Int. J. Human–Computer Studies (1997) **47**, 429–452

Attribution accuracy when using anonymity in group support systems

STEPHEN C. HAYNE
School of Management, Arizona State University West, 4701 W. Thunderbird Rd, Phoenix, AZ 85069-7100, USA. email: hayne@asu.edu

RONALD E. RICE
School of Communication, Information & Library Studies, Rutgers University, New Brunswick, NJ 08903-5067, USA. email: rrice@scils.rutgers.edu

This study explores the taken-for-granted assumption that "anonymous" comments posted on a group support system (GSS) are *socially* as well as *technically* anonymous. It analyses the accuracy of, and influences on, attributions of authors' identities in seven field groups with considerable work history after they used the system to enter technically anonymous comments about salient topics during a brainstorming session. GSS participants made attributions about authors' identities, but overall these attributions were about 12% accurate (ranging from 1 to 29%). The expected predictors of accuracy (an individual's total communication with the group, network centrality, and length of membership in the group) were inconsistent influences across the seven groups.

1. Introduction

Organizational work style and pace are changing, with greater reliance on the use of teams, cross-functional and inter-organizational interactions involving groups of skilled workers with diverse (Keltner, 1989). To maintain coordination and communication amongst these teams, companies hold meetings. While such meetings are the most critical communication forum for many organizations, frequently they are time-consuming and unproductive experiences (Hackman & Kaplan, 1974; Mosvick & Nelson, 1987; Rice & Shook, 1990; Pollard & Hayne, 1996).

A fairly new set of technologies designed to assist groups is called group support systems (GSS) (Jessup & Valacich, 1993). GSS support a range of different group activities such as brainstorming, categorization, decision-making and voting (Hackman & Kaplan, 1974; DeSanctis & Gallupe, 1987; Dennis, George, Jessup, Nunamaker & Vogel, 1988; Zigurs, Poole & DeSanctis, 1988; Nunamaker, Dennis, Valacich, Vogel & George, 1991; Hiltz & Turoff, 1993).

One difference between a GSS meeting and a face-to-face meeting is that a GSS can provide the option for group members to contribute comments or messages *anonymously*, i.e. without revealing the author's identity. Anonymity in GSS is interesting for several reasons. First, prior research indicates that anonymity is an important variable in studying, evaluating and managing GSS. Second, there is increasing interest in examining anonymity as a determinant of a variety of group outcomes, in particular idea

Efforts by many researchers (Hackman & Kaplan, 1974; Kerr & Bruun, 1981; Siegal, Dubrovsky, Kiesler & McGuire, 1986; George, Easton, Nunamaker & Northcraft, 1990; Jessup *et al.*, 1990*a*; Pinsonneault & Kraemer, 1990; Connolly *et al.*, 1991; Jessup & Tansik, 1991; McLeod, 1992; Jessup & Valacich, 1993) have generally found an increase in production and satisfaction when anonymous group brainstorming is used. Other advantages of anonymous participation include decreased evaluation apprehension, decreased member domination, decreased conformance pressure and decreased status competition, which can lead to increased exploration of alternatives and surfacing of assumptions.

However, potential disadvantages include social loafing, failure to listen, disinhibition, deindividuation and poor socialization, which can lead to decreased effectiveness and group dissatisfaction. Anonymity may not affect the ability to contribute, but it does seem to affect what a participant is willing to say, and how it is said.

Valacich, Jessup, Dennis and Nunamaker (1992) provide a major review and conceptual framework which discusses the role of anonymity in GSS. They propose two types of anonymity: process and content. Process anonymity promises a participant that others will not know whether or not the particular participant is in fact participating. (Process anonymity is not often available during same time, same place GSS meetings, because participants can readily see who is there.) Content anonymity prevents others from knowing who contributes particular messages. Content anonymity may be either local (during the process) or global (extended even to later analysis of transcripts or logs). An individual can take advantage of process anonymity, but inadvertently or intentionally reveal themselves with their message content. Both of these forms of anonymity focus on the technical aspect of anonymity.

