9-29-2020

Distributed Cognition in Teams Is Influenced by Type of Task and Nature of Member Interactions

R. S. Tindale
Loyola University Chicago, rtindal@luc.edu

Jeremy R. Winget
Loyola University Chicago

Verlin B. Hinsz
North Dakota State University

Author Manuscript

This is a pre-publication author manuscript of the final, published article.

Task Type, Member Interaction, and Distributed Cognition in Teams

R. Scott Tindale & Jeremy R. Winget

Loyola University Chicago

Verlin B. Hinsz

North Dakota State University

Draft of chapter to appear in …

Abstract

In contemporary organizations, many if not most teams work on cognitive or information processing tasks (Hinsz, Tindale, & Vollrath, 1997). The past 50 years of research has taught us much about how information is accessed, created, attended to, and processed as groups attempt to complete various tasks. However, many of the information processing effects that have been observed are at least somewhat task specific, yet little research has focused specifically on tasks and how their information processing requirements differ. In this chapter, we discuss how task differences can impact how groups use and process information and how different information distribution patterns across member might impact performance. In addition, we address how constraints on the amount and type of interactions among the team members will affect performance in different task domains. We hope our discussion will demonstrate the importance of task differences for understanding team information processing and point out areas where greater research focus is needed.

The study of behavior in and by groups has a long history in the social and behavioral sciences (Triplett, 1897; Sherif, 1932; Lorge & Solomon, 1955). Due largely to military funding after WWII, many researchers began studying how groups performed in a variety of different contexts and under varying conditions. This work was integrated and summarized in a landmark volume by Steiner (1972). Two of Steiner's main conclusions were that groups rarely perform up to their full potential and that both the actual and potential performance of groups was heavily influenced by the type of task on which they worked. These conclusions remain relevant to more recent attempts at theory and research on group performance. Steiner also focused on issues of both coordination and motivation in explaining group performance, which are also still present in current work in this area. However, the term "cognition" was rarely if ever used to describe the work that Steiner reviewed.

Shortly after Steiner's review, the types of tasks on which groups worked began to change from mostly physical tasks to more cognitive or information processing tasks. In addition, the social and behavioral sciences were being swept up in the "cognitive revolution" (Lockman, Lockman, and Butterfield, 1976; Newell and Simon, 1972). Individuals (and later groups) began to be seen as information processing systems and the theories that guided performance research began to focus both on information and the ways in which it was processed. Several review articles on groups published in the 1990s began to reflect this shift in emphasis (Hinsz, Tindale, & Vollrath, 1997; Larson & Christianson, 1993; Thompson & Fine, 1999). Much of the research from this period focused on how information that was distributed among the group members was processed or used by the group (Stasser & Titus, 1985; 1987). Consistent with one of Steiner's (1972) main conclusions, groups rarely performed up to their potential as information processors. However, the focus on information tended to overshadow

Steiner's emphases on other concepts like motivation and coordination. Research since the 1990s has begun to reintroduce such notions into models on group information processing (De Dreu, Nijstad, & van Kippenberg, 2008; Abele & Stasser, 2010).

Another recent trend in group research has focused on how simply aggregating individual judgments (sans interaction or communication) can lead to quite accurate group judgments. The power of groups (or more colloquially, the "Wisdom of Crowds", Suroweiki, 2004) has led to the use of "big data" to help organizations make several different types of decisions (Tetlock, 2015). However, this has led some researchers to argue that group interaction is superfluous or even detrimental, to group performance (Armstrong, 2006). This would imply that using groups to gain information is useful, but the processing of such information should be done elsewhere.

Our goal in the present chapter is to review and integrate research on groups as cognitive or information processing systems, taking into account how such processes function for or are affected by different types of tasks and different amounts of group member interaction. It will be a targeted review, attempting to highlight key task features and key aspects of group member interaction (or lack thereof). One of the key task distinctions we will make draws from Steiner's (1972) distinction between unitary and divisible tasks. Unitary tasks are those where all group members are basically working on the same task together. For example, a group of programmers trying to fix a bug in a computer program would be a unitary task. Divisible tasks are those where each group member (or different subgroups of members) are working on ostensibly different tasks that when combined with the tasks performed by other members (or subgroups) will lead to some collective goal. An example of a divisible task world be an organization launching a new product, but some of the team members are working on marketing the product, others are working on productions, and others still are working on staffing, etc. This distinction is in some senses arbitrary since many unitary tasks could be broken up into subtasks. However, the role of information exchange and interaction is different depending on how easily divisible the task tends to be. Finally, we will attempt to use current theory and research findings to make suggestions on how best to use groups as information processing systems.

