A Simulation for Designing Online Community: Member Motivation, Contribution, and Discussion Moderation

Yuqing Ren
Carlson School of Management
University of Minnesota
Minneapolis, MN 55455

Robert E. Kraut
Human-Computer Interaction Institute
Carnegie Mellon University
Pittsburgh, PA 15213

ABSTRACT

This article describes and validates an agent-based model that integrates social psychological theories on collective effort, group identity, and interpersonal bonds to understand trade-offs in designing online communities. The model is then used to examine when different types of moderation in online communities will be valuable: no moderation in which all members are exposed to all messages, community-level moderation in which off-topic messages are deleted for everyone in the group, and personalized moderation in which people see different messages based on their interests. Compared to a no-moderation control, personalized moderation is effective in increasing members' contribution and commitment, especially in topically broad communities and those with high message volume. In contrast, community-level moderation increases member commitment but not contribution. By removing off-topic messages, community-level moderation increases members' information benefits at the expense of their opportunities for developing online relationships. This research demonstrates the value of computational modeling to synthesize narrow theories to describe behaviors in a complex system and to inform online community design.

Key words: online community, agent-based modeling, design, motivation, simulation

1. Introduction

As of 2007, approximately 70% of American adults use the Internet (Pew Internet 2007), and of these, 84% or about 90 millions Americans participate in an online group (Pew Internet 2001). They are used by the general public to discuss hobbies, sports, politics, finances or other topics of interest, leverage social networks, exchange technical information and social support, build collaborative artifacts, conduct business transactions or play games (Preece and Maloney-Krichmar 2003). They are used by business to exchange information with customers and suppliers, to improve operations internal by providing repositories of professional and technical information, communication spaces among employees, and platforms for new businesses (e.g., Gu et al. 2007, Ma and Agarwal 2007, Wasko and Faraj 2005).

Although many online communities are highly successful, many others fail. For example, the vast majority of open source development communities at SourceForge.com have only a handful of members and no activity (Crowston and Howison 2005). Across a wide range of Usenet groups, greater than 60% of the newcomers who post a message in any given month in a group are never seen again (Arguello et al. 2006).

A rich theory base in social psychology, organizational behavior, sociology and economics address many of the issues crucial to the success of an online community. For example, scholars have used theories of group identity and interpersonal bonds to examine the development of members' attachment to an online community (Ren et al. 2007, Sassenberg, 2002), public goods theories and theories of social loafing to analyze problems of under-contribution (Kollock, 1999, Ling et al. 2005), and economic signaling theory to understand how reputation influences community success (e.g., Friedman and Resnick 2001, Ma and Agarwal, 2007). The success of these analytic exercises demonstrates the applicability of traditional social science research and theory to the new phenomenon of online communities.

Even though social science theory is being used to describe existing communities, it is rarely used prescriptively, as the basis for designing them. A major reason is that the logic of design, which attempts to manage trade-offs among tens or hundreds of parameters that can influence community members' behavior, is at odds with the logic of the social sciences, which attempt to examine the influence of a small set of variables holding everything else equal. This *ceteris paribus* paradigm for developing and testing social science theory produces theories that are often too simple for the purpose of social engineering. Much research, for instance, focuses on the effects of a small set of variables, while attempting to hold other variables constant in regression models or short-lived experiments. Quantitative studies, if they examine many variables simultaneously, rarely examine higher order interactions. In contrast, social engineering requires theory that describes the behavior of a large set of factors varying simultaneously and their influences over a long time period. Making decisions about design trade-offs requires the examination of interactions.

For example, consider attempts to understand and remediate the problem of under-contribution, which is endemic in many online groups (Ling et al. 2005). Contribution, like most outcomes of interest, has multiple causes and each cause is treated by separate social science theories. Social psychologists developing the collective effort model to explain social loafing, for example, concentrate on contributors' identifiability, uniqueness of contribution and liking for the group (Karau and Williams 1993). Public-good economists emphasize expected utilities (Ledyar 1993). Contribution is also influenced by the commitment of individuals whom the community recruits and retains, the presence of explicit goals, group norms, and many other factors. Thus multiple theories are needed to model contribution in online communities and to design effective interventions to increase contribution. A single design choice can have cascading effects on contributions and other unforeseen outcomes. For example, the collective effort model proposes that people will contribute more to groups that they like. Therefore, increasing the homogeneity among group members may cause some members to contribute more to help similar others whom they like. However, the collective effort model also proposes that people will contribute more to a cause they believe their efforts are needed for group success. As a result, they may contribute less in a homogeneous group because they feel their efforts are redundant. Theory-driven design requires a broad navigation map that synthesizes insights from multiple theories to identify the pathways through which particular design choices may have positive and negative effects on the different outcomes that designers aim to achieve.