Such emphasis on technical anonymity tends to ignore group members' cognitive and social processes, seeing individuals instead as minimediators of larger group processes. This is inappropriate, because individuals bring with them a myriad of information that can be brought to bear upon anonymously entered text to uncover cues and then associate text with other group members.

Further, Licker (1992*a*, *b*) proposed that anonymity is a state of the individual, experienced either as *Identitylessness* or *Source Dissociation.* Identitylessness is the perception that one does not have an identifiable locus in a setting—i.e. others do not know I am a participant or what my role is—whereas Source Dissociation is a feeling that others cannot identify one as the source of specific messages—i.e. others do not know I entered this comment. It is the source dissociation aspect of anonymity, a *social* and not a *technical* condition, that we find interesting.

However, most GSS researchers adopt a definition of anonymity that treats anonymity as a purely *technical* feature. This techno-centric definition hides the more social nature of anonymity in communication settings; we feel anonymity is not just a switch that can be turned on and off. First, people are seldom unidentified in every way. Second, it is generally artificial to be an anonymous participant in most kinds of non-GSS meetings (George *et al.*, 1990), although we acknowledge that short-lived anonymous situations such as expressing preferences through blind balloting do exist. Moreover, Mennecke, Hoffer and Wynne (1992) suggest that even during GSS sessions, "extra-meeting" influences, such as those discussions during coffee breaks, may serve as inputs

Researchers have attempted to go beyond the basic premise of attribution theory to specify the information (or social cues) that humans use in inferring the causes of an observed event (or statement), and the consequences of such inferences for the individual. Kelley (1973) postulates that individuals employ certain informational cues—"consensus" (conforming behavior implies situational influences; unique behavior implies individual characteristics as the influence), "distinctiveness" (highly distinctive behavior is attributed to environmental conditions, while consistent behavior is attributed to individual personality) and "consistency" (consistent individual behavior implies the causes of that behavior are also stable)—to determine the causes to which behavior is attributed. These three sources of information are obtained through observation of behavior across people, situations and time, respectively (McArthur, 1972; Hewstone & Jaspars, 1987; Hilton & Jaspars, 1987; McGill, 1989).

Iacobucci and McGill (1990) have meta-analysed published data to demonstrate that several effects are particularly robust. Even over varying paradigms and procedures, consensus and distinctiveness contribute to person *attributions*, distinctiveness contributes to stimulus *attributions*, and consistency contributes to circumstance *attributions*. Thus, the above three major effects involve *making attributions about people.*

Hayes and Hesketh (1989) argue that attribution theory can be further enhanced by considering the evidence that a wide range of cognitive biases and errors exist (Kahneman & Tversky, 1973; Tversky & Kahneman, 1974; Hamil, Wilson & Nisbett, 1979; Henrion, 1980; Nisbett & Ross, 1980). That is, people make attributions based on a variety of cognitive processes that may involve considerable bias and inaccurate inference. Worse, Arkes (1981) has stated that merely explaining a bias to people and then instructing them not to be influenced by it is not an effective strategy for removing the bias. Therefore, more "active" strategies have been suggested (Koriat, Lichtenstein & Fischoff, 1980; Nisbett & Ross, 1980; Arkes, 1981).

Why is this important to our current argument? As summarized in the prior section, use of an anonymous GSS alters the attribution process. *Anonymity in GSS comments supposedly removes personal biases by removing the identity of comments' authors from the content of those comments during the process of decision making.* However, (1) people may continue to attribute statements made in the GSS even in anonymous settings (as people seem disposed to constantly make attributions, frequently, about others), (2) these attributions may not be accurate and may be highly biased and so (3) individual decisions, contributions, discussions and judgments may be affected by mis-attributions in conditions of anonymous input. Note we are concerned with individual-level issues, but in the context of group work.

