, decision-making expertise can be concentrated in one person and information gathering shared among team members. However, in the face of limited time and much information, the communication channels in such a system may be quickly overloaded and some information may be lost in transmission while other information is repeatedly conveyed to the decision-maker from multiple sources. Indeed, effective and efficient communication in such a system requires coordination in order to deliver a comprehensive set of information to the central decision-maker in a timely fashion.

The present study focused on the development of coordinated communication in seven-person teams that are hierarchically organized with a central decision-maker. Each team completed 32 episodes of a military situation assessment scenario. During each episode, the team had to make a binary decision (fire or not

fire) within a limited amount of time. Twenty items of diagnostic information were distributed among the team members such that four items were available to all members (shared), another eight items were available to two members (partially-shared), and the remaining eight items were available to only one member (unshared). The central decision-maker could not directly assess the 16 items of partially-shared and unshared information and had to rely on others in the team to access and communicate these items.

Of primary interest was the coordination of team communication over repeated episodes. The distribution of access to information resulted in a structural imbalance such that, other things being equal, partially shared information had more opportunity to reach the decision-maker than did unshared information. This simulated structural imbalance was intended to mimic situations in which some information is more widely available, either by design or happenstance, than other information. In these situations, opportunities to communicate information covaries positively with the number of team members who can access the information (Larson, Foster-Fishman, & Keys, 1994). Teams can potentially counteract such structural imbalance by coordinating their information search and communication activities and giving higher processing priority to items that are less widely disseminated.

Team coordination often depends on individual members anticipating what others are going to do and adjusting their own behavior accordingly. Predicting others' actions depends partly on knowing their capabilities (what can they do given their abilities and access to resources) and partly on knowing what they will do within the domain of possible actions (Wittenbaum, Stasser & Merry, in press). In intact teams that have repeated experience working together, one salient source of information about others' likely actions is past behavior. The present research examines two general propositions. First, team coordination is facilitated by the knowing how task-relevant resources are distributed among members (knowing what can others do). Second, the development of coordination over time is dependent on the clarity and immediacy of feedback about others' actions (knowing what others have done).

Information Distribution and Communication

The effect of information distribution on communication has been studied extensively in face-to-face discussions (Engel, 1992; Gigone & Hastie, 1993; Larson, et al., 1994; Stasser & Titus, 1985, 1987; Stasser, Taylor, & Hanna, 1989; Stasser & Stewart, 1992; Stasser, Stewart & Wittenbaum, 1995; Stewart & Stasser, 1995). The general finding is that groups are more likely to discuss information that is available to all members before discussion (shared information) than to discuss information that only one member can access (unshared information). For example, Stasser et al. (1989) gave university students descriptions of three candidates for student body president. These descriptions contained position statements

about campus and community issues. The descriptions were constructed so that some information (shared) was read by all members of a group before discussion whereas other information (unshared) was read by only one member. Participants then discussed the candidates in either three- or six-person groups and decided which candidate was best suited to the position. Discussions contained, on average, 46% of shared items but only 18% of the unshared items. This difference was greater for six-person groups than for three-person groups. Additionally, a procedural intervention designed to increase the amount of information discussed also increase disparity between percentages of shared and unshared items mentioned. For example, six-person groups whose discussions were structured mentioned 67% of the shared information but only 23% of the unshared information.

Stasser and Titus (1985, 1987) proposed that a collective information sampling (CIS) model provides one way of understanding why groups often omit unshared information from their discussions. Based on earlier formulations in the group performance literature (Lorge & Solomon, 1962; Shiflett, 1979; Steiner, 1972), the model represents the mentioning of an item of information during discussion as a disjunctive task: the group will fail to discuss an item only if all members fail to mention it. More specifically, the probability that the group will discuss an item, $p(D)$, is a function of the number of members who can mention it, n , and the likelihood that any one of these members will mention it, $p(M)$:

$$p(D) = 1 - [1 - p(M)]^n. \quad (1)$$

For example, if each member of 4-person group mentions 40% of the information that he or she received before discussion, the model predicts that the probability of discussing shared information is $p(D_s) = 1 - [1 - .4]^4 = .87$. In contrast, a piece of unshared information can by definition be mentioned only by the member who received it: $p(D_u) = 1 - [1 - .4]^1 = .4$. Thus, even though shared and unshared items are equally likely to be mentioned by each individual, shared information in this example has a sampling advantage at the level of the group.

