Communication Cost, Belief Development, and Structural Change:
A Dynamic Model of Networked Communities of Practice

Brian S. Butler
(bbutler@katz.business.pitt.edu)

Katz Graduate School of Business
University of Pittsburgh
Pittsburgh, PA 15213

November 10, 2000

The work presented in this paper was supported by NSF Grant # CSS9711548.

Communication Cost, Belief Development, and Structural Change: A Dynamic Model of Networked Communities of Practice

Abstract

Communities of practice (COPs) play an important role in the use and transfer of knowledge within and between organizations. Increasingly, new technologies, such as electronic mail and the Web, are being applied to create alternative infrastructures for these communities. This paper presents a model of dynamic communities of practice (DYNACOP) that integrates the processes of member development, structural change, and communications costs. The DYNACOP model was calibrated and validated with empirical data from a sample of Internet listservs. The model predicts that infrastructures that reduce communication costs will ultimately result in slower development of focused, stable communities of practice. While reduced communication costs increases the efficiency of communication, it also alters the processes by which members form beliefs about the community, which in turn affects the structural development of a community's membership. Hence networked communities of practice are expected to have larger and more diverse, but also less stable and focused, memberships than traditional face-to-face associations.

Communication Cost, Belief Development, and Structural Change:
A Dynamic Model of Communities of Practice

The development of communities of practice (COPs) is an important element of management initiatives that focus on enhancing the learning capabilities of organizations. Communities of practice are social structures in which the members share knowledge about a common set of problems or interests (Brown and Duguid 1991). A successful COP leverages the knowledge of the members, providing an environment that supports individual learning (Lave and Wegner 1991). An organization's ability to successfully manage knowledge is dependent on the development of effective COPs (Allen 1984, Goodman and Darr 1999).

Communication technologies such as e-mail and the Web have the potential to significantly affect an organization's ability to support the development of communities of practice. Interpersonal communication technologies allow individuals to exchange information across the boundaries of time and space. The broadcast facilities of networked communication infrastructures allow people to ask questions to larger audiences (Sproull and Kiesler 1991) and generally be more aware of what is going in an organization (Kraut and Attewell 1997). Intranets, organizational memory systems, and other archival facilities support communities by capturing communication activity for use by others in the future. Networked communication technologies are expected to significantly affect the operation of COPs.

The choice of communication infrastructure affects the development of communities of practice. Face-to-face meetings, print, telephones, and electronic mail are all technologies that support communication among members of a COP. Each type of infrastructure has different implications for the mechanics of communication within a COP. However, when assessing the impact of a technology, it is also important to consider second-order consequences, or changes in individuals' behavior which arise as a result of the new technology (Sproull and Kiesler 1991).

Communication is only part of the operation of social collectives; members join and leave, individuals' beliefs about the value of membership change, and from these processes the focus and activities of a community develop. Thus when selecting a communication infrastructure for a COP, it is necessary to consider how the technology will affect both the communication behavior of individuals and the development of the community as a whole.

The proposed model of dynamic communities of practice (DYNACOP) incorporates elements of individual and structural perspectives on social structure development with a representation of communication costs to examine the role played by communication technology in the development of these structures. Individual change, in the form of member belief shifts, and structural change, in the form of membership movement, are integrated to describe the development processes of COPs. Communication technology determines the communication costs incurred by members, costs that impact a COP's development by affecting the way members process information about the community. Communication activity acts as a signal, providing information about the emergent features of the community. Through the communication infrastructure, that signal affects individuals' beliefs about the COP; those beliefs, in turn, underlie their decision to continue or end membership in the community. In this fashion, the linked processes of individual belief change and member movement interact with communication costs to dynamically determine the emergent features of a COP.

Background

Existing streams of research independently consider the role of member change, at the individual level, and membership movement, at the group level, on the development of social structures (Tuckman 1965, 1977; McPherson 1983a, 1983b, 1990). The individual approach is characterized by a concern with the perceptions, beliefs, and behavior of individuals within groups. Studies of group development focus on how members change over time (c.f. Gersick

1988), and social groups are seen primarily as a context in which to consider individuals' mental states and processes (Moreland, Hogg, and Hains 1994). Working from this perspective, some theorists have proposed models that characterize the development of groups in terms of stages (Tuckman 1965, Tuckman and Jensen 1977, Hare 1976, LaCoursiere 1980, McGrath 1984), describing each stage in terms of prototypical member behaviors or attitudes.

However, sequential stage models are limited in that they are essentially descriptive models, providing a set of snapshots of prototypical member characteristics (Tuckman 1965, Hare 1976, Poole 1983). They say little about the mechanisms which underlie the observed changes (Gersick 1988) and the sequential structure is generally a poor representation of the development process observed in naturally occurring groups (Fisher 1970; Scheidel and Crowell 1964; Poole 1981, 1983). As a result, recent work in this area has begun developing models of the social processes which link individual attitudes and behaviors to developing features of social structures (Gersick 1988, 1989; Moreland and Levine 1982; Worschel 1996). However, because they have historically focused on small groups, consisting of a few members meeting face-to-face, these models do not directly address questions about the role of communication and communication technologies in the development processes of social structures.

Under the structural approach, social structures are seen not as environments for member change, but as dynamic social entities that arise within larger social systems (c.f. Carley 1991). Structural theorists characterize these systems many ways, including institutional (Blau 1967, Etzioni 1964, Simmel 1955), ritual (Goffman 1959), and competitive (Hannen and Freeman 1977, McPherson 1983b, McPherson and Rotolo 1996) models. These studies focus on the processes, such as membership movement and resource contribution, which underlie the formation and continued existence of structural entities. However, the models typically do not link individual behaviors and structural development processes (Carley 1991). Consequently, with the exception of recent work with computational models of emergent social structures (Kaufer and Carley 1993; Carley and Wendt 1991; Carley 1995a, 1995b; Huberman and Hogg 1995) structural models, as with individual level studies, do not directly consider the role of communication or communication technology. Thus while the individual and structural perspectives provide a theoretical foundation for modeling the dynamics of COPs, they provide little guidance regarding the consequences of different communication infrastructures.

Although traditional studies of social structure have generally not considered the impact of alternative communication infrastructures, there is a growing literature about technology-supported collectives. Studies of group decision support systems (GDSS) and electronic meeting systems (EMS) have generally adopted the individual approach, considering how the use of new communication technologies changes individuals' perceptions, and behaviors in small groups (for a review see McGrath and Hollingshead 1994). This work provides insight into the social behaviors of individuals in networked environments. However, as with traditional studies taking this approach, it often unclear how (or if) changes in individual members will affect the overall development of the structure (McGrath and Hollingshead 1994). In contrast, field studies of networked collectives have generally adopted a structural perceptive, focusing on the processes and structure of online collectives (for a review see Butler 1999). Through rich description, these studies provide insights into the features of social structures in networked environments. However, much of this work is focused on demonstrating that 'real' social activity can occur in online contexts. As a result, this work also fails to systematically consider how alternative technologies might differentially impact the development of social structures such as COPs.