2.4. SOCIAL CUES AND MEDIATED COMMUNICATION

All cues that individuals provide, or other individuals perceive, during an interaction may both stimulate, and be used as evidence for, one's attributions. Obviously, we pay close attention to the content of what a person is saying or doing. However, other specific cues, such as attractiveness, dress or speech may be especially salient due to context, expectations or attention. These salient cues are then the basis for forming impressions and evaluating others (Taylor & Fiske, 1978). Salient stimuli encourage the use of simple

by decoding text-based cues and derive psychological-level knowledge about other actors from CMC interaction" (Walther, 1992, p. 67). Walther argues (1992), and finds (1993) that impression formation through CMC may increase over time to approach the initially higher levels of face-to-face impression formation, and that this will be more likely for ongoing CMC sessions and for smaller groups. With groups that already have considerable experience communicating through a variety of channels with other group members, users may be able even more easily to interpret and decode textual entries, and associate relational impressions to stylistic characteristics and stated positions available in others' comments, even though the comments are technically anonymous. Lea and Spears (1991) note that paralanguage (the use of typographical and punctuation symbols to convey tone and emotion) is one means by which social information is communicated in CMC. The meaning of such paralinguistic marks is dependent on the group or individual communication context. They also found that deindividuated subjects for whom group identity had been made salient would evaluate users of paralanguage more positively than when group salience was low, supporting their argument that social context could be highly influential in mediated communication.

Even if nonverbal cues (visuals, tone, inflection, eye-gaze, appearance, etc.) are "filtered out", it may not be the case that the majority of evaluative content and social cues is, in fact, carried in the nonverbal portion. Krauss *et al.* (1981) argue that the high percentage of communication of emotion carried in nonverbal modes commonly cited are derived from atypical experiments. Their own study compared the extent to which people evaluated, and correctly evaluated, the emotional affect of a presenter, through several different channels in two studies: full-channel (audiovisual), verbal written (transcripts), vocal (content-filtered speech) and visible (silent video). In both studies, verbal content was the best predictor of evaluative judgments. "Overall, no support was found for the widespread assumption that nonverbal channels (vocal or visible information) form the primary basis for the communication of affect" (p. 312). They also noted several other studies that found similar results, but the authors did agree with the proposition that nonverbal cues provide the context for accurate interpretation of verbal content.

What all these propositions and results mean is that readers of text-only comments in a technically anonymous GSS might perceive enough cues to make attributions, and that those attributions might be at least somewhat accurate. For example, an author might always use an accepted abbreviation or jargon (consensus), might use a unique grammatical style (distinctiveness), and might repeatedly state a position that a reader knows the author holds (consistency). In this case, accurate attribution may be possible.

2.5. PRIOR COMMUNICATION

We suspect that participants might, in fact, be able to identify the authors of such technically anonymous comments, especially in groups that have a considerable history of interaction. Knowledge of the other participants' identities brings to bear prior knowledge of relevant cues such as status, expertise, authority, conflict or social role, thereby reducing the benefits of anonymity. This could give some participants an advantage, biasing the evaluation of comments or even allowing some individuals to

3. Hypotheses

Thus, the two general research questions for this study are: (1) *can we accept the assumption that using the technical "anonymity" feature of a GSS will truly preserve "anonymity" and* (2) *what might influence attributions of identity even under conditions of "anonymity"*? Three sets of directional hypotheses follow from the prior discussions:

H1. Group members can accurately identify the source of technically anonymous comments.

H2a. Longer membership in a group is positively associated with one's mean percent correct identification of authorship of technically anonymous comments.

H2b. The total amount of one's prior communication within a group is positively associated with one's mean percent correct identification of authorship of technically anonymous comments.

H2c. One's position in the prior communication network, as measured by betweenness centrality, is positively associated with one's mean percent correct identification of authorship of technically anonymous comments.

H3. The prior pairwise communication among group members is positively associated with the percentage correct identification of authorship of technically anonymous comments between those pairs of group members.

Hypothesis H1 is conceptually independent from the other hypotheses. That is, if people cannot accurately make attributions, the attribution process itself might still be predictable.