Unitary Tasks

Simple Aggregation – No Interaction

Although the basic finding has been known since the 1950s (Steiner, 1972), Suroweiki (2004) brought the notion of the "wisdom of crowds" to the forefront of both science and popular culture. The basic idea behind the wisdom of crowds is that an aggregation of many individual judgments will tend to be more accurate than a randomly selected individual judgment and will often be more accurate than a single judgment from an expert. This phenomenon has been replicated many times in a variety of different judgment task domains (Larrick & Soll, 2006; Surowiecki, 2004). Ariely et al. (2000) showed that, assuming pairwise conditional independence and random individual error distributions (although rare in many decision contexts), the average of J probability estimates (J = the number of estimators) will always be better than any of the component individual estimates and that as J increases, the average will tend toward perfect calibration diagnosticity (accurate representation of the true state of affairs), even when information provided to the various estimators is less than optimal. In addition, Johnson et al. (2001) empirically showed the accuracy of the average probability estimate to be robust over several conditions, even when individual estimates were not independent. Recent work on forecasting has shown that a simple average of multiple independent forecasts will perform better than individual experts and often perform as well as more sophisticated aggregation techniques (Armstrong, 2001).

Larrick and Soll (2006) have explained the advantage of simple averages over individual judgments using the concept of "bracketing". If the group member judgments are independent, different members will make somewhat different estimates with some of the estimates above the "true score" and others below it. Thus, the estimates "bracket" the true score. When this is true, it can be mathematically shown that the average of the multiple estimates will always be more accurate than the average individual judge. If the true score is well bracketed by the multiple estimates (near the median or average), the aggregate accuracy will be far superior to the typical individual judge. However, even if the true score is closer to one of the tails of the distribution, the average will still outperform the typical individual, though not to the same degree. Larrick and Soll (2006) also show that even when the true score is not bracketed by the estimates, the group (average) will do no worse than the typical individual judge.

From an information processing perspective, bracketing is a function of distributed information. Different group members have different information about the particular judgment context. This different information access leads to judgments that vary as a function of that information. Assuming the generally available information is not biased toward a particular tail of the distribution, the judgments should randomly vary around the true score. Thus, the wisdom of crowds can be viewed as a function of the natural distribution of information across members.

Although central tendency aggregation models have been shown to do quite well in a number of situations (Larrick & Soll, 2006), a number of researchers have attempted to improve aggregate forecasts by modifying the aggregation procedure. Budescu and Chen (2013) formulated a method for improving group forecasts by eliminating group members whose forecasts detract from the group performance. They had group members make probabilistic forecasts for a variety of events and then assessed whether the group's forecast was better or worse when each group member was included in (or removed from) the aggregate. By only including those members whose forecasts showed a positive influence on accuracy, they consistently improved the accuracy of the group forecasts relative to the simple average and other less effective weighting schemes, and the improvements persisted for future judgments not used to define the inclusion criteria (see also Mellers et al., 2014).

Mannes, Soll, and Larrick (2014) suggest a select-crowd strategy, which ranks judges based on a cue to ability (e.g., the accuracy of several recent judgments) and averages the opinions of the top judges (e.g., the top five). Through both simulation and an analysis of 90 archival data sets, results show select crowds of five knowledgeable judges yield very accurate judgments across a wide range of possible settings—the strategy is both accurate and robust (Mannes et al., 2014). Following this, they examine how people prefer to use information from a crowd. The authors' findings demonstrate people are drawn to experts and dislike crowd averages, but importantly, they view the select-crowd strategy favorably and are willing to use it. The select-crowd strategy is accurate, robust, and appealing as a mechanism for helping individuals tap collective wisdom.

Aggregation with Limited Information Exchange

Although simple aggregation tends to produce fairly accurate decisions, there is little chance for members to share information or defend their positions. In addition, group members often remain unaware of others' positions and the final group product. Although there is evidence that often little is gained by member exchanges (Armstrong, 2006; Lorenz, Rauhutb, Schweitzera, & Helbing, 2011), it is difficult for members with insights or valuable information to have influence without some type of interchange among group members (Kerr & Tindale, 2011). Obviously full group deliberation (a topic discussed later) would allow for members to share and defend their position. However, there is evidence that the most influential members in freely interacting groups (based on status or confidence) are not always the most accurate or correct (Littlepage et al, 1997). Thus, various approaches at compromise procedures where some information exchange is allowed but pressures toward conformity and incidental influence would be minimized.