1.1. Agent-Based Modeling and Its Advantages

In this article, we present an agent-based model to express, synthesize, and extend social psychological theories that are relevant to motivation and contribution in online communities. Our goal is to understand trade-offs in designing online communities. Agent-based modeling is a way to capture the behaviors of complex adaptive systems from ground-up (North and Macal 2007). The emergent properties of a complex social system (e.g., a financial market, beehive or online community) are examined by simulating the behaviors of the agents that comprise the collective (e.g., the traders, the bees or the members). Agent-based modeling is typically used to understand connections between individual behavioral rules and system-level patterns and to predict potential outcomes of future actions.

The agent-based model described in this article simulates the behaviors of individual members of an online community to understand the dynamics of the community and to understand how various interventions affect community performance. Agents in the model are animated using principles derived from a set of well-established social science theories: the collective effort model of contribution to small groups (Karau and Williams 1993), theories of group identity and interpersonal bonds as the basis of commitment to groups (Prentice et al. 1994), information overload theory (Rogers and Agarwala-Rogers 1975), and public goods theory from economics (Ledyard 1995). Synthesizing multiple theories enables us to examine the multiple paths through which a design choice may affect members' motivation. Agent-based modeling also enables us to understand the complex, reciprocal interdependencies between member behaviors and community dynamics as a community develops and evolves over time. A design choice has both immediate, first-order effects (e.g., identifying members increasing their contributions, e.g., Williams et al. 1981) and longer-term, second-order effects (e.g., members' contributions increasing information overload and driving members away, e.g., Jones et al. 2004).

An agent-based model can serve as a test bed for running what-if experiments, which allow researchers to construct a mid-level theory to inform the design of online communities. Because existing theories are often too abstract for use in design, integrating and concreting them in an agent-based model enables one to identify places where theories agree, disagree, or are independent of each other, and to pin down factors that community designers could manipulate to produce desirable outcomes. Finally, the use of an agent-based model also opens up the black box between design choices and behavioral outcomes, by modeling intervening processes, such as the informational or social benefits that members receive when they read or write a message, and the way these processes mediate the relationship between design and visible outcomes.

1.2. Discussion Moderation as a Design Decision

To demonstrate the usefulness of simulations models of online communities, we apply our model to examine the effects of different types of discussion moderation on members' commitment and the community's growth. At the core of many online communities are members exchanging interests or engaging in conversations. Members converse to ask and answer questions, exchange opinions and social support, and to get to know each other. Without conversation, these communities would vanish. Even in online games like World of Warcraft or production-oriented communities like Wikipedia, members depend upon conversations to coordinate their work and to develop commitment to the group.

Even though communication is central to most online communities, too much communication or the wrong kind can threaten them. Sustaining conversations can be threatened by information overload fueled by high message volume and heterogeneity of discussion topics or member interests (Butler 2001, Jones et al. 2004). High message volume can be even more problematic when much of the conversation is off-topic. Although many communities are organized around specific topics, as members meet and interact with one another they often engage in personal, off-topic conversations that have nothing to do with the nominal topic. For example, a large minority of messages in a SeniorNet discussion group organized around about depression has nothing to do with depression (http://discussions.seniornet.org, Wright 2000) and messages in investment discussion groups have little to do with finances (Gu et al. 2007). For members only interested in getting information, off-topic messages are an irritation that can drive them away. High message volume can be especially problematic in communities that encourage conversation across a wide range of topics. As Butler (1999) notes, in communities with diverse interests, messages interesting to some members are likely to be off-topic and uninteresting to others.

To deal with problems of high message volume and off-topic conversation, designers and managers of online communities have introduced moderation techniques to manage conversations. Common practices include (1) community-level moderation (e.g., Yahoo! groups), in which human moderators or software agents block or remove inappropriate or off-topic messages, (2) personalized moderation, in which different users get a different subset of messages matched to their interests (e.g., Harper et al. 2007), (3) collaborative moderation (e.g., Digg.com or Slashdot.com), in which members rate messages so that others can use these ratings to guide their reading behavior (Lampe and Johnston 2005) and (4) partitioning of the community, by segmenting the community into smaller, homogeneous sub-forums. This paper contrasts community-level and personalized moderation.