DYNACOP: A Model of Dynamic Communities of Practice

This model of dynamic communities of practice (DYNACOP) describes the development of a social system that consists of a set of N individuals and a single community of practice (COP). The COP is a socio-technical structure that supports bounded many-to-many communication among a collection of members who individually choose to contribute and be exposed to the communication activity. Unlike some recent theoretical studies of social structure, in which groups are conceptualized as purely emergent structures (Carley 1991, 1995a, 1995b; Epstein and Axtell 1996; Huberman and Hogg 1995), in the DYNACOP framework the community of practice is characterized as an independent social entity which is known to the individuals within the social system. Individuals treat the COP not as a collection of distinct individuals, but as a composite social agent (Moreland and Levine 1982).

Like individual agents, the COP has some features and processes, such as message impact and noise cost, which are set as parameters. However, unlike individuals, the COP also has features, such as membership and communication activity, that are the result of aggregation of other agents' choices and actions (Figure 1).

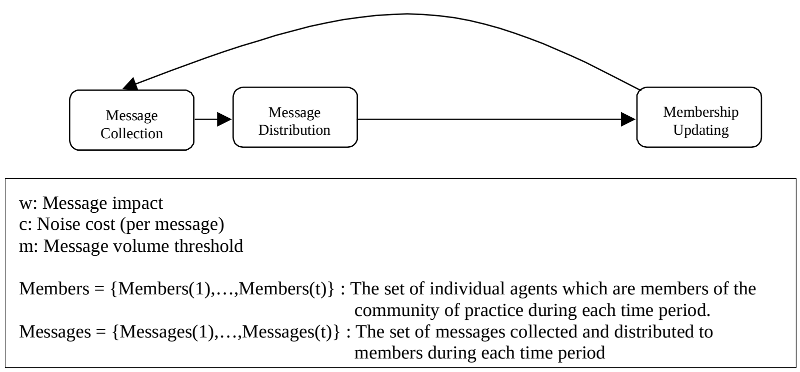

Figure 1: Community of Practice (COP) Processes, Parameters, and Internal State

A COP's membership is the set of individuals (members) who are exposed to and have access to the community's infrastructure. In a traditional COP, such as a professional society, membership might involving attending meetings, while in an online COP it may involve having access to an e-mail distribution list or website. In this analysis of the DYNACOP model all individuals are initially linked the to the COP and there is no inflow of new members. As individuals chose to leave the COP, its membership changes. Thus the COP's membership list (Members), implemented as a binary matrix in which $Members_{it}$ is 1 if individual *i* is a member during time period t and 0 otherwise, is an emergent consequence of the choices made by individual agents.

Communication occurs in the form of messages, discrete units that each have a single topic and, in the DYNACOP model, are distinguishable one another only in terms of that topic. Message topics are indicated by values drawn from the circular [0,1) interval which represents the COP's topic space. Members participate in the COP by contributing messages. Within a time period, each member makes a decision about participating, choosing whether to send a single message. The messages are then distributed, exposing each member to the communication activity. The message list is implemented as a matrix in which the entries ($Message_{it}$) are set as [0,1] values indicating a message topic if individual *i* sent a message during time period *t* and –1 otherwise. The set of non-negative values in a column of the Message matrix ($Message_t$) describes the communication activity within the COP during time period *t*. After the members have processed the messages, updated their evaluation of the COP, and made a decision about maintaining their membership, the COP's membership list is updated. Members are only removed from the list if they explicitly request termination of their membership. Consequently, an individual agent remains a member and sees all messages, until it takes explicit action to terminate its connection with the COP. The processes of message gathering, distribution, and membership adjustment (Figure 1) form a cycle that the COP completes once per time period.

Alternative infrastructures enable different forms of communication, which in turn affect the costs of communication and individuals' reactions to that communication. The DYNACOP model captures the effects of alternative communication infrastructures on COP development by varying the following parameters: *noise cost* (c), *message threshold* (m), and *message impact* (w). Noise cost (c) is the cost incurred by a COP member in processing a message that is outside the individual's interests (i.e. a valueless message). Noise cost is described relative to normalized signal benefit, a value that represents the maximum net benefit individuals receive from processing an interesting message. In the current implementation of the DYNACOP model, the normalized maximum signal benefit is fixed at 1. Hence, a noise cost value of 0.3 indicates that the cost incurred by an individual processing a non-interesting message is 30% of the net benefit derived from the most beneficial messages. Infrastructures that lower the per unit cost of communication, either by reducing overall communication costs or enabling smaller units of communication activity, lower the incremental noise costs incurred by COP members.

In addition to noise cost, the DYNACOP model uses message threshold (m) to represent the effects of alternative infrastructures on the costs and benefits of communication for individual COP members. Message threshold (m) is the number of messages of interest a member can receive in a day before the benefits of those messages are outweighed by the costs of processing them. The DYNACOP model posits that within a time period, on average, each successive message within an individual's area of interest (i.e. 'valuable' messages) contributes less benefit than the one before. This structure of declining marginal benefit, coupled with relatively stable costs of message processing, implies that marginal net benefit will also be declining. Thus, infrastructures that lower the costs of communication will increase the 'optimal' number of messages by shifting the point at which the incremental net benefit becomes negative.

Alternative infrastructures also affect the impact that messages have on individuals' beliefs about the COP. In the DYNACOP framework, the message impact (w) parameter is used to represent this effect. In addition to their explicit substantive content, messages implicitly provide information about the COP itself. This information plays a role in members' recurrent decisions to continue or end membership. Message impact, or the weight given to information about the COP derived from a single unit of communication, is represented by a [0,1] value. COP infrastructures that enable lower cost communication result in lower message impacts because the cost of "sampling" a COP is lower; that is, individuals assign less weight to the information provided by any single message because of the lower costs of processing each one.

Although a COP is a distinct social entity, it also has emergent features. Communication activity results not from an internal COP process but rather from actions taken by members. Membership is maintained not by dictate of the COP, but by the choices of independent individuals. COP members contribute messages, process other messages, and, based on their perceptions and expectations, decide whether to continue or maintain their membership. Thus the DYNACOP framework also includes a representation of individuals (Figure 2).

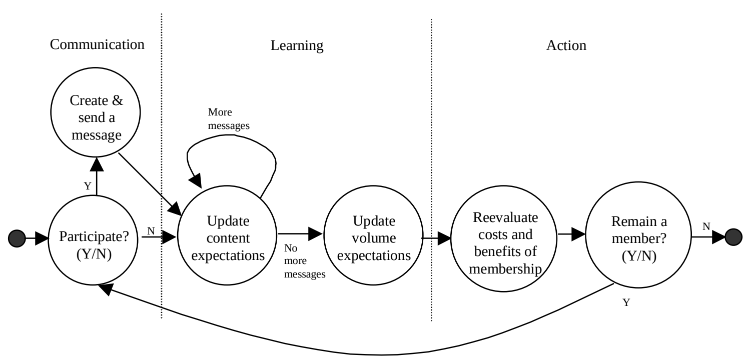

CE_{it}: Individual i's expectations at time t regarding the content of the COP's communication activity
VE_{it}: Individual i's expectations at time t regarding the volume of the COP's communication activity within a single time period
pp_i: Individual i's participation probability
[INT_{iL}, INT_{iH}]: The range of individual i's content interests

Figure 2: Individual Agent Processes, Parameters, and Internal State

This three stage cyclic model of communication, learning, and action is based on those proposed by Turner (1988) and Carley (1991) in their studies of social structure development. Individuals participate in a COP by constructing messages that are then sent to the other members.