Probably the most famous of these procedures is the Delphi Technique (Dalkey, 1969; Rowe & Wright, 1999; 2001). This technique has been used in idea generation and forecasting most often, but it has also been adapted to other situations as well (Rohrbaugh, 1979). The procedure starts by having a group of (typically) experts make a series of estimates, rankings, idea lists, etc. on some topic of interest to the group or facilitator. The facilitator then compiles the list of member responses and summarizes them in a meaningful way (mean rank or probability estimate, list of ideas with generation frequencies, etc.). The summaries are given back to the group members and they are allowed to revise their initial estimates. The group members are typically anonymous, and the summaries do not specify which ideas or ratings came from each member. This procedure allows information from the group to be shared among the group members but avoids conformity pressure or undue influence by high status members. The procedure can be repeated as many times as seems warranted but is usually ended when few if any revisions are recorded. The final outcome can range from a frequency distribution of ideas to a choice for the most preferred outcome or the central tendency (mean or median) estimate. A number of related techniques (e.g., Nominal Group Technique, Van de Ven. A.H., & Delbecq, A. L., 1974) use similar procedures but vary in terms of how much information is shared and whether group members can communicate directly. Overall, the purpose of these procedures is to allow for some information exchange while holding potential distortions due to social influence in check. Research on the Delphi technique has tended to show positive outcomes. Delphi groups do better than single individuals and do at least as well as, if not better than, face-to-face groups (Rorhbaugh, 1979). They have also been found to work well in forecasting situations (Rowe & Wright, 1999; 2001).

A more recent technique is the use of prediction markets (cf. Wolfers & Zitzewitz, 2004). Much like financial markets, prediction markets use buyers' willingness to invest in alternative events (e.g., Obama will win vs. McCain will win in the 2008 US Presidential election) as a gauge of their likelihood. They typically do not prohibit direct communication among forecasters/investors/bettors, but in usual practice there is little such communication (if any). However, since the value placed on the assets is typically set in an open market of buyers and sellers, those already in (or out) of the markets can be informed and swayed by various market indicators (e.g., movements in prices, trading volume, volatility), and thus mutual social influence can occur through such channels. The simple "initial forecasts–group aggregation–final forecast" sequence does not really apply to this method very well; it is a much more dynamic and continuous aggregation process, where bids and offers can be made, accepted, and rejected by multiple parties, and the collective expectations of the "group" can continue to change right up to the occurrence of the event in question (e.g., an election). Except for those with ulterior motives (e.g., to manipulate the market, or to use the market as a form of insurance), investments in such markets are likely to reflect the investors' honest judgments about the relative likelihood of events. Members can use current market values to adjust their thinking and learn from the behavior of other members. However, such investment choices are not accompanied by any explanation or justification. Indeed, such investors may even have incentives to withhold vital information that would make other investors' choices more accurate (e.g., that might inflate the price of a "stock" one wants to accumulate). Thus, in terms of opportunities for mutual education and persuasion, prediction markets fall somewhere between statistical aggregation methods (which allow none) and face-to-face groups (which allow many).

There is now a growing body of evidence supporting of the accuracy of prediction markets (Forsythe, Nelson, Neumann, & Wright, 1992; Rothchild, 2009; Wolfers & Zitzewitz, 2004). They sometimes overestimate the likelihood of very rare events, but they have done extremely well at predicting presidential elections results in the USA over the past three election cycles. There is also experimental evidence that group members can learn from participating in market-type environments. Maciejovsky and Budescu (2007) had people participate in a competitive auction bidding for information in order to solve the Wason card task, which requires testing a hypothesis using evidence. Their results showed that participants were better at solving such problems (chose the appropriate evidence in an efficient manner) after having participated in the auctions. Thus, even with very minimal exchange, groups can be very accurate decision makers and their members can gain expertise during the process.

Vroom and Yetton (1968) argued that one of the ways managers make decisions is through consultation; the decision is made by the manager but only after getting advice from key members of the team or organization. Vroom and Yetton argued that consultation is optimal when managers do not have all the information at their disposal to make a good decision. Thus, they take advantage of the information that is distributed among the rest of the team members. Sniezek and Buckley (1995) referred to this mode of social decision making as the "Judge – Advisor" systems approach. The judge is responsible for the final decision but he/she seeks out suggestions from various advisors. Such systems have recently received a fair amount of research attention (see Banaccio & Dalal, 2006 for a review). Based on the research discussed above, unless the judge had far more expertise than an advisor, the judge should weight the advice equal to their own opinion. Although receiving advice usually does improve judges' decisions relative to when they receive no advice, a vast amount of research has shown that judges tend to weight their own opinions more than twice as much as the advice they receive (Larrick, Mannes, & Soll, 2012). This has been referred to as "egocentric advice discounting" (Yaniv, 2004; Yaniv & Kleinberger, 2000). This effect has been found to be extremely robust and has been replicated in many decision situations with several types of judges and advisors (Banaccio & Dalal, 2006).