The foregoing analyses are based on the assumptions that individual members are equally likely to mention shared and unshared information [i.e., $p(M_s) = p(M_u)$] and that recall among members is functionally independent. The model's predictions have been reasonably accurate when members cannot distinguish *a priori* shared from unshared information and cannot coordinate their information processing activities (e.g., as in Engel, 1992; Parks, 1991; Stasser and Titus, 1985, 1987; Stasser et al, 1989).

These assumptions are likely to be violated when members develop impressions of how information access is distributed among members. For example, Stewart and Stasser (1995) found that members of a decision-making group were more likely to mention unshared, than shared, information when they were aware that they had access to more items in a category of information than did other members (see also, Stasser, Stewart, and Wittenbaum, 1995). They interpreted their results in terms of

Wegner's theory of transactive memory (Wegner, 1986; Wegner, Erber, & Raymond, 1991). The theory posits that members of a team develop a coordinated system of encoding, storing, retrieving, and communicating information. The responsibility for domains of information is assigned, often implicitly, based on two types of considerations. First, members can be assigned responsibility for items that are associated with real or perceived expertise, where expertise is defined broadly to include not only skills and abilities but also interests and preferences. Second, and more germane to the present work, *circumstantial knowledge responsibility* depends on the environmental dynamics that determine how information is encountered by the team. For example, a member who first receives information may be implicitly assigned the responsibility for similar information in the future. Similarly, members who have more opportunity to access one type of information may be assigned responsibility for monitoring and remembering it.

An effective transactive memory system depends on a shared representation of how responsibility of members is distributed over relevant categories of information. For example, assignment of responsibility for items in category A to member X must be accompanied by an acceptance of that assignment by X. If X does not accept the responsibility, critical items of information in A may be unavailable to the team. Moreover, two or more members acceptance of responsibility for items in a category may result in unnecessary duplication. Under high information loads, such duplication may result in inadequate attention to other categories of critical information.

Although transactive memory theory allows for the possibility that assignments of responsibility can be explicit, many of the examples and applications of the theory suggest that the transactive memory systems often develop based on unspoken assumptions about what others will and should do. That is, members tacitly coordinate their information processing activities (Wittenbaum, et al., in press). The suggestion that coordination of information processing often develops tacitly is consistent with several ideas emerging from the group performance literature. For example, expectation states theory asserts that power structure emerge from unspoken assumptions (expectations) about others' abilities based on prior experience, diffuse status cues (e.g., gender, appearance, etc.), and early behavioral exchanges (Berger, Rosenholtz & Zeldich, 1980; Ridgeway, 1984). Once the power structure is formed, it affects the dynamics of the group although it is never explicitly discussed nor endorsed by the group. Similarly, performance strategies and interaction patterns often develop quickly and with little or no discussion in task-oriented groups (Hackman & Morris, 1965; Gersick, 1988). Even when the situation permits members to communicate freely, much of their interaction, including coordination and task strategies, is shaped tacitly.

Overview and Hypotheses


Seven person teams completed 32 trials of a simulated military, situation assessment scenario. Teams were told that they were to play the roles of a ship's crew that was monitoring the approach of a plane which could be either "friend or foe". On each of the 32 trials, the team decided to fire or not fire on approaching plane within a limited amount of time. The communication links among team members simulated a hierarchical organization as depicted in Figure 1. In the present study, decision-making was centralized such that the Commander at the top of the hierarchy had to make the decision to fire or not fire. Access to 20 items of diagnostic information was distributed among the team members such that four items were available to all members. The remaining 16 items could be accessed only by members at the bottom of the hierarchy. Thus, the central decision-maker could not directly assess these 16 items and had to rely on others in the team to assess and communicate these items. Eight of these items were available to two members (partially-shared), and the remaining eight items were available to only one member (unshared). Of primary interest was the ability of teams to convey to the Commander the partially-shared and unshared items within the limited time available for the Commander to make a decision.

Two features of the decision environment were systematically varied. Members in half of the teams were told how access to the 20 items of diagnostic information was distributed among members (elaborated knowledge of information distribution), whereas each member of the remaining teams only knew which items of information she/he could and could not access (limited knowledge of information distribution). That is, under limited knowledge, a member did not know whether other members could also access items that he/she could access and did not know who or how many members could access items that were not available to her/him.