One of our goals in creating a computation model was to illustrate how a theory based model can be used to synthesize theory and to answer practical questions about designing an online community. More specifically, we aim to answer three questions with respect to discussion moderation. (1) How does the style of moderation affect a community's viability and its members' commitment? (2) To what extent are the effects of moderation contingent upon community characteristics such as topical breadth and message volume? (3) How does the style of moderation affect the trade-offs among the various benefits that members receive from participating in an online community? Below we describe community-level and personalized moderation in more detail. Section 2 then describes the conceptual framework for the model, its theoretical background, and how we implemented and calibrated the model. Section 3 describes simulation experiments comparing three types of discussion moderation in different types of communities and our main findings. Section 4 discusses the practical implications of our findings and how agent-based modeling informs critical trade-offs in online community design.

1.2.1. Community-level moderation. Most online communities moderate messages at the community level: a message is available either for everyone visiting the site or for no one (Figallo 1998, Lampe and Johnston 2005). Community-level moderation can be performed ex ante, by approving or rejecting messages before they can be posted or ex post, by removing message after they have been posted. The goal is to prevent spam, trolling messages, anti-social flames or off-topic messages. Community-level moderation can be less effective in communities that attract members with diverse interests or ones that encourage diversity in content. In such broadly defined community, nominally on-topic messages might be of no interest to a large proportion of members. For example, in the movie discussion forum rottentomatoes.com, a message one evaluating a new action movie is likely to be of no interest to the many members who dislike action movies. Conversely, nominally off-topic conversations, such as one describing high school romances consummated in movie theatres, may be of great interest to some members. Under either scenario, community-level moderation leads to sub-optimal experience.

1.2.2. Personalized moderation. Some communities attempt to deal with the problem of message overload by personalizing the content, to match messages that a particular member sees with that member's interests. For example, in Raging Bull, an investor site, individual members can hide messages from particular posters and in Slashdot.com, users can surface messages based on ratings of quality, humor and other attributes. In e-commerce sites, personalization increases users' satisfaction by decreasing the total number of items to be processed and thus reducing information overload, while at the same time increasing each item's average fit to users' interests (Tam and Ho 2005, Liang et al. 2007, Shchafer et al. 2001). We believe personalized recommendation may have the same effect on user satisfaction and motivation in the context of online discussion groups. Previous research has shown that personalizing the messages that members are exposed to leads to increased participation, including reading and posting messages (Harper et al. 2007).

2. The Conceptual Framework for the Agent-Based Model

Figure 1 is the conceptual framework underlying the agent-based model of motivation in online communities. It is based on social psychological theories of motivation and voluntary contribution. Karau and William's collective effort model (1993) suggests that individuals contribute to a group to the extent that they believe their efforts will directly or indirectly lead to personal benefits. The model includes three types of benefits: (1) information benefits from accessing information or sharing information with others; (2) social benefits from identifying with a group or interacting with group members; and (3) other benefits from recreation and reputation. Consistent with the collective effort model, motivation to participate and contribute is calculated as a weighted sum of the three types of benefits a member expects to receive, with weights indicating how much that member values each type of benefit.

Insert Figure 1 about Here

Due to the complexity of the model, we describe the details of the model in three steps. (1) The decision rules that determine whether an agent reads or posts a message in the community; (2) calculations of the benefits the agent receives from membership, which determine the agent's motivation to read and to post messages; and (3) methods for model implementation and for running simulation experiments. For convenience, we describe how the model operates for a movie discussion forum. The model, however, applies broadly to text-based, conversationally-oriented online communities.

2.1. Member Actions: Reading and Posting Messages

Table 1 provides an overview of the decision rules an agent in the model uses to decide whether to take various actions. Following Butler (2001), we define participation as an action that members take to be exposed to communication activity, such as reading messages. We define contribution as an action that members take to actively engage in community activity, such as posting messages. Following the utility-like logic underlying the collective effort model, we assume that a member (1) logs in to read messages when expected benefit from participation exceeds expected cost, and (2) posts messages when expected benefit from contribution exceeds expected cost.