Judges do take the expertise of the advisors into account when re-evaluating their position. Thus, judges discount less when the advisors are known experts, or their past advice has proved to be accurate (Goldsmith & Fitch, 1997). Judges are also more likely to use advice when making judgments in unfamiliar domains (Harvey & Fischer, 1997), and they learn to discount poor advice to a greater degree than good advice (Yaniv & Kleinberger, 2004). However, judges are not always accurate in their appraisals of advisor's expertise. Sniezek and Van Swol (2001) have shown that one of the best predictors of judges use of advice is advisor confidence, which is poorly correlated with advisor accuracy. Discounting has been found to be less for advice that is solicited by the judge as compared to advice simply provided (Gibbons, Sniezek, & Dalal, 2003). In addition, judges discount less when the task is complex (Schrah, Dalal, & Sniezek, 2006), when there are financial incentives for being accurate (Sniezek & Van Swol, 2001), and when they trust the advisor (Van Swol & Sniezek, 2005). However, discounting is present in virtually all JAS situations and it almost always reduces decision accuracy.

Much of the research on Judge-Advisor systems have only allowed advisors to provide judgments or judgments with confidence ratings (Banaccio & Dalal, 2006). This does not allow judges to hear arguments in support of particular positions or estimates. In addition, most of the judgment tasks used for this research (and for the simple aggregation research discussed previously), are what Laughlin (1980) would call "judgment", rather than "intellective" tasks. Judgment tasks do not allow group members to actually “demonstrate” the accuracy or correctness of their judgments. For intellective tasks, group members should be able to convince other members that their position is correct or most accurate using the information available to the group. However, to do this, they would need to be able to discuss and share information relevant to performance on the task. These limited interaction strategies tend not to allow for such interactions. Kerr and Tindale (2011) showed that limited interaction strategies will do well when the correct answer is bracketed by the group member preferences or when the correct solution is the most popular. But when correct solutions are only preferred by a minority of the members or are far from the mean or median member positions, such strategies lead to performance much below levels expected for interacting groups. It is toward fully interacting groups that our attention now turns.

Fully Interacting Groups

Most of the research on group decision making has focused on groups where the members meet face-to-face and discuss the decision problem until they reach consensus. Early research in this area tended to focus on member preferences as the major feature predicting group decision outcomes (Davis, 1973; Kameda, Tindale, & Davis, 2003). More recent research has focused on how groups process information (Hinsz, Tindale, & Vollrath, 1997) and the degree to which available information is used by the group (Brodbeck, Kerschreiter, Mojzisch, & Shulz-Hardt, 2007; Lu, Yuan, & McLeod, 2012). More recently still, the degree to which motivational aspects of groups and group members have begun to receive attention (De Dreu, Nijstad, & van Knippenberg, 2008). We will focus mainly on the two more recent areas in the sections below.

Group Information Processing. A popular approach to studying interacting groups working on a unitary task utilizes the hidden profile paradigm. This paradigm is marked by a biased pattern of information distribution in which, prior to group discussion, some information is common to all group members and other information is unique to individual members (Stasser & Titus, 1985). The common information favors a suboptimal decision alternative, whereas all the unique information combined reveals the optimal alternative. Ultimately, this “hides” the optimal decision choice from the group as a whole. It can only be discovered when each individual shares their unique information and the group uses this information to inform its decision.

Research on hidden profile tasks has shown groups generally do not exchange information efficiently and decision quality suffers as a result. A meta-analysis of 65 studies using the hidden profile paradigm (101 independent effects, 3,189 groups) showed (1) groups mention more pieces of common information than unique information; (2) hidden profile groups are less likely to find the solution than are groups having full information; (3) information pooling (i.e., percentage of unique information mentioned out of total available information, percentage of unique information out of total discussion) is positively related to decision quality; and communication medium (i.e., computer mediated communication vs. face-to-face) does not affect (4) unique information pooling or (5) group decision quality (Lu, Yuan, & McLeod, 2012). However, group size, total information load, the proportion of unique information, task demonstrability, and hidden profile strength moderated these effects.