Following prior work on participation in discussion groups (Skvoritz 1988) and information sharing in decision-making teams (Wittenbaum and Stasser 1996), each individual's message contribution behavior is modeled as the result of an independent stochastic process. An individual chooses to contribute a single message in a time period with a given probability (pp_i), which may vary between individuals but does not change over time. Upon choosing to participate, an individual creates a message, selecting a topic from his interests. In the current implementation, a message is constructed by selecting a random value from a uniform distribution of the individual's interests as described by the continuous interval [INT_L, INT_H] within the COP's topic space. The message is then passed to the COP for distribution.

An individual's relationship with the COP develops over time (Moreland and Levine 1982). As members are exposed to activities within the COP, their expectations about the content (CE_{it}) and volume of activity (VE_{it}) change. Content expectations (CE_{it}) are [0,1] values that represent individuals' expectations in time period *t* regarding the probability that future COP messages will be of interest. Volume expectations (VE_{it}) are positive values representing individual *i*'s expectations at time *t* about the message volume in the COP in a single time period.

In addition to their primary function of transferring topical information, COP messages also serve a secondary purpose of providing information about the community's topical focus. As a member, an individual is exposed to COP messages. Based on this exposure, individuals adjust their expectations about the community's topical focus and level of activity. This activity is implemented as a reinforcement process (Hunter, Danes, and Cohen 1984). The change in an individual's content expectations (ΔCE_{it}) is described by:

$$\Delta CE_{it} = rw[CE_{i(t-1)}][\ 1 - CE_{i(t-1)}]$$

where *w* is the message impact parameter, which is inherited from the COP agent, and *r* denotes the individual's reaction to the message (r = 1 if the message is of interest, and r = –1 otherwise). An individual deems a message interesting when its topic falls within the arc defined by the individual's interest parameters (INT_{iL} and INT_{iH}). When multiple messages are distributed in a time period, the change in an individual's content expectations is determined separately for each message. After all messages for a *t* have been received, volume expectations for the next time period (VE_{it}) are set to the observed mean message volume for all previous time periods.

Individuals' expectations about the future activities of the COP are an important factor in their assessment of the benefit of continued membership (Moreland and Levine 1982). After adjusting their beliefs in response to being exposed to a set of messages, individuals assess the expected costs and benefits of continued membership. This assessment takes into account the expected benefit from messages that are of interest to the individual, a benefit that is subject to limits. Individuals have strict limits on the time available to them. As more time is 'spent' processing messages, the remaining time is more valuable. Thus as the number of interesting messages distributed within a single time period increases, the incremental net benefit of an additional interesting message in the same time period is lower. The assessment also considers the costs incurred as a result of processing uninteresting, or noise, messages. Expectations of the content and volume of activity (CE_{it}, VE_{it}) and the known costs of processing messages within the collective's communication infrastructure (m,c) are combined in each individual's assessment of the expected costs and benefits of continued membership. This assessment is described by:

$$E_{it} = E(CE_{it},VE_{it}; c,m) = (-1/2m)(CE_{it}VE_{it})^2 + (CE_{it}VE_{it}) – c((1- CE_{it})VE_{it})$$

The first two terms $[(-1/2m)(CE_{it}VE_{it})^2 + (CE_{it}VE_{it})]$ indicate the total expected net benefit due to messages which are expected to be of interest. The final term $[– c((1- CE_{it})VE_{it})]$ is the expected costs due to noise messages. Individuals' assessment of the expected costs and benefits determines their willingness to remain in the COP. If the expected net benefit is positive, the individual will choose to remain a member; otherwise she will choose to terminate membership.

The individual (Figure 2) and COP (Figure 1) processes are different components of a social system represented by the DYNACOP model which contains a single community of practice with an initial membership of N individuals (Figure 3).

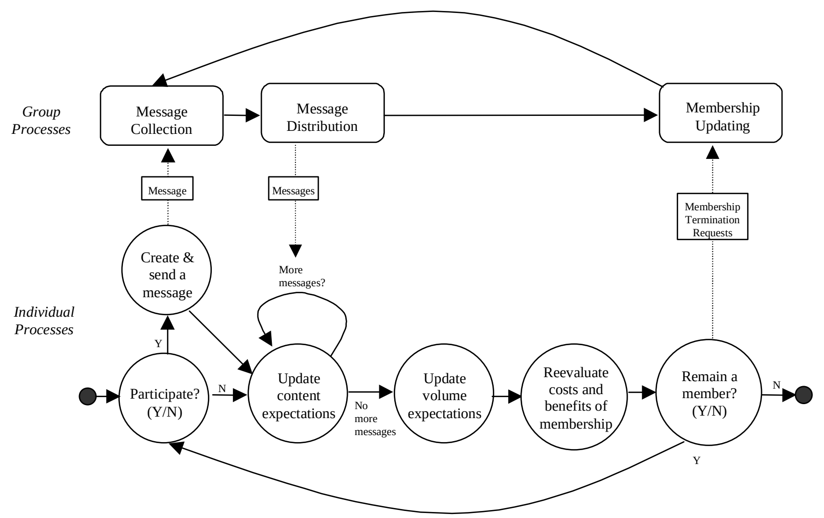

N: Initial membership size (i.e. number of individual
c,w,m : COP infrastructure parameters
ppDistribution: The distribution of participation probabilities among the individual agents
IntDistribution: The distribution of interests among the population of individual agents

Figure 3: DYNACOP Social System Processes and Parameters

Within the social system individuals and the COP interact, and from those interactions arise the emergent structural features of the COP. The parameters of the DYNACOP model are the infrastructure parameters of the COP, noise cost (c), message impact (w), and message threshold (m), and the exogenously determined characteristics of the population of individual agents, the distribution of participation probabilities (pp_i) and the distribution of individual interest ranges.

The current implementation of the DYNACOP framework is a lockstepped computational model.   All individuals perform an action before any one agent moves to the next step in the process. Within a given step, the individuals take action in a fixed order. After the individuals have completed a step, the COP performs any relevant activities. All individuals make a choice about contributing messages, create messages, and send them to the COP before the messages are distributed. Likewise, all individuals decide about continuing membership before the COP updates its membership list and starts accepting new messages.

Calibration and Validation of the DYNACOP Model

Model calibration is the process of adjusting a computational model to reflect the features of empirical data (Carley 1996). The result of calibration is a set of parameters which best represent one instance of the phenomena of interest, providing a conceptual reference point for validation, analysis, and interpretation of the computational model. Model validation is the process of comparing the results of a computational model with empirical data for the purposes of testing the model (Carley 1996). Unlike calibration, in which the parameters of the model are adjusted to the most appropriate settings, during validation the proposed parameters and processes are taken as given, the model is run, and the outcomes are compared with empirical data. The results of validation provide information about how well, and under what conditions, the computational model accurately represents the intended phenomena.