Additionally, half of the teams were given explicit communication feedback after each trial. This feedback told all members which items of information reached the Commander before the end of the trial and which items were redundantly communicated (i.e., the Commander received the information two or more times). The remaining teams did not received this communication feedback. All teams received performance feedback after each trial the indicated the correct decision and whether the Commander made the correct choice.

It was expected that teams with limited knowledge of information access would convey to the Commander more partially-shared than unshared information. This is, these teams would not have adequate knowledge to counteract effectively the structural bias due to partially-shared items having more opportunities than unshared items to enter and survive in the communication network. In contrast, teams with elaborated knowledge were expected to give high priority to accessing and communicating unshared information and, as a result, to convey more unshared than partially-shared information to the Commander. In the rubric of transactive memory theory, these teams had unambiguous

Figure 1. Communication links and distribution of information.



circumstantial cues for assigning the responsibility for unshared information, but the appropriate assignment of responsibility for partially-shared information was not obvious.

Communication feedback was expected to attenuate over trials these biases favoring unshared information under limited knowledge and partially-shared information under elaborated knowledge. Under limited knowledge of information distribution, it was expected that members would learn from the feedback that certain items within their domain of access were failing to reach the Commander and shift their attention to these items. Because the omitted information would be predominantly unshared information, this shift of attention due to feedback would tend increase the communication of unshared items over trials. Similarly, under elaborated knowledge of information distribution, members who received communication feedback would learn that certain partially-shared items were failing to reach the commander. That is, individual members could claim responsibility for omitted partially-shared information and relinquish responsibility for partially-shared items that were redundantly communicated to the Commander.

Method

Design and Subjects. The design can be represented as a 2 X 2 X 4 mixed factorial of information distribution knowledge (limited versus elaborated), communication feedback (available versus not), and trial blocks. Each trial block contained eight trials. A team participated in one of four conditions as defined by knowledge of information distribution and communication feedback. University students were recruited via announcements in the student paper, campus bulletin boards, and meetings of campus service organizations to participate for up to two hours. The announcements described the general nature of the research and the minimum amount of pay for participation (\$10.00). Three hundred and ninety-two students participated in 56 7-person teams. Fourteen teams participated in each of the four conditions created by the knowledge of information distribution and communication feedback factors.

Experimental Task. The experimental task simulated a situation which requires the decision-making team to choose between two courses of action within limited time. An incorrect choice resulted in monetary loss whereas a correct choice yielded monetary gain. Each trial block contained 8 trials: 5 of these were "normal" trials with time limits of 90 sec and payoffs set at the \$.05 level, 1 was a high payoff trial with a 90-sec time limit, 1 was a low payoff trial with a 60-sec time limit and 1 was a high payoff trial with a 60-sec time limit. The occurrence of high payoff and short trials within a block was randomly assigned with the restriction that the first trial was always a normal (low payoff/90-sec) trial. The task was implemented on a computer network with team members seated at separate work stations.

The simulation program emulated situational assessment in a military environment. Participants were told that they were to

play the roles of a ship's crew that was monitoring the approach of a plane which could be either "friend or foe". If the plane intended to bomb the ship, there was a limited time for the crew to launch a preemptive strike. On the one hand, shooting down a "friendly" plane was costly and to be avoided. On the other hand, failing to launch a preemptive strike on an enemy plane was also costly. The costs were represented as a loss of money that team members were paid for their participation.

It was explained that, as is typical of such situations, they would not be able to ascertain with certainty the actual intent of the approaching plane until it was too late to react. However, they would have information that allowed them to estimate the likelihood of an actual attack. For example, the plane's velocity, altitude, and heading could provide useful diagnostic information that either raised or lowered the judged likelihood of attack. Moreover, the ship's location (contested versus uncontested international waters), current weather conditions, recent intelligence reports, and the like, may provide additional indicators of the likelihood of hostile action by the approaching plane.

There were 20 items of information that were divided into three sets. Four items of common information were always accessible by all team members. Each of eight partially shared items was directly accessible by two of the four team members at the bottom of the communication hierarchy. Each of the eight unshared items could be obtained by only one member. Thus, each of the four lower members could access two partially shared items and two unshared items (plus the four items of common information).