Most of the current research findings have been nicely encapsulated by Brodbeck et al. (2007) in their Information Asymmetries Model of group decision making. The model categorizes the various conditions that lead to poor information processing in groups into three basic categories. The first category, negotiation focus, encompasses the various issues surrounding initial member preferences. If groups view the decision-making task mainly as a negotiation, members negotiating which alternative should be chosen tend to focus on alternatives and not on the information underlying them. The second category, discussion bias, encompasses those aspects of group discussion that tend to favor shared vs. unshared information (e.g., items shared by many members are more likely to be discussed). The third category, evaluation bias, encompasses the various positive perceptions associated with shared information (e.g., shared information is more valid, sharing shared information leads to positive evaluations by other group members). All three categories are good descriptions of typical group decision-making and can lead to biased group decisions and inhibit cross-fertilization of ideas and individual member learning (Brodbeck et al. 2007).

A key aspect of the Brodbeck et al. (2007) model is that the various aspects of information processing in typical groups only lead to negative outcomes when information is distributed asymmetrically across group members, as when a hidden profile is present. Although such situations do occur, and groups can make disastrous decisions under such circumstances (Janis, 1972; Messick, 2006), they are not typical of most group decision environments. In situations where members have independently gained their information through experience, the shared information they have is probably highly valid and more useful than unique information or beliefs held by only one member. Thus, the fact that members share preferences and information in many group decision contexts is probably adaptive and has generally served human survival well (Hastie & Kameda, 2005; Kameda & Tindale, 2006). In addition, groups are often (but not always) sensitive to cues in the environment that indicate that information is not symmetrically distributed (Brauner, Judd, & Jacquelin, 2001; Stewart & Stasser, 1998). Although minorities often are not very influential in groups, if minority members have at their disposal critical information that others do not have and that implies the initial group consensus may be wrong, other group members will pay attention to them.

However, such minority effects may only be realized when groups are (or think they are) working on Intellective tasks. Several studies have shown moderation effects of task demonstrability. Lu and colleagues (2012) found the likelihood of a manifest profile (i.e., all members have access to all information) over a hidden profile group finding the optimal task solution increased when working on tasks with high (vs. low) demonstrability (i.e., odds ratios of 15.18 vs. 2.46, respectively). Their results indicate hidden profile tasks without a clear preferred solution are most detrimental to information sharing and decision quality, whereas highly demonstrable tasks increase information sharing (Lu et al., 2012). These findings are consistent with other research showing information pooling is more predictive of decision quality (Mesmer-Magnus & DeChurch, 2009) and group discussions are less likely to focus on common information during high demonstrability tasks (Reimer, Reimer, & Czienskowski, 2010).

Other research provides converging evidence for these claims. Specifically, Laughlin, Bonner, and Minner (2002) had 82 four-person cooperative groups and 328 independent individuals solve a random coding of the letters A–J to the numbers 0–9. On each trial the group or individual proposed an equation in letters (e.g., A + D = ?), received the answer in letters (e.g., A + D = B), proposed one specific mapping (e.g., A = 3), received the answer (e.g., True, A = 3), and proposed the full mapping of the 10 letters to the 10 numbers. Researchers found groups needed fewer trials to find the solution, proposed more complex equations, and identified more letters per equation than each of the best, second-best, third-best, and fourth-best individuals. In this experiment, the nature of the task had a clear solution (i.e., it was intellective rather than judgmental), which required demonstrable recognition of correct answers, demonstrable rejection of erroneous answers, and multiple insights into effective collective information processing strategies.

According to the theory of combinations of contributions, the outcomes of group interaction on a task can be predicted by two components: the contributions and the combinations (Hinsz & Ladbury, 2012). The contributions refer to the inputs group members bring with them to the task situation (e.g., cognitive skills, processing goals, etc.). The combinations refer to the aggregation principle by which the contributions are combined to lead to the group outcomes (e.g., strategies to pool, share, and integrate information). Importantly, contributions and combinations directly relate to the cognitive processes involved in how group inputs result in team outcomes on a task (Hinsz & Ladbury, 2012).

Groups always exist in a context, and they are sensitive to this context. Thus, the combinatorial rule that summarizes the processes by which inputs are transformed into outcomes in dependent on the context as well. One of the key findings concerning how groups process information is the common knowledge effect; that is, information shared by many group members plays a larger role in group process and performance than unshared information (Stasser & Titus, 1985). Given this finding, it seems that to increase the amount of information sharing within a group, all group members should have access to all the information available. Indeed, despite the benefits of such manifest profile groups (e.g., Lu et al., 2012), in such information-rich environments, assigning all information to all members may overload each member's cognitive capabilities.