Formatted: Font: (Default) Times New Roman, Not Bold

Deleted: Table 1

Insert Table 1 about Here

2.1.1. Which messages to read? Typically, a member views messages in reverse chronological orders and stops viewing when he runs out of time, gets bored, or has finished reading all messages. We assume that the total number of messages a member views on a particular day depends upon the total number of messages available that day and how motivated the member is to read. We calculate messages an agent will view on a specific day as proportional to the amount of benefit he has received in the past from reading messages, capped by the total number of messages available to read. Because most people read in reverse chronological orders and messages get stale with time, members are more likely to view and respond to recent messages (i.e., messages posted within a day or so) and have a lower probability of reading older and less active messages (Arguello et al. 2006, Kalman et al. 2006) . To post a message, an agent makes two additional decisions: the topic of the message and, if he decides to post a reply, which message to respond to. For simplicity, we assume that the agent is equally likely to start a new thread or to reply to an existing thread, and sensitivity analyses indicate that our results remain robust when we vary the percentage of new threads from 50% to 30% and 70%.

2.1.2. What is the topic? The community can be organized broadly, relevant to any movie-related subjects including movie critics and celebrities, or more narrowly around a single topic, such as fantasy movies or Harry Porter. An agent can be interested in one or more of the topics. We assume that agents' interests remain static and do not change during the experimental period. We assume that each message concerns only one topic although the analysis is the same if each message refers to several topics. When an agent posts a thread-starting message, the topic of this message is a joint function of the agent's interests and the topics of messages the agent has recently viewed. When an agent posts a reply, the topic is a joint function of the agent's interests, topics of messages the agent has recently viewed, and the topic of the replied-to message. Thus, a fantasy movie lover is likely to initiate or reply to messages about fantasy movies, and this tendency will be greater in the fantasy movie forum than the general movie discussion forums. In communities with little off-topic discussion, members are less likely to bring up off-topic subjects for fear of violating group norms (Sassenberg 2002). Theory also suggests that newcomers are more likely to post on-topic messages than old-timers (Ren et al. 2007). Thus, we assume that agents posting for the first-time always begin with on-topic messages.

2.1.3. Which message to reply to? Theory and empirical evidence (Faraj and Johnson 2005, Fisher et al. 2006) suggest three common patterns of interaction among community members: (1) preferential attachment, in which members respond to popular messages or posters, (2) reciprocity, in which members respond to those who have written to them in the past, and (3) interest matching, in which members respond to messages that match their interests. Of course, people respond only to messages they have read. The agent in the model responds to messages based on a weighted sum of (1) the number of replies a potential to-be-replied message has received; (2) the number of times the poster of the message has responded to the agent; and (3) the match between the topic of the message and the agent's interests.

2.2. Member Benefits and Costs

Table 2 provides an overview of how information, social, and other benefits are implemented in the model, including the theories used to make assumptions, the rules used to calculate benefits, and key parameters in the benefit functions.

Formatted: Font: (Default) Times New Roman, 12 pt, Not Bold

Deleted: Table 2

Insert Table 2 about Here

2.2.1. Benefit from information exchange. We model two types of benefits related to information exchange: benefit an agent receives from accessing information and benefit the agent receives from providing information to others.

Benefit from accessing information. According to information overload theory, (1) only messages that match members' interests provide information benefit, and (2) benefit from accessing information is a marginally decreasing function of the number of messages viewed (Gu et al 2007, Jones et al. 2004). We calculate the benefit from reading messages as a joint function of the quantity and quality of messages that an agent reads. On average, the more messages that an agent views that match his interest, the greater information benefit he receives, with diminishing returns, because of information redundancy or information overload. The first graph in Table 2 illustrates the information access benefit function. The parameters were fixed based on an experimental study of Internet news recommendation. Liang et al. (2007) found that an increase from 20 to 40 news items caused information overload and led to a reductions in user satisfaction. Compared with news items, messages are shorter and less complex. Thus, we increased the value at which marginal benefit starts decreasing from 20 news items to 40 messages.

Reading messages takes time and effort. We assume that reading cost is proportional to the total number of messages an agent reads. In addition, having to evaluate and discard uninteresting messages increases the cost of reading (Gu et al. 2007). We thus calculated reading cost as a function that is proportional to the total number of messages the agent viewed divided by the signal-to-noise ratio, that is, the number of messages that match the agent's interests divided by the number of messages that fail to match his interests.