The patterns of information values across trials were constructed randomly. For 4 trials in each block, the probability that each item had a value indicating "friendly" was .70. On the other 4 trials in each block, the probability of a "friendly" value was .30. The probability that the approaching plane was actually "friendly", $p(F)$, was determined by the overall pattern of information values:

$$p(F) = .5 + \sum .025V_i \quad (2)$$

where $i = 1, 2, 3, \dots, 20$ and $V = 1$ if the value of the i th is "friendly" and -1 , otherwise. Optimal team performance occurred when a team considered all of the information and decided to fire only when the number of items indicating "enemy" exceeded the number indicating "friendly". Because of the probabilistic nature of the task, optimal use of information leads to an expected performance of 70% correct decisions. Chance performance is 50% correct decisions.

Procedures. Participants were scheduled in seven-person teams. Upon arriving, they were assigned randomly to computer workstations; each workstation had a designated role (e.g., Ship Commander, Radar Captain, Air Position Radar, etc.) that was associated with a particular node in the communication network. Introductory materials described the features of the task and the communication network. The types and meaning of the diagnostic

information were reviewed thoroughly. In addition to the introductory instructions, each team member had a station template that depicted the communication structure, summarized the types and meaning of information, and indicated what items that member could and could not access. Examples of these templates for the Radio Operator position are given in Figures 2 and 3. Under limited knowledge of information distribution, each team member was given a template that indicated which items that he/she could and could not access. An example of this template for the Radio Operator position is given in Figure 2. Under elaborated knowledge of information distribution, the template additionally indicated who, if anyone, among other team members could also access information that was available to a member's position. Figure 3 displays this template for the Radio Operator position. After team members had finished the introduction to the task, they completed 4 practice trials and then had the opportunity to discuss with the experimenter any questions or concerns about the procedures, the task, or their participation in the study.

The team completed 32 trials that were divided into four, eight-trial blocks. A trial was terminated either by the Commander making a decision or by expiration of allotted time. If time expired before the Commander made a decision, the default decision was not to fire. (Teams were aware of this default.) After each trial, each member reported his/her individual decision and then received feedback about the Commander's decision and the correct decision. Accumulated earnings were also displayed. In the communication feedback conditions, members received an additional display after each trial that indicated which of the 20 items the Commander had received before the end of the trial. Also, this display indicated which items the Commander received more than once. In the no communication feedback conditions, this additional display was not presented.

Each trial began with a display of the available time for that trial; the shorter 60-sec trials were also signaled by a bell. The trial clock was updated continuously throughout the trial. Additionally, high payoff trials were accompanied by a flashing dollar sign that was displayed under the trial clock in the computer monitor.

During each trial, team members could take the following actions: assess the values of items to which they had direct access; retrieve information sent to them by other members; send information that they had either directly assessed or received from another member; review items already assessed or received from another member; and request specific items from another specified team member. Members could send, receive, and request information only through the operative communication links in the network. When a message was sent, the recipient was notified that his/her mail file contained a message and it remained in the mail cue until it was retrieved. Messages sent to node had to be retrieved in the order that they were sent. A request for information had to be acknowledged in that the request remained

Figure 2. Example of station template for the Air Position Radar station under limited knowledge of information distribution.





Figure 3. Example of station template for the Air Position Radar station under elaborated knowledge of information distribution.



on the screen and other actions were suspended until the "space bar" on the keyboard was pressed. However, the target of the request did not have to comply with the request (and could not comply if he or she had not already assessed or received the desired information).

The Commander had the additional option of making a decision. A decision could be made at any time and once it was made the trial was terminated. Each station was notified of this action when it occurred and all other actions were suspended at that point. However, they were not told what the decision was until after they had registered their individual decisions.

Between trial blocks, team members were also given the opportunity to send a message to any other team member. The recipients of these inter-block messages were not constrained by the communication links that were operative during the trials. At the end of the session, team members responded to two open-ended questions that were designed to obtain verbal reports of strategies that they had employed and strategies that they would recommend to others. The first question asked them to describe the strategy (or strategies) that they used to select information to send to other members. The second question asked them what advice they would give to members of a new team.

Results and Discussion

Information Communicated


The average number of partially-shared and unshared items reaching the Commander by the end of each trial are displayed in Figure 4 as a function of experimental condition and trial block. These data were analyzed in a mixed factorial analysis of variance of the following factors: knowledge of information distribution (limited versus elaborated), communication feedback (available versus not), information type (unshared versus partially-shared), and trial blocks. The information type and trial blocks factors were repeated measures.