Tindale and Sheffey (2002) examined ways to optimally assign information to group members. Following a model proposed by Zajonc and Smoke (1959), the researchers assessed the effects of information assignment redundancy and group interaction on group memory performance. Participants in five-person groups either received a full list of consonant-verb-consonant non-word trigrams to memorize, or a partial list with each trigram distributed to two group members. Groups recalled trigrams as either coacting or interacting groups. In terms of correct recall, coacting groups outperformed interacting groups, and partial redundancy produced better recall than total redundancy. However, intrusion errors were greatly reduced by group interaction and/or a reduction in the cognitive load on the individual group members (i.e., partial redundancy). Groups in the partial redundancy condition tended to perform near optimal levels.

Motivation in groups has been a topic of interest in social psychology since its earliest days as a field of inquiry (Triplett, 1898). Many studies have focused on how groups affect the amount of effort expended by their members, and both motivation gains and losses have been demonstrated (Kerr & Tindale, 2004; Weber & Hertel, 2007). Motivation has also been an important topic in group as well as individual, decision making, and until recently the basic motivational assumption was hedonism. Many models of collective decision-making use basic game theoretic, or utility maximization, principles to explain how members both choose initial preferences and move toward consensus (Kahn & Rapoport, 1984). Thus, much of the early work on group decision making tended to treat individual group members as players in a utility maximization game (Budescu, Erev, & Zwick, 1999). Game theory approaches are quite prevalent and also quite useful for understanding social behavior (Kameda & Tindale, 2006), but other motives more associated with the group level of analysis have also been found to be important (Levine & Kerr, 2007). In addition, many of these motivations were discovered because social behavior did not follow game theoretic expectations (Dawes, van de Kragt, & Orbell, 1988).

Probably the most heavily researched of these more recent motives in groups involves the ingroup bias (Hogg & Abrams, 1988). There is now substantial evidence that when group members think about themselves as a group (thus they share a social identity), they begin to behave in ways that protect the group from harm or enhance its overall welfare. Many of the implications of this bias are positive for the group, but there are situations where it prevents groups from making good decisions. For example, groups are more likely than individual to lie about preferences and resources in a negotiation setting (Stawiski, Tindale, & Dykema-Engblade, 2009). Probably the most prominent example where protecting or enhancing the group's welfare leads to less than optimal decisions is the inter-individual – intergroup discontinuity effect (Wilshulz et al., 2003). This effect was initially demonstrated by McCallum et al. (1985) who compared individuals to groups when playing a prisoner's dilemma game. The prisoner's dilemma game is a mixed motive game where the dominant or individually rational response is not to cooperate with the other player. However, when both players make the non-cooperative choice, they both do poorly. The only collectively rational choice is for both players to cooperate, which leads to the greatest collective payoff and to moderate positive gains for each player. When two individuals play the game and can discuss the game before making choices, they both end up cooperating better than 80% of the time. However, when two groups play the game and each group must choose between cooperation and non-cooperation, groups quite often choose not to cooperate. Over multiple plays of the game, groups end up locked in the mutual non-cooperation payoff and earn far worse payoffs compared to the inter-individual situation. This effect has been replicated many times using several types of mixed motive game structures and different sized groups (see Wilshutz et al., 2003 for a review).

However, giving he groups the right motivation can help groups to be better information processors. De Dreu, Nijstad, and van Knippenberg (2008) developed a model of group judgment and decision-making based on the combination of epistemic and social motives. Called the "motivated information processing in groups" model (MIP-G), the model argues that information processing in groups is better understood by incorporating two somewhat orthogonal motives; high vs. low epistemic motivation, and pro-social vs. pro-self motivation. Earlier work on negotiation had shown that negotiators that share both high epistemic motivation and a pro-social orientation were better able to find mutually beneficial tradeoffs and reach better integrative agreements as compared to negotiators with any other combination of motives (De Dreu, 2010). Recent research now shows that the same appears to hold true for groups working cooperatively to solve a problem or make a decision. According to the model, high epistemic motivation involves a goal to be accurate or correct which should lead to deeper and more thorough information search and analysis (Kruglanski & Webster, 1996). Work on the information sharing effects has consistently demonstrated that instilling a goal of accuracy or defining the task in terms of solving a problem both increase information sharing (Postmes et al., 2001; Stewart & Stasser, 1992). Members high in prosocial motivation help to ensure that all types of information held by each member are likely to be disseminated, rather than just information supporting the position held by an individual member. Recent research showing that member focusing on preferences rather than information tends to impede information sharing is quite consistent with this assertion (Mojzisch & Schutz-Hardt, 2009). The model predicts that group information processing will only approach optimal levels when group members are high on both epistemic motivation and pro-social orientation. This is because that is the only combination which produces both systematic and thorough processing of information in an unbiased manner. Although the model is fairly recent, it does a good job of explaining several well replicated findings and has fared well in the few direct attempts to test it (Bechtoldt, De Dreu, Nijstad, & Choi, 2010; De Dreu, 2007).