Calibration and validation of the DYNACOP framework was performed by comparing the results of the model with empirical data collected from a random sample of unmanaged e-mail based Internet listservs. These communities utilize Internet-based e-mail and centralized mailing lists to enable members to broadcast text-messages to one another. From a population of more than 70,000 listservs, a random sample of 192 was selected. The sample was stratified with respect to topic to ensure that it spanned a range of interest areas and member populations. For a 130 day period, between July 23, 1997 and November 30, 1997, message archives and membership data were collected for each listserv. From this data, measures of daily communication volume and percentage membership loss were constructed. Mean daily communication volume was calculated by determining the total number of messages distributed to the members during the observation period and dividing that by the number of days (i.e. 130). Percentage membership loss was determined by counting the number of people who left the listserv during the observation period and normalizing it by the number of members present on the first day of the observation period (for additional details of the sample selection and data collection procedures see Butler 1999).

Model Calibration

The DYNACOP model was calibrated by systematically varying parameters and comparing the model predictions with the empirical data for the rate of member loss (structural change) and the average daily message volume (communication activity level). From the sample of listservs, 100 were randomly selected to be the calibration sample; the remaining 92 were used as the validation sample. Calibration was performed in a series of sessions. Each session involved simulating 100 COPs and comparing the resulting distribution of member loss and communication activity with data from the listservs in the calibration sample. In each calibration session, one or more of the model settings were modified to better reflect the features of the observed collectives. Model parameters that could be directly observed in the empirical data, such as initial membership size, were calibrated in early sessions. Parameters that could not be directly observed, such as noise cost or message impact, were estimated in later sessions, after the observable parameters were set. Also, to better represent the initial state of the listservs, all of which were at least four months old, each calibration session included two phases. The first phase, a 100 time period initialization phase, was run to represent a listserv's prior history. The second phase, with 130 time periods, was then run and the results compared with data from the calibration sample. The multi-session calibration process provided a baseline for analysis of the DYNACOP model by identifying the set of parameters that best fit the listservs (Table 1).

Model Setting	Calibrated Value
N	Log-normally distributed [LN(4.17,1.34)]
c	1.0
w	0.02
m	5
ppDistribution	Ratio of participants (i.e. pp > 0) to members is chosen from an exponential distribution with m = 0.17; All agents labeled as participants have the same participation probability – a value which drawn from a log-normal distribution [LN(-4.03,0.53).
Interest range distributions	Individuals' interest range length is chosen from a uniform distribution between 0 and MaximumInterestRange. MaximumInterestRange for the COP is chosen from a uniform distribution between 0.25 and 0.75.
Initial CE_0 distribution	Uniformly distributed between 0.5 and 1
Initial VE_0 distribution	Fixed @ 1

Table 1: DYNACOP Parameters for E-mail Based Internet Listservs

The calibration settings provide a reference point for validation, analysis, and interpretation of the DYNACOP model.

Model Validation

The DYNACOP model was validated using data from the 92 listservs not used for calibration. The model was validated using matched analysis (Law and Kelton 1991), a process that involves simulation of specific cases and comparison of model results and observed behavior with a paired t-test. The empirically observable model settings, initial collective size (N), and participation structure (as described by the participation ratio and participation probability), were set based on the data from a single listserv in the validation sample. The infrastructure parameters (noise cost (c), message weight (w), and message threshold (m)) and initial conditions (CE_0 and VE_0) were set based on the results of model calibration (Table 1). Ten model runs were performed using these settings. For each run, the remaining unobservable model setting, variation in the collective's member interest ranges (determined by MaximumInterestRange), was selected from a uniform distribution between [0.25,0.75]. The mean percentage membership loss and mean daily communication volume for the set of 10 simulated COPs were recorded as the predicted structural outcomes. This process was repeated for each case in the validation sample.

Comparison of the mean predicted and observed outcomes with paired t-tests indicates that the DYNACOP model accurately predicts community membership loss (H_0: $\mu_{predicted} = \mu_{observed}$; $p = 0.263$). There is, however, a significant difference between the predicted and observed values for communication activity (H_0: $\mu_{predicted} = \mu_{observed}$; $p < 0.01$) implying that the model under-predicts communication activity ($\mu_{predicted} = 0.408$ and $\mu_{observed} = 0.795$). OLS regression analysis of the membership loss error (Table 2) indicates that the model's predications are most accurate for communities with lower membership loss, smaller membership and relatively less communications activity – a common type of online community (Butler 1999).

	Dependent Variable			
	Membership Loss Error		Message Volume Error	
	Unstandardized	Standardized	Unstandardized	Standardized
Intercept	0.051***		0.0251	
Initial Size	0.00012**	0.209**	0.0013***	0.242***
Observed Message Volume	0.025***	0.301***	-0.7585***	-1.011***
Observed Membership Loss	-0.818***	-0.882***	-	-
Adjusted R^2	0.572		0.955	
p < 0.05: * p< 0.01: ** p < 0.001 ***				

Table 2: OLS Regression Analysis of Model Error

Similarly, OLS analysis of the difference between predicted and observed levels of communication activity indicates that the DYNACOP model's predications are more accurate for COPs with lower levels of communication activity. This error structure is a consequence of the approximation used to model participation in the current implementation of the DYNACOP framework. Overall, the validation and error analysis results indicate that the DYNACOP framework provides a model of member loss and communication activity in which is applicable for communities of practice in the range of sizes and activity levels that are common in both online (Butler 1999) and traditional settings (McPherson 1983a).

Development of Communities of Practice

Before considering the impact of alternative communication infrastructures on the dynamics of communities of practices, we will first examine how a COP develops within the DYNACOP framework. As members are exposed to communication within a COP, their beliefs about the community change. Developing beliefs alter individuals' assessments of the benefits of being part of the community, changing their level of commitment and possibly causing them to end their membership (Moreland and Levine 1982). Member movement changes the composition and size of a COP's membership, which is reflected in the development of its aggregate interests and the focus of its communication activity. These changes, in turn, affect the development of remaining members' beliefs.

The simultaneous development of member commitment and the structural features of a COP's membership are a result of the cycle of individual and structural change processes. As the following example illustrates, these development processes can be seen in a COP modeled with the DYNACOP framework[i]. The mean commitment level among members converges to a stable value (Figure 4).