The only significant main effect was for trial block, $F(3, 144) = 141.52$, $p < .0001$. As is suggested by the pattern of means in Figure 4, there was a strong linear trend over blocks which is not only statistically significant, $F(1, 144) = 412.71$, $p < .0001$, but also accounts for 97% of the variance due to the block effect. Inspection of the trend in Figure 4 reveals that the linear trend is consistently positive. This pattern indicates increasing communication proficiency as the team gained experience with the task.

As predicted, the interaction of distribution knowledge and information type was significant, $F(1, 48) = 46.83$, $p < .0001$. Under limited knowledge, more partially-shared items ($M = 2.59$) than unshared items ($M = 1.82$) reached the Commander on each trial, $F(1, 48) = 26.56$, $p < .001$. Under elaborated knowledge, the reverse pattern obtained; more unshared items ($M = 2.99$) than partially-shared items ($M = 2.32$) reached the Commander, $F(1, 48) = 20.46$, $p < .001$.

The interaction of information type and trial block was significant, $F(3, 144) = 9.81$, $p < .0001$. The information type

Figure 4. Mean number of unshared (UNSH) and partially shared (SHAR) items reaching the Commander by experimental condition and trial block.



by linear trend across trial blocks accounts for 98% of the variance due to the interaction and is significant, $F(1, 144) = 13.51$, $p < .0001$. The increase across trials in number of items reaching the Commander was greater for partially-shared, than for, unshared information. This two-way interaction was qualified by a significant three-way interaction of information type, communication feedback, and trial block, $F(3, 144) = 4.11$, $p < .01$. The two-way interaction of information type by linear trend across trial blocks was significant for communication feedback conditions, $F(1, 144) = 6.96$, $p < .01$, but was not significant for no feedback conditions, $F(1, 144) = 2.60$, $p > .1$. That is, the differential increase in communication favoring partially-shared information was localized in the communication feedback conditions.

Whereas these significant interactions involving linear trends across trial blocks characterize the adjustments in communication over time, they do not directly test the *a priori* hypothesis that feedback would attenuate initial biases favoring one type of information over the other. To assess this hypothesis, we constructed two *a priori* contrasts: one for the limited knowledge conditions where communication was biased in favor of partially-shared information and one for the elaborated knowledge conditions where the reverse bias obtained. In both cases, the contrast was constructed to test whether the linear decrease in the bias over trials was greater when communication feedback was available.

The assertion that communication feedback would attenuate, over trials, biases favoring one type of information over the other was partially supported. Communication feedback attenuated the bias favoring unshared information in the elaborated knowledge conditions as indicated by a three-way interaction of information type, communication feedback, and linear trend across trial blocks within the elaborated knowledge conditions, $F(1, 144) = 7.38$, $p < .01$. The linear reduction in bias favoring unshared information under elaborated knowledge was significant when communication feedback was available, $F(1, 144) = 17.02$, $p < .0001$, but was not significant when communication feedback was absent, $F(1, 144) = 0.08$, *ns*. Communication feedback did not significantly attenuate the bias favoring partially shared information under limited knowledge, $F(1, 144) = 0.058$, *ns*. Indeed, there was a significant linear increase over trial blocks in the disparity between partially shared and unshared items reaching the Commander both when communication feedback was present, $F(1, 144) = 11.03$, $p < .0001$, and when it was absent, $F(1, 144) = 8.89$, $p < .001$. In short, communication feedback was useful in removing communication biases only when team members had elaborated knowledge of the distribution of information access.

In summary, teams that had limited knowledge of how information access was distributed across members communicated more partially shared than unshared information to the Commander. This finding is consistent with earlier findings for face-to-face

decision-making teams (e.g., Larson, et al., 1994; Stasser, et al., 1989) and can be attributed a structural bias in collective information sampling: partially shared items have more opportunity to be sampled and communicated than do unshared items. We reasoned that communication feedback would counter this bias over repeated episodes by alerting members to systematic omissions in the information reaching the Commander. However, there is no evidence that teams were able to use effectively the feedback in the absence of detailed knowledge of how information access was distributed among the members.