4. Methods

We have engaged in an exploratory study involving a "real-time" survey of GSS use. The groups studied herein are not highly controlled laboratory, student or simulated groups (as critiqued by McGrath & Hollingshead, 1993), but are field groups attempting to resolve actual business problems. Natural groups are often part of a larger organization, are likely to have worked together previously and may be working on several projects simultaneously. Groups do not, of course, have to represent projects or teams—they may be any set of people with regular interaction and a common goal.

4.1. PARTICIPANTS

The seven organizational units involved here are all operations divisions of mid-sized firms—i.e. subjects were staff, managers and executives, not faculty or students. All groups consisted of people who had worked together previously—some for up to 15 yrs—and contained members who worked together daily. All were meeting and using a GSS to resolve real, salient and ongoing business problems. These groups were *not* standing "committees" or "task forces" assigned to assess a particular project; the groups were made up of almost all the members of their respective business unit.

The first four groups ("PE" 1–4) were departments belonging to the physical educa-

4.3. DATA

Three forms of data were collected. The first was a communication network roster administered before each group began its session. This roster simply listed the names of participants, and asked each participant to indicate the extent to which they communicated with each other listed person about any topic, using any communication channel (from 0 = never, 1 = once every few months, to 7 = several times a day).

The second form of data were a sample of the actual comments participants entered. Comments were restricted to those longer than 65 characters in length in order to reduce comments like "Me, too!". Our intention was to provide comments with at least a minimal amount of social cues.

Fifteen minutes before the end of the GSS session, one of the researchers used a simple macro to select up to 50 comments from the issue areas with the most comments entered into them at that time (potentially different across sessions), in order to maximize the chance that the comments had been seen by everyone in the group. The comments were randomly chosen in proportion to the numbers of total comments in each issue window, yet the original order of the comments from Topic Commenter was preserved. Some of the issues received more *attention* than did the others. We assumed that more comments in an issue area might mean that more people had browsed the comments. Note that we have no basis to assert that any *particular* comment was read by any *particular* member. These selected comments were copied into a word-processing file and included in the survey to create and print a "comment" questionnaire that was distributed immediately at the end of the phase.

The comment questionnaire was the third form of data. First it asked several individual-level questions about group membership and attitudes toward anonymity. These were followed by a list of each of the participants with an assigned sequential letter, and a listing of the selected comments for that group. Since Topic Commenter is designed to maintain and guarantee anonymity, it could not record who actually entered each comment. So, the participants were first asked to indicate which comments were theirs (if any), by entering their own letter alongside them. Note that we have no basis to assert that each such self-identification was correct. However, there are two bounds to this assertion: (1) very few comments were claimed by more than one participant and (2) very few people did not claim any comments (see Table 1, from 1 to 9). Then the

TABLE 1
Comment totals

Group	Made by group	Selected	Claimed by more than one person	Not claimed by anyone	Mean percent attributed
PE1	228	52	1	1	21
PE2	352	26	3	2	21
PE3	200	39	3	1	9
PE4	156	50	1	1	23
Consult	134	35	1	3	70
Pipe	159	40	1	9	34

centrality is standardized (according to Freeman's operationalization) to remove the effect of differing group sizes. First the network matrices were symmetrized by putting the average of the $\text{cell}_{i,j}$ and $\text{cell}_{j,i}$ into both cells, then dichotomizing those values at 3 (communicating several times a month). That is, the average extent of communication between two group members had to be at least several times a month to be considered potentially influential. From this symmetric binary matrix, betweenness centrality for each individual in the group was calculated as the average of all the geodesics (the shortest path between each pair of group members) on which each individual was located.

Accuracy of attributions (from post-session survey).

3. The *average normalized attribution accuracy score* (row mean, ignoring the diagonal) for each member, based on the *square, asymmetric matrix of normalized attribution accuracy* by each member of each other member's comments.