Divisible Tasks

Steiner's (1972) definition of divisible tasks probably maps more closely onto most organizational team tasks than the unitary tasks discussed thus far. Unfortunately, much more work on information processing has been done on unitary tasks. This is partially a function of using laboratory studies to follow how information flows through a group. Unitary tasks are easier to use when time is limited, and the implications of information are more clearly defined for unitary tasks. However, recent research in organizational contexts has given cognition a much more prominent role (Salas, Goodwin, & Burke, 2009). Another difference in between unitary and divisible tasks involves how information is used. For unitary tasks, information exchange and processing are usually oriented toward solving a specific problem or choosing a particular course of action. Although information processing can also serve this function, more often it serves a coordination function. When subgroups are working on different aspects of a task that are interdependent, knowing when other subgroups may complete their task is important for knowing how to judge the timing and performance on their own subgroup (Marks, Mathieu, & Zaccaro, 2001). Thus, information processing is critical for performance on divisible tasks but in different ways.

Probably one of the main "cognitive" constructs relevant to divisible tasks are shared mental models (Hinsz, 1996; Cannon-Bowers et al, 1993). Mental models refer to mental representations of the task and the behaviors associated with performing the task (Rouse & Morris, 1985). At the group level, mental models also involve roles and interdependencies among group members. Cannon-Bowers et al. (1993) differentiated between task models and team models. Task models involve the various steps involved in the task and the resources (equipment, etc.) necessary to accomplish it. Group, or team, models involve the information and skills that members have that are relevant to the task and the ways in which their skills and behaviors must be coordinated to move efficiently toward task completion. Such shared cognitive structures help team members to coordinate actions and interpret information from other team members in consistent ways. They allow team members to develop similar causal explanations of the environment and to more effectively communicate implicitly, both of which improve team performance. Team mental models can enhance performance to the degree that the models are accurate and the members all share the same model (Salas, Rosin, Burke, & Goodwin, 2009).

Team training on both task and team models tends to improve performance by insuring that all aspects of both models are shared (Cannon-Bowers, et al. 1993). A well-known team training program, Cockpit Resource Management (Weiner, Kanki, & Helmreich, 1993) shows how training helps to create effective mental models. In an attempt to decrease errors in cockpit crews, Weiner et al. (1993) had each member of the crew train not only on their specific role or task, but also on every other role in the cockpit. This allowed team members to better understand how their role fit in with other roles and how the information that they possessed affected other roles. In addition, team members were trained to feel comfortable communicating the information that they had and to argue for its relevance in the presence of higher status team members. Thus, teams were trained to share both a mental model of the cockpit but also the appropriateness free flowing information exchange across team members and status differences. Teams trained in this way showed substantial reductions in errors and an increase in airline safety. Similar performance enhancements have been shown for surgery teams in hospitals.

However, sharedness for either the task or group model will only enhance performance to the degree that the model is accurate. Stasser and Augustinova (2008) have shown that distributed decision situations often produce better outcomes if information is simply sent up through the system by each group member without requiring any type of intermediary judgments by others. However, many groups assume that allowing judgments from various members is useful and thus use such a model to guide their behavior. Although aggregate judgments by many actors with different types and amount of information tend to be more accurate than judgments made by single individuals (Kerr & Tindale, 2011), in distributed systems where each member has only one type of information, asking all the members to make judgments adds noise to the system. In addition, research has shown that it is better for members not to know that others might have the same information that they do because it reduces their feelings of criticality and decreases the likelihood that they will send all their relevant information forward. Tschan, Semmer, Gurtner, Bizzari, Spychiger, Breuer et al. (2009) have shown that critical information easily available to emergency medical teams is often overlooked because each member assumes that someone else would have discovered and presented the information if it was relevant. Thus, intuitive mental models shared by group members can inhibit performance if they are inaccurate in terms of the task or if they lead to decreased information sharing.