In contrast, teams that had elaborated knowledge of information access communicated to the Commander more unshared than partially shared items. Thus, members apparently gave higher priority to processing information to which they had unique access in order to counter the structural bias favoring partially shared information. Moreover, communication feedback apparently permitted members to adjust their information processing over time to equalize the chances of unshared and partially shared information reaching the Commander.

If members were selectively focusing on unshared items under elaborated knowledge, this focus could have impacted their selection of items to communicate and their selection of information to access initially.

Information Sent

On average, the members of a team sent 12.2 items of information to another member in a trial. Reflecting the overall structural bias in the task, partially shared information was sent more often ($M = 13.3$) than unshared information ($M = 11.0$), $F(1, 48) = 23.16$, $p < .0001$. This bias favoring partially shared information was more pronounced in the later than in the earlier trials. The number of items sent increased steadily across trial blocks ($M = 7.7, 11.4, 13.8$, and 15.8 for blocks 1, 2, 3, and 4, respectively), $F(1, 144) = 183.87$, $p < .0001$. The increase was greater for partially shared ($M = 7.7, 12.3, 15.4$, and 17.9 for blocks 1, 2, 3, and 4, respectively) than for unshared information ($M = 7.7, 10.5, 12.1$, and 13.8 for blocks 1, 2, 3, and 4, respectively), $F(1, 144) = 25.90$, $p < .0001$.

This overall pattern of communication favoring partially shared information was substantially muted in the elaborated knowledge conditions giving rise to a significant interaction of distribution knowledge and information type, $F(1, 48) = 69.68$, $p < .0001$. In the limited knowledge conditions, partially shared information was sent much more often ($M = 15.0$) than was unshared information ($M = 8.8$), $F(1, 48) = 31.31$, $p < .0001$. In contrast, there was a nonsignificant trend to send more unshared ($M = 13.3$) than partially shared ($M = 11.6$) in the elaborated knowledge conditions, $F(1, 48) = 2.32$, $p > .1$.

Information Assessed

The four members at the bottom of the hierarchy could access and assess the values of two unshared and two partially shared information. Thus, if every member had assessed all of the items available to him or her, the team would have

collectively accessed unshared items eight times and partially shared items 16 times. On average, 11.1 items were assessed per trial. Number of assessments increased over trial blocks ($M = 9.4, 11.0, 11.8, \text{ and } 12.2$, for blocks 1, 2, 3, and 4, respectively), $F(1, 144) = 94.13, p < .0001$. Nonetheless, individual members apparently never assessed all of the items that were available.

There was a strong information type effect reflecting partly the structural bias favoring partially shared information, $F(1, 48) = 616.85, p < .0001$. Partially shared items were assessed 14.1 times per trial whereas unshared items were assessed 8.1 times. This main effect was qualified by an information type by distribution knowledge interaction, $F(1, 48) = 13.37, p < .001$. Teams with elaborated knowledge assessed unshared items ($M = 8.6$) more often than teams with limited knowledge ($M = 7.5$), $F(1, 48) = 11.70, p < .001$. There was a marginally significant reverse trend of teams with elaborated knowledge assessing partially shared items ($M = 13.8$) less frequently than teams with limited knowledge ($M = 14.4$), $F(1, 48) = 3.06, p < .1$.

In summary, detailed knowledge of information distribution affected both what information was acquired (assessed) by members at the bottom of the hierarchy and what information was communicated to others higher in the hierarchy. Noticeably absent are effects of communication feedback on members' acquisitions of information and selections of items to communicate. This absence of feedback effects on members' actions is somewhat puzzling given the modulating effect of feedback on the amounts of unshared and partially shared information reaching the Commander in the elaborated knowledge condition. One implication is that this modulating effect is not attributable to a simple adjustment of processing priority from unshared to partially shared information over time. Rather members were apparently using the feedback in this condition in more sophisticated ways, perhaps tacitly negotiating responsibility for specific partially shared items while maintaining a high priority for processing unshared items.

Communication Feedback

Feedback about a team's past behavior was identified as one way that member's could tacitly coordinate their actions. By knowing what teammates did and did not accomplish in the past, each member could potentially shift attention away from overstaffed to understaffed activities. Whereas teams in this study who had elaborated knowledge of information distribution were apparently able to use communication feedback in this way, the communication of teams who did not have detailed knowledge of information distribution was seemingly not affected by communication feedback. One possibility is that these teams simply ignored the feedback. Another possibility is that they were unable to effectively use the feedback. The use of queries suggest that they did not ignore the feedback.