Each cell in each group's accuracy matrix is filled with the result from the equation

$$\text{cell}_{i,j} = (\text{sum}[(a_{ij}) - (k_j)])/n$$

where i, j is the number of participants, $i \diamond j$, a_{ij} the attribution by i that comment k was made by j (1 if made, 0 otherwise), k_j the comment k made by j (0 if true, 1 otherwise) and $n = k$ from 0 to n, the total number of attributions made by i. For each comment listed on the comment questionnaire, *if* person i makes an attribution of authorship to person j *then* that cell i, j gets a 1 if correct, 0 if incorrect. If i did not make an attribution for a comment, the value is set to missing, so that nonattributions do not count either as accurate or not accurate. That is, this formula does not presume that i *should* make attributions to every comment; it only considers when i *does* make attributions. This is done for each, and summed over all, of i's attributions to comments made by j. This number is then divided by the number of attributions i made. There was no correction for random (as opposed to well-intentioned) guessing. Thus, the attribution accuracy score for each i, j relation theoretically ranges from 0 to 100%.

4. *Number of attributions* made by each participant.

5. *Percent attribution accuracy*, the number of correct attributions made, divided by the total number of attributions made, for each participant. This is set as *missing* for those who made *no* attributions.

Individual-level variables (from post-session survey).

6. *Duration of group membership* (how many months they had been a member of this particular group) used in many studies mentioned in Jessup and Valacich (1993).

7. *Perception of anonymity* (how anonymous they felt they actually were during the Topic Commenter part of the meeting, from 1 = not at all to 7 = completely) (Jessup & Tansik, 1991).

8. *Importance of anonymity* (how important it was to them for the Topic Commenter part of this meeting to be anonymous, from 1 = unimportant to 7 = crucial) (Jessup & Tansik, 1991).

9. *Presence of cues* as indicated by responses to an open-ended question asking whether there were any events during the meeting that made it easy for the researchers to

5. Results

The percent of all comments attributed by the group ranged from a low of 9% for PE3 to a high of 70% for the Consult group, with a similar range in mean number of attributions by individuals from 3.6 in PE3 to 24.6 in Consult. That is, groups ranged widely in the percentage and frequency of attributions. People felt they were mostly anonymous ($M = 5.4$), and about equally concerned that these comments should be anonymous in these meetings ($M = 5.3$). There was no overall statistically significant difference among the groups on these two perceptions. On average, members belonged to these groups for almost 5 yr ($M = 55.9$ months), with only PE2 having a significantly higher mean.

Very few people provided any indication (other than the content of the comments) that there were events that could make it easy to identify comments: most either left the question blank or wrote "no" or "none". Five people mentioned that it was occasionally possible to see others' screens. The following remarks came from the Consult participants: "certain comments have been made in the office and appear here"; "fluency in typing"; "through joking discussion or the last person was still typing"; "the person uses many words not normally used by others"; "recognize some people's English or lack thereof" and one person in the Consult group apparently added his initials to his GSS (not on the survey) comments, which two other respondents from the group mentioned. As noted below, this group had the second-highest level of perceived anonymity, a medium level of anonymity importance and was by far the most accurate group! The remaining comments came from the PE groups (each of which, remember, included people from the other three PE areas): "how people express themselves"; "strong opinion that serves their area"; "people naming their areas" and "just the type of answers that were given led me to assume I knew the person who answered". Each of these comments about cues, except for being able to see another's screen and seeing who's left typing, presumes prior knowledge of the other person from communication cues or associated issue positions.

Notice that we are not able to determine what percentage would be below, equal to or above chance guessing. Due to the nature of this commercial anonymous GSS, we cannot capture the number of comments each participant contributed, and it would be unwise to assume that the proportion of comments claimed by a group member would hold for the totals in the session. Thus, we are not assessing *guesses or errors of omission*, only the accuracy of attributions that are *made*.

5.1. HYPOTHESIS TEST: INDIVIDUAL INFLUENCES

Mean accuracy ranged from 10% for PE3 and Oil group to 29% for the Consult group, with an overall average of 12% accuracy (with the Consult and Pipe groups significantly more accurate than the other groups). When our respondents indicated attempts at authorship attribution, they were wrong almost 90% of the time. Hypothesis 1 must be rejected (note, however, that they were *right* more than 10% of the time!). Furthermore, there is no correlation between a group's accuracy and the group size or the number of comments selected for the survey.