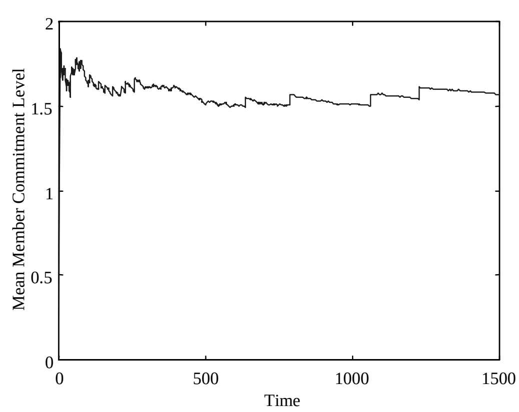

Figure 4: Mean Member Commitment

Changes in the mean level of member commitment within a COP arise from two processes. The abrupt shifts are the result of peripheral members' terminating their membership, while the steady progression of the mean toward a stable point is the result of core members' commitment converging on a common value. Jumps in the mean level of member commitment occur as individuals who have lower evaluations of the COP, and hence lower commitment, choose to end their membership. These abrupt shifts occur in tandem with the convergence of remaining members' commitment to a common level, a result of the development of strong beliefs about the positive value of future activity within the COP.

The combination of individual and structural change also results in several emergent phenomena that are consistent with prior conceptual models of social collective development. For example, the minimum level of member commitment, another indication of a COP's development, is subject to sudden shifts (Figure 5).

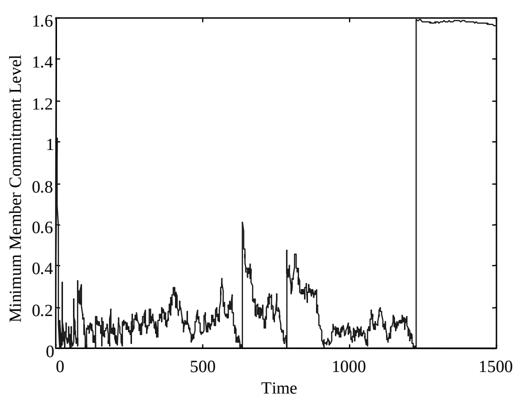

Figure 5: Minimum Commitment Level with the Collective

This shift in the minimum level of commitment may underlie significant changes in the operation of the collective, similar to the development phases proposed by Tuckman (1965) or the transitions observed by Gersick (1988, 1989).

Another feature of social structure development seen within the DYNACOP framework is the emergence of an interest focus within a COP. Although it is common to refer to a COP's 'interests', especially in discussions of online social structures (Baym 1993, Kollock and Smith 1996), even in formally managed settings, collective communication activity is actually the aggregate result of members' individual choices. The focus of interests of a COP is not a monolithic construct, but an emergent one that arises from the members' actions. As a COP's membership develops, the distribution of interests among the members changes (Figure 6).

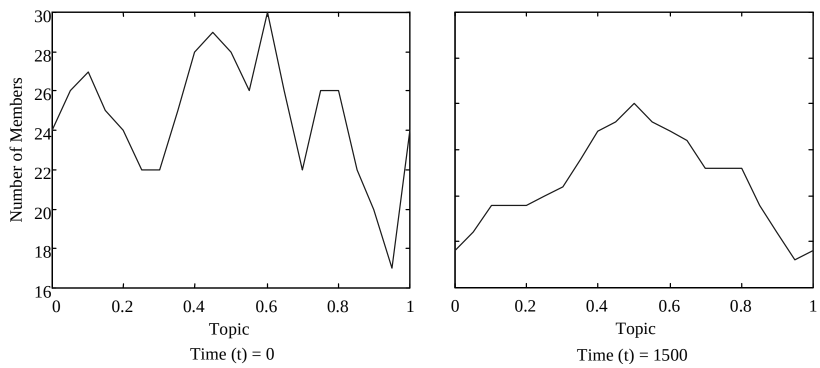

Figure 6: Development of a Collective's Interest Distribution

Ultimately a COP's membership reaches a stable point. When stability is reached the distribution of interests[ii] is peaked (Figure 6). Thus without a formally defined interest specification, a COP still develops an interest focus as a result of member and structural change processes.

Alternative Infrastructures and the Development of Communities of Practice

One of the most often discussed features of network infrastructures is their ability to reduce communication costs (Sproull and Keisler 1991). By reducing the costs of message transmission, networks can enable messages to be sent that in more costly infrastructures, such as group meetings or print publications, would have gone undistributed. Computer-mediated systems also change the way people process communication, affecting both the incremental costs of receiving a desirable message and the costs incurred as a result of noise. Network infrastructures also reduce many economies of scale, allowing communication to take place in smaller units. Rather than having a multi-day conference or sending a multi-page newsletter, with a network infrastructure, communication can occur in more frequent, shorter messages.

Different infrastructures lead to different communication costs for the members of a community of practice. Different costs affect the development of a COP in several ways. Communication costs and structures affect the development of individuals' perceptions by altering the impact of individual units of communication activity on the individuals' beliefs about the community. If communication activity is large-grained and expensive, the impact of any given unit on a member's beliefs will be greater. Lower processing costs also affect COP development by altering individuals' assessments of expected net benefits of continued membership. This, in turn, may change when, or if, individuals choose to end their membership. Thus changing communication costs have the potential to affect collective development by affecting both individual and structural change processes.

To analyze the DYNACOP models predictions regarding the impact of alternative infrastructure on COP development, a virtual experiment was performed. The experimental conditions were created by systematically varying the communication parameters, noise cost (c), message weight (w), and message threshold (m). The remaining model settings and the initial conditions were set as determined during calibration (c.f. Table 1). Noise cost was set at 0.33, 1.0, or 3.0. Message weight was set at 0.005, 0.02, or 0.1. Message threshold was set at 2, 5, or 8. The values were chosen to represent a range of communication infrastructures. To anchor the analysis, the experimental settings included as their center points the values identified during calibration as most appropriate for e-mail based Internet listservs.

For each condition, one hundred COPs were modeled for one model 'year'. To simulate a year, runs were 365 time periods long, in addition to an initialization stage of 100 time periods. For each condition 100 COPs were simulated and measures of membership size and intermediate stability were examined. Size, measured in terms of the number of members, is an important structural feature of communities of practice and other social structures (Butler 1999).

Intermediate stability is linked to the likelihood of individuals leaving a COP. Individuals end their membership when they expect the costs of continued membership to outweigh the benefits. COPs can be seen as regularly experiencing shocks, both internal and external, that can alter the membership of the community. A COP is more stable if its membership is less likely to leave in the face of these events. Examination of simulation results indicate that in the DYNACOP model differences in individual members' evaluation of a COP were due primarily to differences in their expectations about the content of activity (Figure 2). The lower members' expectations about future message usefulness, the lower their assessment of future membership benefits will be, and the higher the chances that they would leave if a shock occurs. Intermediate stability can thus be measured in terms of the minimum content expectations among a COP's members, with higher minimum content expectations indicating greater intermediate stability.

The COP features, size and intermediate stability, were measured for each of the simulated COPs and differences between the conditions were examined with a series of ANOVA analyses. Infrastructures with lower costs and message weights are predicated to result in larger COPs. Communities operating in infrastructures with lower relative noise costs see less membership loss, and as a result they are larger than those in which the relative cost of processing noise is higher ($F = 42.514$; $p < 0.001$) (Figure 7a).