Members could request specific information from particular teammates via queries. Although this action was one means of

explicitly affecting communication patterns, it was very inefficient in that it required a series of actions to obtain the desired item. To be successful a query had to be sent, the receiver had to acknowledge the query, and then comply by sending the desired item. If the target of a query did not have the desired item, he/she either had to access it if possible or to query another member for the item in order to comply with the request. Thus, without extensive knowledge of how information access was distributed across the team, many queries were likely to be unsuccessful.

Number of queries per trial were analyzed in a mixed factorial analysis of variance of the following factors: knowledge of information distribution (limited versus elaborated), communication feedback (available versus not), and trial blocks. The interaction of distribution knowledge and communication feedback was significant, $F(1, 48) = 4.21, p < .05$. The number of queries in the limited knowledge/feedback condition ($M = 10.5$) exceeded substantially the number in the other conditions ($M = 6.8$); the associated contrast is significant, $F(1, 48) = 8.26, p < .01$.

There is also a significant block effect, $F(3, 144) = 16.55, p < .0001$. The linear component of this block effect is highly significant, $F(1, 144) = 46.19, p < .0001$, and accounts for 92% variance due to the overall block effect. Neither the quadratic nor cubic component is significant, $F(1, 144) = 3.25, p > .05$ and $F(1, 144) = 0.59, ns$, respectively. The increasing use of queries over blocks is larger in the limited knowledge/feedback condition than in the other conditions. A Scheffe post hoc test of the contrast of the linear trend in the limited knowledge/feedback with the trend in the other conditions confirms this impression, $F(1, 144) = 8.21, p < .05$.

In sum, teams who received communication feedback but had limited knowledge of information distribution sent an unusually large number of queries, particularly during the later trial blocks.

Decision Performance

Because the task was designed so that partially-shared and unshared information was equally diagnostic, communication biases that favored either type of information should not have systematically affected performance in this study. Nonetheless, teams that were able to get information to the Commander quickly should have enjoyed a slight advantage.

The proportion of correct decisions is bounded, given the probabilistic nature of the task, by .5 (chance performance) and .7 (highest expected performance). Overall, teams made the correct decision on 60.2% of the trials. Teams with elaborated knowledge of information access were correct slightly more often ($M = .628$) than teams with limited knowledge ($M = .576$), $F(1, 52) = 5.24, p < .03$. This effect was most evident in the final trial block during which teams with elaborated knowledge were right 65.6% of the time and teams with limited knowledge were right 56.2% of the time, $F(1, 52) = 4.18, p < .05$.

Summary and Conclusions

The results support the predictions regarding the effects of information distribution knowledge on communication of partially-shared and unshared information. When the information distribution was not known, partially-shared information was more likely than unshared information to reach the Commander. When the information distribution was known, the pattern was reversed in that unshared items were more likely to reach the Commander than were partially-shared items. However, the availability of feedback virtually eliminated this bias in favor of unshared information by the fourth trial block. Thus, feedback was efficacious when members also knew who could access what information. Without this knowledge, feedback had little systematic effect.

Without the benefit of knowing how information access was distributed among members, teams were seemingly unable to implement a strategy that allowed them to address the disproportionate omission of unshared items. However, there is indirect evidence that they detected a problem. In the feedback without distribution knowledge condition, requests from one team member to another for specific items of information increased by threefold over the trial blocks and the number of these requests was 60% higher in trial blocks 3 and 4 than in the other conditions.

Several applied implications follow. First, when information is not uniformly available to all members, increasing the number of members who have access to particular items will increase the likelihood that they will be successfully transmitted through a communication network. Thus, consistent with intuition, information that is thought to be relatively more important should be immediately accessible to more team members. Second, however, if team members are aware of who, and how many, can access particular types of information, this recommendation may change. Under these conditions, each member is more likely to access and communicate information that is uniquely available to him or her. The implication is that the more people who are thought to have access to an item, the less responsibility any one of them will accept for accessing and communicating the item. Finally, feedback (debriefing) regarding the information actually reaching the decision-maker is not sufficient by itself to ensure that teams can address the communication biases introduced by the patterns of information access across team members, but feedback coupled with awareness of who has access to what information does permit teams to correct communication biases with experience.