The duration of group membership was significantly correlated with attribution

5.2. MULTIPLE REGRESSIONS

Because there were some significant correlations among the independent variables, a final summary stepwise multiple regression was run for each group, including network row total, network column total, network centrality, months in the group, perceived anonymity and importance of anonymity (see Table 5).

For two groups (PE1 and PE4), number of months in the group was the only significant positive predictor of accuracy. In two other groups (PE2 and Consult), months approached significance as a predictor but was *negatively* associated with accuracy. For the Pipe group, months was a significant positive predictor, and network centrality was a significant negative predictor.

5.3. HYPOTHESIS TEST: NETWORK LEVEL INFLUENCES

The test for Hypothesis 3 compares each group's comment accuracy matrix to the group's communication network matrix, based upon pairwise patterns of similarity between i, j in each matrix. Traditional statistics are inappropriate for testing the significance of network matrices, because the rows and columns of a who-to-whom communication matrix are likely not independent. The generally accepted approach is to first reorganize each matrix into a vector, and then compute the correlation between the two vectors, to produce a familiar measure of the strength of association, the Pearson r. Then the quadratic assignment procedure (QAP) (Hubert & Schultz, 1976) is used to determine the *significance* of the association between the two matrices, by means of a nonparametric gamma association significance test (derived from a distribution of possible associations between one matrix and all permutations of the other matrix). However, because there are asymmetric missing values in the accuracy matrix (when j made a comment but i did not make an attribution), even this significance test will not

TABLE 5
Regressions of accuracy on network, duration and anonymity variables

			Equation parameters	
Group	Standardized betas		Adj. R2	*F*-ratio
PE1	Months	0.85****	0.69	25.0****
PE2	Months	− 0.37 n.s.	0.08	2.5 n.s.
PE3	—	—	—	n.s.
PE4	Months	0.59*	0.21	3.5*
Consult	Months	− 0.44 n.s.	0.13	2.9 n.s.
Pipe	Months	0.65****	0.46	10.9****
	Net Cent	− 0.42***		
Oil	—	—	—	n.s.

Note: Months = number of months as member in group.

Second, it may be that *technical* anonymity is fairly "effective"—i.e. authorship is separated from content and people cannot easily identify authorship. It may also be the case that while we asked about the extent of prior communication based on all channels, the GSS comments are of course textual, which represents only one of the channels used in prior communication and, thus, can provide only a subset of possible social cues used as the basis for attributions. Future studies might look at prior communication through written media only. This potential threat to the validity of the results presumes, however, great stylistic and cue differences for individuals between oral and written communication and that these are not correlated.

It is also plausible that some users would craft their comments to conceal their identity and thus maintain anonymity. This might decrease the accuracy of attributions, but would not affect attempts at attribution. Another possible contribution to Type-II error might be that we asked participants to attribute statements to people they have worked with but may not know that well. However, participants were asked to make attributions of comments (not people), were asked only to make attributions if they thought they knew who the author of the comment was, and had worked with each other for many years. Further, the pairwise network matrix comparison specifically tests for differences in strength of prior interactions.

Third, different naturally occurring groups may have different abilities to evaluate cues from and attribute authorship to technically anonymous cues in a GSS. The results did not reveal a pattern of relationships that was consistent across every group. It may be that, while technical anonymity "works" (it separates content from authorship identity), low social anonymity confounds, or reassembles, these components, leading people away from an absence of any attributions to a presence of false ones. The more people communicate with others, the more they may think they "know" others' thoughts and cues, leading them to overestimate their confidence in being able to attribute (mediated) content to (technically anonymous) authors. In general, however, the small sample sizes (in the statistical sense) of the separate groups increase the likelihood of Type-II error.

To the extent that prior communication with one's group members does influence one's ability to detect and interpret cues about the identities of others making technically anonymous GSS comments, specific pairwise relations and specific network locations seem to be better predictors—both conceptually and empirically—than simple amount of prior communication with the group. Thus, attempts to assess the effect of anonymous comments should take a network-theoretic approach.