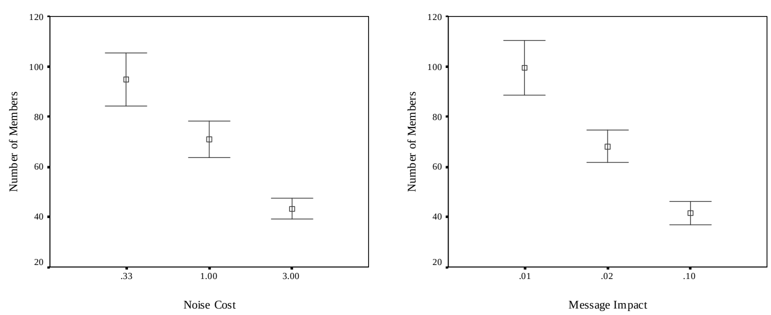

(a) (b)

[95% confidence intervals for collective size after 365 time periods]

Figure 7: Effects of Communication Features on Community Size

The effect of relative noise cost on COP size is a consequence of altering the minimum acceptable signal to noise ratio. When faced with lower relative costs of processing noise, individuals are willing to tolerate a higher proportion of unwanted messages. Thus the threshold at which individuals choose to end their membership is lower. This slows down the rate at which members filter out, resulting in larger communities of practice. Infrastructures in which the impact of individual messages on member beliefs is lower are also larger ($F = 54.220$; $p < 0.001$) (Figure 7b). Lower message impacts reduce the rate at which individuals' beliefs about a COP change. This, in turn, slows down the rate at which members either leave the community or become fully committed to the community.

There is also an interaction between noise cost and message impact (Figure 8). In infrastructures with lower noise costs, the impact of decreasing message weights on COP size is greater.

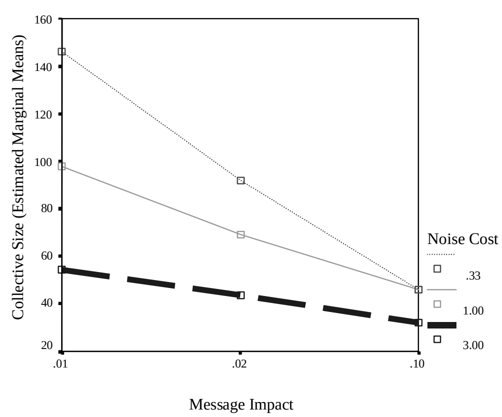

Figure 8: Interaction of Message Impact and Noise Cost on Community Size

This interaction indicates that message impact and noise cost are expected to have effects on the development of COP stability that are non-additive.

Infrastructures that have lower noise costs and message impacts also result in COPs that are less stable. The DYNACOP model predicts that communities of practice operating in infrastructures with lower noise costs will be incrementally less stable ($F = 2761.603$; $p < 0.001$)

(Figure 9a).

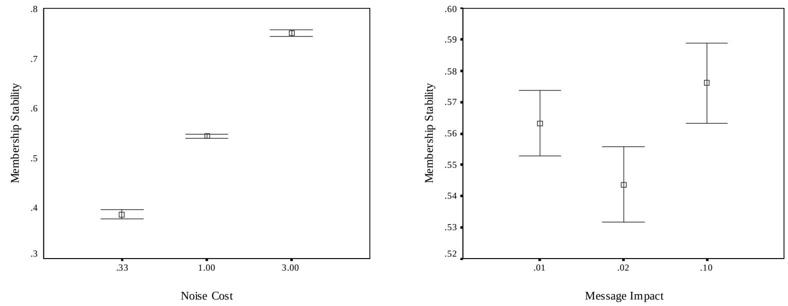

(a) (b)

[95% confidence intervals for collective stability after 365 time periods]

Figure 9: Effects of Communication Features on Community Stability

However, the effects of message weight on community stability are more complex (Figure 9b). Based on the trajectory of their belief development, individuals can be classified in one of three post hoc categories: leavers, committed core, and peripheral members. Leavers, whose interests do not match the developing interests of the COP, are characterized by a steadily declining content evaluation trajectory. Given enough time, these individuals all leave. The committed core consists of those members who have consistently high and increasing evaluations of the communication content. Their content evaluation trajectories are steadily increasing. Peripheral members are characterized by evaluation trajectories that first decrease, then increase. Early on, peripheral members' interests are a poor match with those of the community as whole. As a result, their content evaluation decreases over time. However, as individuals leave and the COP's interest focus develops, the match with peripheral members' interests improves and with it, their evaluations of the benefits of membership. Given enough time, peripheral members' evaluations increase to the level of that of the committed core.

Early in the development of a COP, the members all have relatively high expectations. As time passes, leavers and peripheral members' expectations of the content drop while those of the committed core increase. As the leavers work their way out, their commitment drops and with it the overall stability of the COP. Thus, early on leavers undermine a community's intermediate stability. Furthermore, after these individuals leave, peripheral members' temporarily low levels of commitment serve to keep stability low. Then as the peripheral members' evaluations of the COP recover, their expectations increase and with them the stability of the community. Reduced message impact decreases the rate at which individuals' beliefs change, slowing down the rate at which leavers exit, the committed core forms, and peripheral members' evaluations fall and recover. As a result, reducing message impact alters the rate at which a COP's stability develops.

The non-linear pattern seen in Figure 9b is the result of COPs reaching different stages of development in the simulated year. In the low message impact condition, a reduced rate of belief change results in many leavers having not yet worked their way out of the communities. In the high impact condition, stability is higher because the leavers have ended their membership and the peripheral members have recovered. In the moderate condition, the simulated COPs have lost the leavers, but because the rate of belief change is reduced, there has not yet been time for the peripheral members to recover. As a result, stability in these communities is lower than the low impact infrastructures, because the peripheral members have had time to 'reach bottom', and lower than the high impact infrastructures, because they have not had enough time to recover.

Discussion

The DYNACOP framework draws from prior research on belief change and social structure to model the effects of alternative infrastructures on the intertwined processes of member and structural development in communities of practice. Over time, the processes of individual attitude change and membership movement combine to shape both a COP's true interests and member perceptions of that focus. Communication features, such as message impact and noise cost, significantly affect both member and structural development processes in communities of practice. In contexts with high communication costs, the costs of noise messages and the impact of messages on belief formation will be greater. When the impact of a message is high, individual beliefs form rapidly; when noise costs are high, members will leave rather than tolerate communication activity that falls outside their interests. Consequently, in COPs with high cost infrastructures, such as face-to-face meetings, both member and structural development proceed more quickly, resulting in smaller, more stable communities. In contrast, by providing lower cost communication, networked technologies (such as e-mail or Web conferencing systems) slow down the individual and structural change processes that underlie community development. Lower communication costs, as well as smaller units of communication, reduce the pressure on individuals to reach confident conclusions about continuing membership. Under such conditions, then, individuals remain in the community who would have left quickly a high communication cost situation. As a result, networked COPs are expected to be larger, more diverse, and less stable, having lower minimum and average member commitment, than communities that rely on traditional communication technologies, such as face-to-face meetings or print.