Another type of shared mental model that has received a fair amount of attention in the literature is “transactive memory” (Wegner, 1987; More recent review). Using an individual - level metaphor, Wegner argued that team members encode, store, and retrieve information much like single individuals do. However, unlike individuals, teams have multiple information storage units each associated with a different member. Thus, the memory capacity of a team is considerably larger than that of any given team member. However, for the group to be able to use the additional memory storage efficiently, different team members must encode and store different information. For teams working on divisible tasks, the different aspects of the task often define which member will be responsible for encoding and storing certain types of information. For example, pilots may be responsible for knowing flight plans and schedules, whereas copilots maybe responsible for knowing current protocols for final safety checks. The copilot does not need to remember specific details on the flight plan because he/she could always retrieve them from the pilot, and vice versa. Although groups working on unitary tasks can divide up relevant information about the task and form transactive memory systems (Stewart & Stasser, 1995), such systems tend to form naturally for groups working on divisible tasks (Baumann et al, 2005). Transactive memory systems allow for the efficient storage and retrieval of information and also increase team memory capacity. However, for a transactive memory system to work, the members must share a model of who knows what.

Liang, Moreland, and Argote (1995) showed that training groups together as a team, rather than training each individual separately, naturally leads to the formation of transactive memory systems and that such systems lead to better performance. Three-person teams were trained on how to assemble a small radio. The assembly consisted of three component parts and each team member was trained on each component. Half of the teams were trained together and practiced the different components as a team. The other half involved each individual member being trained individually and then the three individuals were brought together to work as a team. Teams trained together performed better than teams trained as individuals and the development of a transactive memory system (which team members were better at and knew more about different components) accounted for the difference. These results have now been replicated a number of times (Moreland, Argote & Krishnan, 1998; Moreland, 1999) and have generalized to training groups in more natural settings (need a reference here).

Although Steiner (1972) referred to divisible tasks within groups where different members perform different subtasks, recent theorizing on teams in organizations has conceptualized parts of organizations as "multi-team systems" (Mathieu, Marks, & Zaccaro, 2001). A multi-team system involves "two or more teams that interface directly and interdependently in response to environmental contingencies toward the accomplishment of collective goals" (Mathieu et al., p. 290). Thus, the larger organizational task is divided among different teams. De Church and Mathieu (2009) argue that such systems can be interdependent in at least three domains: inputs, processes, and outputs. As long as the teams share at least one over-arching goal and show interdependence in at least one domain, the teams can be seen as forming a system within the larger organization. Many of the same issues associated with teams working on divisible tasks also appear for multi-team systems. The types of interdependencies will delineate what types of interaction are necessary among the teams. Findings from one team may serves as inputs to another or their actions may require a certain degree of coordination over time to insure efficient system functioning. A recent research focus by groups and teams researcher that would fall within this conceptualization is multidisciplinary science teams (need reference). Many societal problems involve issues that span different levels of analysis and scientific disciplines. For example, understanding global warming and finding ways to slow it down involves meteorology, chemistry, psychology, and sociology if not additional disciplines as well. Thus, future research helping to further explain how groups/teams or team systems can operate most efficiently may help to solve other societal problems as well.

Summary and Conclusions

Research to date shows quite clearly that distributed cognition is one of the strongest aspects of group or team performance. The ability to combine information from multiple sources to solve a problem or choose a course of action is what allows groups to perform better than individuals working alone (Hinsz, Tindale, & Vollrath, 1997). Recent research on aggregation has shown that even without inter-member communication, the diversity of knowledge across members allows groups to make accurate judgments (Larrick & Soll, 2006). However, for groups to maximize their potential and to ensure that unique information, other cognitive and motivational factors must be involved.

First, in addition to unique, distributed knowledge, groups/teams need a core base of shared knowledge so that all members can see the relevance of the unique information when it is shared (Laughlin & Ellis, 1986; Cannon-Bowers, et al., 1993; Hinsz, 1995). Thus, appropriate shared background knowledge or accurate shared mental models allow distributed cognition to aid in group performance. Such shared cognitions also allow team members to better coordinate their efforts. Second, groups/teams need to have high epistemic and well as social motivations in order to fully use the information at their disposal (De Dreu et al., 2008). Research has consistently shown that viewing tasks as having a correct or optimal solution leads to better information sharing (Stewart & Stasser, 1995). In addition, recent work on team forecasting has found that being open to the opinions of other group members is one of the key predictors to team success and forecasting improvement over time (Mellers, et al., 2014). When distributed cognition is yoked with accurate shared knowledge, appropriate motivation, and well-coordinated action, teams should be able to use their distributed knowledge to its fullest extent.