Apart from not being able to predict the level of attribution accuracy, there's a more disturbing issue involved here. It might seem that the technical anonymity of *actual* identities in GSS groups is essentially preserved in spite of attempts to attribute, because the users were largely inaccurate in their attributions. Note, however, *that this is not the same as maintaining technical anonymity, because users may still be making attributions.* That is, implementing technical anonymity is conceptually and pragmatically *independent* from preventing attributions (whether accurate or inaccurate). It is plausible that for groups with a salient task and prior history, GSS participants may well be quietly attempting to identify authors of technically anonymous comments. They would be applying status cues, etc., derived from prior communications (especially their location in the communication network) and from on-line textual style and word choice, to the

ing in GSS. Consequently, the level of inaccuracy reported here may not be as high in technically anonymous computer-mediated communication in general.

While our results are clearly not definitive, they do raise and clarify some previously hidden assumptions about GSS anonymity, and arouse concerns about subtle consequences of the use of GSS anonymity features. It is apparent from this preliminary study that the *social* aspect of GSS anonymity needs to be explored more fully; it is *not* the same thing as *technical anonymity*. We plan on attempting to resolve several problems with this study in future experiments that achieve the following.

1. Develop ways for the researcher to accurately identify an author's comments without compromising the natural group's comfort with the system.
2. Include other measures of prior relationships, such as time spent working with each other member.
3. Assess the extent to which people think they informally attempt to identify authors of comments.
4. Assess the extent to which they think their attempts are correct.

We feel that this initial study has uncovered some assumptions in prior GSS anonymity research that should be reconsidered.

We would like to acknowledge the intellectual contribution and access to meetings provided by Paul Licker (University of Calgary, Alberta), and the insightful comments provided by Joe Walther (NorthWestern University, Illinois) and various anonymous reviewers on an earlier version of this paper.

References

ARKES, H. (1981). Impediments to accurate clinical judgment and possible ways to minimize their impact. *Journal of Consulting and Clinical Psychology*, **49**, 323–330.

BERNARD, H., KILLWORTH, P. & SAILER, L. (1982). Informant accuracy in social-network data V. An experimental attempt to predict actual communication from recall data. *Social Science Research*, **11**, 30–66.

BLANK, R. (1991). The effects of double-blind versus single-blind reviewing: experimental evidence from *The American Economic Review*. *The American Economic Review*, **81**, 1041–1067.

CONNOLLY, T., JESSUP, L. & VALACICH, J. (1991). Effects of anonymity and evaluative tone on idea generation in computer-mediated groups. *Management Science*, **36**, 689–703.

CULNAN, M. & MARKUS, M. (1987). Information technologies: electronic media and intra-organizational communication. In F. JABLIN, L. PUTNAM, K. ROBERTS & L. PORTER, Eds. *Handbook of Organizational Communication*, pp. 420–443. Newbury Park, CA: Sage.

DAFT, R. & LENGEL, R. (1984). Information richness: a new approach to managerial behavior and organization design. *Research in Organizational Behavior*, **6**, 191–233.

DAFT, R. & LENGEL, R. (1986). Organizational information requirements, media richness and structural design. *Management Science*, **32**, 554–571.

DENNIS, A., GEORGE, J., JESSUP, L., NUNAMAKER, J. & VOGEL, D. (1988). Information technology to support electronic meetings. *MIS Quarterly*, **12**, 591–624.

DESANCTIS, G. & GALLUPE, B. (1987). A foundation for the study of group decision support systems. *Management Science*, **33**, 589–609.

FREEMAN, L. (1979). Centrality in social networks: conceptual clarification. *Social Networks*, **1**, 215–239.