These findings suggest that COPs that make use of networked infrastructures will tend to be less "community-like" than COPs that use traditional face-to-face communication. Within the topical bounds of individual COPs that use networked infrastructures, an interest focus will be less likely to emerge. Hence, networked COPs will be less likely to spontaneously develop a well-defined community identity. In addition, members will tend to be less certain about their beliefs about the community, less positive in their evaluations, and less committed. This arises because, while there continues to be a core of individuals who are strongly committed, the proportion of marginally interested and committed members will be significantly higher.

The notion that networked communities will differ significantly from traditional social structures is not a unique feature of the DYNACOP model; previous research has pointed to this observation as well. Rather, what distinguishes the DYNACOP model from prior work in this area is its explanation of the differences between networked and traditional communities. Previous discussions about the differences between these two types of communities have largely explained these differences as a result of drawbacks to networked communication infrastructures. During the past decade, many researchers have considered the possibility that online communities find it difficult to attain certain aspects of traditional communities (Roberts 1998, Sproull and Keisler 1991). The primary argument for this claim has been that because networked infrastructures lacked features such as social presence and the ability to transmit multiple cues, they are poorly suited for supporting the type of communication required to develop member commitment and community identity. Thus these studies proposed that differences between networked and traditional communities result from networked communities' inability to provide aspects of communication that are present in traditional communities.

While prior discussions have focused on "missing" features of networked infrastructures as the cause of the community development challenges, the DYNACOP model suggests that these differences result from a reduction in communication costs, a positive feature of networked infrastructures. Lower communication costs reduce the motivation for individuals to develop firm beliefs about and a strong commitment to the COP. They also slow down the matriculation of marginal members – a social process that plays a role in the formation of community identity and higher overall commitment among members. The DYNACOP model thus indicates that many of the challenges faced by networked communities result not from deficiencies in the communication infrastructures they employ, but, seemingly paradoxically, as a consequence of one of the main advantages of those infrastructures. This alternative explanation of the differences between networked and traditional COPS runs counter to previous discussion of these issues, and future research on the development of communities, both within and outside organizations should further consider these very different explanations in order to better inform designers of community infrastructures.

Another implication of the DYNACOP model is that the impact of lower cost infrastructures on behavior will vary among individuals. Individuals whose interests do not coincide with the emergent interest focus will still leave the COP but will do so much more slowly in the presence of low communication costs. Other members, who interests are on the margin of the emerging focus, will remain in the COP longer, and hence will be more likely to ultimately become part of the core of the COP. Individuals who in high cost situations would be part of the core will remain committed in the presence of lower cost infrastructure. Identifying differential technology impacts are significant because they can help explain how a supposedly "uniform" change (i.e. lowering of communication cost) can result in the complex changes in social structure which have been observed in networked COPs. It also suggests that caution should be used when relying on individual reports to characterize the impact of new infrastructure.

Furthermore, the differential impact suggested by the DYNACOP model may be at the root of the tendency to focus on the adoption of new infrastructures as an issue primarily of mechanistic efficiency. Because of their high commitment and expected benefits, core members are most likely to be the ones making choices about a COP's infrastructure. However, the DYNACOP model suggests that core members' behavior with respect to COP membership will be the least affected by the adoption of lower cost infrastructure. Hence, they are likely to see the choice of infrastructure as decision affecting *how* members will participate in the COP, not *whether* people will participate. However, the DYNACOP model suggests that doing so is misleading because it ignores the consequences of the infrastructure on the development of the community as a whole. Reducing the cost increases the size and diversity of the membership by significantly altering the behavior of individuals who are tangentially interested and marginally committed. This, in turn, can significantly change the nature of the community. Thus while core members tend to see infrastructure design as a question of communication mechanics, they also need to be aware that the differential impact of lower communication costs can significantly impact the community's development.

Conclusion

The computational model considered in this paper is presented for purposes of theory development. With this in mind, its validity should be judged in terms of the balance between the purpose (theory generation), the relationship of the model features to that purpose, and the use of empirical data (for calibration and validation) relative to that purpose (Burton and Obel 1995). Although inclusion of the additional features of a COP within the DYNACOP framework might increase the "realism" of the model, the features included in the current implementation support the primary purpose of integrating member development, structural change, and communication costs into a model of dynamic communities of practice. Further development of the participation model, explicit consideration of the inflow of new members, and examination of the operation of larger systems of multiple communities would build on the foundation presented here to further the study of the role of communication infrastructures in the development of effective, beneficial communities of practice.

The DYNACOP framework presented here suggests that integration of individual and structural change processes provide a theoretical foundation for examining the effects of alternative infrastructures on developing communities of practice, as well as presenting a foundational process model of COP development. This theory focuses on the dynamic aspects of social structures which play an important role in the flow of information within and between organizations (Goodman and Darr 1998, Von Hippel 1988). It also considers the impact of alternative technologies on these processes. As organizations, both public and private, invest in information technology with the goal of facilitating information flow, it becomes increasingly important to understand how features of the technology and the social processes of COP development interact.

As network infrastructures such as the Internet become widely available, there is a tendency to assume that providing the ability to efficiently communicate will itself necessarily support and encourage interaction. Prior research on computer-mediated communication has focused primarily upon demonstrating the many ways that individuals *can* interact effectively through networked infrastructures. This research has thus served to define a universe of possibilities for designers and developers of networked infrastructures. However, much of this early work has failed to recognize that, while individuals *can* behave in many ways, how they *will* behave is significantly affected by social structures. While the first-order effects of new technologies on mechanistic efficiency may be the most visible, it is incorrect to assume that they are the most important consequences of introducing new communication infrastructures (Sproull and Kiesler 1991). However, without dynamic models that take into account individual, structural, and technological features of social environments, our ability to understand, and ultimately predict the impact of new communication technologies is limited. By combining empirical data and computation modeling, the DYNACOP framework provides a foundation for future research about how technology has, and will continue, to affect communities of practice both within and outside organizations.

References

Allen, T., *Managing the Flow of Technology*, MIT Press, Boston, 1984.

Baym, N., "Interpreting Soap Operas and Creating Community: Inside a Computer-Mediated Fan Culture," *Journal of Folklore Research*, 30 (1993), 143-176. [Also appears in S. Kiesler (Eds.), Culture of the Internet]

Blau, P. M., *Exchange and Power in Social Life*, Wiley, New York, 1967.

Burton, R. M. and B. Obel, "The Validity of Computational Models in Organizational Science: From Model Realism to Purpose of the Model," *Computational and Mathematical Organization Theory*, 1, 1 (1995), 57-71.

Butler, B. S., *The Dynamics of Cyberspace: Examining and Modeling Online Social Structure,* Ph.D. Dissertation, Carnegie Mellon University, Pittsburgh, PA, 1999.

Brown, J. S., and P. Duiguid, "Organizational Learning and Communities of Practice: Toward a Unified View of Working, Learning and Innovation," *Organizational Science*, 2 (1991), 40- 57.

Carley, K. M., "A Theory of Group Stability," *American Sociological Review*, 5-6 (1991), 331-354.

Carley, K. M., "Communicating New Ideas: The Potential Impact of Information and Communication Technology," *Technology in Society*, 18, 1 (1995a), 1-12.