References

- Berger, J., Rosenholtz, S. J., and Zeldich, M., Jr. (1980). Status organizing processes. Annual Review of Sociology, 6, 479-508.
- Engel, B. T. (May 1992). Group structure and the sharing of information during decision making. Paper presented at the meeting of the Midwestern Psychological Association, Chicago, IL.
- Gersick, C. J. G. (1988). Time and transition in work teams: toward a new model of group development. Academy of Management Journal, 31, 9-41.
- Gigone, D., and Hastie, R. (1993). The common knowledge effect: Information sampling and group judgment. Journal of Personality and Social Psychology, 65, 959-974.
- Hackman, J. R., and Morris, C. G. (1975). Group tasks, group interaction process and group performance effectiveness: A review and proposed integration. In L. Berkowitz (Ed.), Advances in experimental social psychology (Vol. 8, pp. 45-59). New York: Academic Press.
- Hutchins, E. (1991). The social organization of distributed cognition. In L. B. Resnick, J. M. Levine, and S. D. Teasley (Eds.), Perspectives on socially shared cognition (pp. 382-307). Washington, D.C.: American Psychological Association.
- Larson, J. R., Jr., Foster-Fishman, P. G., and Keys, C. B. (1994). Discussion of shared and unshared information in decision-making groups. Journal of Personality and Social Psychology, 67, 446-461.
- Lorge, I., and Solomon, H. (1962). Group and individual behavior in free-recall verbal learning. In J. H. Criswell, H. Solomon, and P. Suppes (Eds.), Mathematical methods in small group processes (pp. 221-231). Stanford, CA: Stanford University Press.
- Parks, C. D. (1991). Effects of decision rule and task importance on willingness to share "unique" information. (Doctoral dissertation, University of Illinois, 1991). Dissertation Abstracts International, 52 (07), B3948.
- Radner, R. (1986). Information pooling and decentralized decision making. In B. Grofman and G. Owens (Eds.), Information pooling and group decision making (pp. 195-217). Greenwich, CT: JAI Press.
- Ridgeway, C. L. (1984). Dominance, performance and status in groups: A theoretical analysis. In E. J. Lawler (Ed.), Advances in group processes (Vol. 1, pp. 59-93). Greenwich, CT: JAI Press.
- Shiflett, S. C. (1979). Toward a general model of small group productivity. Psychological Bulletin, 86, 67-79.
- Stasser, G. (1988). Computer simulation as a research tool: The DISCUSS model of group decision making. Journal of Experimental Social Psychology, 24, 393-422.
- Stasser, G., Taylor, L. A., & Hanna, C. (1989). Information sampling in structured and unstructured discussions of

- three- and six-person groups. Journal of Personality and Social Psychology, 57, 67-78.
- Stasser, G., & Stewart, D. (1992). Discovery of hidden profiles by decision-making groups: Solving a problem versus making a judgment. Journal of Personality and Social Psychology, 63, 426-434.
- Stasser, G., Stewart, D. D., & Wittenbaum, G. M. (1995). Expert roles and information exchange during discussion: The importance of knowing who knows what. Journal of Experimental Social Psychology, 31, 244-265.
- Stasser, G., & Titus, W. (1987). Effects of information load and percentage of shared information on the dissemination of unshared information during group discussion. Journal of Personality and Social Psychology, 53, 81-93.
- Stasser, G., & Titus, W. (1985). Pooling of unshared information in group decision making: Biased information sampling during discussion. Journal of Personality and Social Psychology, 48, 1467-1478.
- Steiner, I. D. (1972). Group process and productivity. New York: Academic Press.
- Stewart, D. D., & Stasser, G. (1995). Expert role assignment and information sampling during collective recall and decision making. Journal of Personality and Social Psychology, 69, 619-628.
- Wegner, D. M. (1986). Transactive memory: A contemporary analysis of the group mind. In B. Mullen and G. Goethals (Eds.). Theories of group behavior (pp. 185-208). New York: Springer-Verlag.
- Wegner, D. M., Erber, R., and Raymond, P. (1991). Transactive memory in close relationships. Journal of Personality and Social Psychology, 61, 923-929.
- Wittenbaum, G. M., Stasser, G., & Merry, C. J. (in press). Tacit coordination in anticipation of small group task completion. Journal of Experimental Social Psychology.