LEA, M. & SPEARS, R. (1991). Computer-mediated communication, de-individuation and group decision-making. *International Journal of Man–machine Studies*, **34**, 283–301.
LICKER, P. (1992*a*). Knowing ourselves and knowing others: exposing anonymity in GSS. *Proceedings of the SW Decision Sciences Institute Conference*, San Antonio, TX, pp. 184–186.
LICKER, P. (1992*b*). You said what? Understanding the role and function of anonymity in GSS. *Proceedings of the Administrative Sciences Association of Canada Conference*, Quebec City.
LIM, F. & BENBASAT, I. (1991). A communication-based framework for group interfaces in computer-supported collaboration. *Proceedings of the 24th Hawaii International Conference on Systems Sciences*, Kauai, Hawaii, pp. 610–620.
MCARTHUR, L. (1972). The how and what of why: some determinants and consequences of causal attributions. *Journal of Personality and Social Psychology*, **22**, 171–193.
MCGILL, A. (1989). Context effects in judgments of causation. *Journal of Personality and Social Psychology*, **57**, 189–200.
MCGRATH, J. & HOLLINGSHEAD, A. (1993). Putting the 'group' back in group support systems: some theoretical issues about dynamic processes in groups with technological enhancements. In L. M. JESSUP & J. S. VALACICH, Eds. *Group Support Systems: New Perspectives*, pp. 78–96. New York, NY: Macmillan.
MCLEOD, P. (1992). An assessment of the experimental literature on electronic support of group work: results of a meta-analysis. *Human–Computer Interaction*, **7**, 257–278.
MENNECKE, B., HOFFER, J. & WYNNE, B. (1992). The implications of group development and history for group support system theory and practice. *Small Group Research*, **23**, 524–550.
MOORE, C. (1987). Group techniques for idea building. In *Nominal Group Technique*. Newbury Park, CA: Sage.
MOSVICK, R. & NELSON, R. (1987). *We've Got to Start Meeting Like This! A Guide To Successful Business Meeting Management*. Glenview, IL: Scott, Foresman.
NISBETT, R. & ROSS, L. (1980). *Human Inference: Strategies and Shortcomings of Information Judgment*, Englewood Cliffs, NJ: Prentice-Hall.
NUNAMAKER, J., DENNIS, A., VALACICH, J., VOGEL, D. & GEORGE, J. (1991). Electronic meeting systems to support group work. *Communications of the ACM*, **34**, 40–61.
PINSONNEAULT, A. & KRAEMER, K. (1990). The effects of electronic meetings on group processes and outcomes: an assessment of the empirical research. *European Journal of Operational Research*, **46**, 143–161.
POLLARD, C. & HAYNE, S. (1996). The reality of meetings and the use of automated meeting tools, *Proceedings of the Americas Information Systems Conference*, pp. 250–253.
RICE, R. (1984). Mediated group communication. In R. E. RICE & Associates. *The New Media: Communication, Research and Technology*, pp. 129–154. Beverly Hills, CA: Sage.
RICE, R. (1987). Computer-mediated communication and organizational innovation. *Journal of Communication*, **37**, 65–94.
RICE, R. (1993). Using network concepts to clarify sources and mechanisms of social influence. In W. RICHARDS, Jr. & G. BARNETT, Eds. *Advances in Communication Network Analysis*, pp. 43–52. Norwood, NJ: Ablex.
RICE, R. & LOVE, G. (1987). Electronic emotion: socio-emotional content in a computer-mediated communication network. *Communication Research*, **14**, 85–108.
RICE, R. E. & SHOOK, D. (1990). Relationships of job categories and organizational levels to use of communication channels, including electronic mail: a meta-analysis and extension. *Journal of Management Studies*, **27**, 196–229.
SHORT, J., WILLIAMS, E. & CHRISTIE, B. (1976). *The Social Psychology of Telecommunications*. London: Wiley.
SIEGAL, J., DUBROVSKY, V., KIESLER, S. & MCGUIRE, T. (1986). Group processes in computer mediated communication. *Organizational Behavior and Human Decision Process*, **37**, 157–187.
SPROULL, L. & KIESLER, S. (1986). Reducing social context cues: electronic mail in organizational communication. *Management Science*, **32**, 1492–1512.
TAYLOR, S. & FISKE, S. (1978). Salience, attention and attribution: top of the head phenomena. In L. BERKOWITZ, Ed. *Advances in Experimental Social Psychology*, Vol. 11, pp. 249–288. New