Carley, K. M., "Communication Technologies and Their Effect on Cultural Homogeneity, Consensus, and the Diffusion of Ideas," *Sociological Perspectives*, 38, 4 (1995b), 547-571.

Carley, K. M. and K. Wendt, "Electronic Mail and Scientific Communication," *Knowledge:*

Creation, Diffusion, and Utilization, 12, 4 (1991), 406-440.

Carley, K. M. "Validating Computational Models", CASOS Working Paper, Carnegie Mellon University, Pittsburgh, PA, 1996.

Etzioni, A., *Modern Organizations*, Prentice Hall, Englewood Cliffs, NJ, 1964.

Epstein, J. M. and R. Axtell, *Growing Artificial Societies,* Washington, DC, Brookings Instituion Press, 1996.

Fischer, B.A., "Decision emergence: Phases in Group Decision-Making" *Speech Monographs*, 37 (1970), 53-66.

Gersick, C. J.G., "Time and Transition in Work Teams: Toward a New Model of Group Development," *Academy of Management Journal*, 31, 1 (1988), 9-41.

Gersick, C. J.G. "Marking Time: Predictable Transitions in Task Groups," *Academy of Management Journal*, 32, 2 (1989), 275-309.

Goodman, P. S. and E. D. Darr, "Computer-Aided Systems and Communities: Mechanisms for Organizational Learning in Distributed Environments," MIS Quarterly, 22,4 (1998), 417-440.

Goffman, E., *Presentation of Self in Everyday Life*, Doubleday/Anchor, New York, 1959.

Hannan, M. T. and J. Freeman, "The Population Ecology of Organizations," *American Journal of Sociology,* 82 (1977) , 929-940.

Hare, A.P. *Handbook of Small Group Research (2^{nd} ed.),* Free Press, New York, 1976.

Huberman, B A. and T. Hogg, "Communities of Practice: Performance and Evolution", *Computational and Mathematical Organizational Theory*, 1, 1 (1995), 73-92.

Hunter, J. E., Danes, J. E., and S. H. Cohen, *Mathematical Models of Attitude Change (Volume 1),* Academic Press, New York, 1984.

Kaufer, D. S. and K. M. Carley, *Communication at a Distance: The Influence of Print on Sociocultural Organization and Change*, Lawrence Erlbaum, Hillsdale, NJ, 1993.

Kollock, P. and M. Smith, M. "Managing the Virtual Commons: Cooperation and Conflict in Computer Communities," in S.C. Herring (Eds.) *Computer Mediated Communication: Linguistic, Social, and Cross-Cultural Perspectives*, Philadelphia: John Benjamins Publishing, 1996.

Kraut, R. E. and P. Attewell, "Media Use in a Global Corporation: Electronic Mail and Organizational Knowledge," in S. Kiesler (Ed.) *Culture of the Internet*, Lawrence Erlbaum, Mahwah, NJ, 1997.

LaCoursiere, R.B., *The Life Cycle of Groups: Group Development Stage Theory*, Human Sciences Press, New York, 1980.

Lave, J. and E. Wegner, Situated Learning: Legitimate Peripheral Participation. Cambridge University Press, 1991.

Law, Averill M. and W. D. Kelton, *Simulation Modeling and Analysis (2nd ed.)*, McGraw-Hill, New York, 1991.

McGrath, J. E., *Groups: Interaction and Performance*, Prentice-Hall, Englewood Cliffs, NJ, 1984.

McGrath, J. E. and A. B. Hollingshead, *Groups Interacting with Technology: Ideas, Evidence, Issues and an Agenda*, Thousand Oaks, CA, Sage, 1994.

McPherson, J. M., "The Size of Voluntary Organizations," *Social Forces*, 61, 4 (1983a), pp. 1045-1064.

McPherson, J. M., "An Ecology of Affiliation," *American Sociological Review*, 48 (1983b), 519-532.

McPherson, J. M., "Evolution in Communities of Voluntary Organizations," in. J.V. Singh (Ed.). *Organizational Evolution New Directions*, Sage, Newbury Park, CA, 1990.

McPherson, J. M. and T. Rotolo, "Testing a Dynamic Model of Social Comparison: Diversity and Change in Voluntary Groups," *American Sociological Review*, 61 (1996), 179-202.

Moreland, R. L. and J. M. Levine, J. M., "Socialization in Small Groups: Temporal Changes in Individual-Group Interactions," *Social Psychology*, 15 (1982), Academic Press, New York.

Moreland, R. L., Hogg, M. A., and S. C. Hains, "Back to the Future: Social Psychological Research on Groups," *Journal of Experimental Social Psychology*, 30 (1994), 527-555.

Poole, M. S., "Decision Development in Small Groups I: A Comparison of Two Models," *Communication Monographs*, 48 (1981) , 1-24.

Poole, M. S. "Decision Development in Small Groups II: A Study of Multiple Sequences of Decision Making," *Communication Monographs*, 50 (1983), 206-232.

Roberts, T. L., “Are newsgroups virtual communities?” Proceedings of CHI'98, Los Angeles, CA, 1998.

Scheidel, T. and , L. Crowell, "Idea Development in Small Discussion Groups," *Quarterly Journal of Speech*, 50 (1964), 140-145.

Simmel, G. *Conflict and the Web of Group Affiliations*, Translated by K. Wolff and R. Bendix. Free Press, Glencoe, IL, [1908], 1955.

Skvoretz, J., "Models of Participation in Status-Differentiated Groups," *Social Psychology Quarterly*, 51, 1 1988, 43-57.

Sproull, L. and S. Kiesler, *Connections: New Ways of Working in the Networked Organization.* MIT Press, Boston, 1991.

Tuckman, B., "Developmental Sequence in Small Groups," *Psychological Bulletin*, 63 (1965), 384-399.

Tuckman, B. and M. Jensen, "Stages of Small-Group Development," *Groups and Organizational Studies*, 2 (1977), 419-427.

Turner, J. H., *A Theory of Social Interaction,* Stanford University Press, Stanford, CA, 1988.

Von Hippel, E., *The Sources of Innovation,* Oxford University Press, New York, 1988.

Wittenbaum, G. M. and G. Stasser, "Management of Information in Small Groups," in J. L. Nye and A. M. Brower (Eds.) *What's Social About Social Cognition? : Research on Socially Shared Cognition in Small Groups*, Sage, Thousand Oaks, CA, 1996.

Worschel, S., "Emphasizing the Social Nature of Groups in a Developmental Framework," in J. L. Nye and A. M. Brower (Eds.) *What's Social About Social Cognition? : Research on Socially Shared Cognition in Small Groups,* Sage, Thousand Oaks, CA, 1996.

---

[i] The results presented in this illustrative example were the result of a single run of the model with the following parameters: N = 75 ; PartProb = 0.0749 ; w = 0.05 ; c = 0.1 ; m = 5 ; INTRange = 0.6132. These results were chosen as representative after running a virtual experiment which systematically varied w,c,and m and randomly varied N, PartProb, and INTRange.

[ii] A COP's interest distribution is assessed by counting the number of members with interests at twenty equally spaced points within a [0,1] topic space.