Benefit from sharing information. In many online communities, a small proportion of members often engage in altruistic behaviors such as answering questions (Fisher et al. 2006), or performing community maintenance tasks such as promoting and policing the site (Butler et al. 2007). According to the collective effort model (Karau and Williams 1993), social loafing in a group is greatly reduced when people perceive group tasks as interesting or when they identify with the group or like other members. The collective effort model also suggests that members are less willing to contribute if they believe that (1) other group members are already contributing or (2) if the group is large. Public goods theory also suggests less contribution as group size grows (Ledyard 1995).

Formatted: Font: (Default) Times New Roman, 12 pt, Not Bold

Deleted: Table 2

The pseudo code in Table 3 shows how we implemented these rules in the model. If the agent is interested in the messages or feels strongly attached to the group or its members, we calculate two components – one is non-zero when the group is perceived as at risk of failing (operationalized as hosting fewer than 100 messages) and the other is non-zero when others are perceived as under-contributing. We assume that agents who have a history of contributing ten times more than community average tend to perceive others as under contributing and therefore compensate for others' lack of contribution. Finally, we divide the sum of all components by a marginally decreasing function of group size or the total number of others who are present to contribute to capture the diffusion of responsibility effect.

Formatted: Font: (Default) Times New Roman, 12 pt

Deleted: Table 3

Insert Table 3 about Here

2.2.2. Benefit from social attachment. Prior literature shows that both identification with the group as a whole (i.e., a sense of belonging) and interpersonal bonds with particular members (i.e., friendship) can lead members to become committed to groups (e.g., Prentice et al. 1994, Sassenberg 2002). We model these identity-based attachment and bond-based attachment separately because they have distinct antecedents and consequences on members' attitudes and behaviors (Ren et al. 2007).

Benefit from identity-based attachment. Social identity theory suggests that assigning a member to a group, the presence of an out-group, and similarity among group members all lead to stronger attachment to the group. Shared interests and similarity in preferences have been used to manipulate and measure identity in laboratory experiments (Amichai-Hamburger 2005, Postmes and Spears, 2000). To simplify the model, we assume that people who share a common interest with the community identify with it. For example, a movie lover feels a stronger sense of belonging to a movie discussion group if other members are also movie lovers and if the conversation is about their shared interests than if the forum is full of discussion of jobs, love, politics, or other off-topic stories. We operationalize benefit from group identity as a function of the similarity between an agent's interest and the community's interest, calculated as the percentage of viewed messages that correspond to the agent's interests. The higher the percentage, the greater level of identity-based attachment the agent feels to the community.

Benefit from bond-based attachment. Small groups research suggests that repeated interactions lead to interpersonal attraction (Festinger et al. 1950); as the frequency of interaction between two persons increases, their liking for one another also increases (Cartwright and Zander 1953). Studies of Usenet groups suggest that getting a quick reply after posting seems to encourage members of an online community, especially newcomers, to return and participate in community discussion (Kraut et al. 2007). We speculate that replies from other members signal the likelihood of forming relationships with other in a community. We calculate benefit an agent receives from interpersonal bonds as a joint function of the number of other agents with whom the agent has developed a social relationship through repeated, mutual interaction (i.e., the two agents have responded to each other at least twice) and the number of responses the member received during the last period of interaction, whichever is higher. We assume benefit from interpersonal bonds also has a marginally decreasing function – as illustrated in Table 2 – the first few relationships an agent develops bring greater social benefit than subsequent ones.

2.2.3. Benefit from recreation. A third motivation that leads people to join online communities is recreational, that is, the enjoyment they derive from reading and posting online (Ridings and Gefen 2004). Several studies have identified stable individual differences in the extent to which people think online behavior is fun (e.g., Cotte et al. 2006). For instance, posters enjoy online interaction more than lurkers do (Preece et al. 2004). Our model captures these individual differences by drawing an agent's interest in reading and posting randomly from a right-skewed gamma distribution (as illustrated in Table 2). With a gamma distribution, the majority of members have a moderate level of interest in reading and posting and a small proportion of members have a high level of interest.

Formatted: Font: (Default) Times New Roman, 12 pt, Not Bold

Deleted: Table 2

2.2.4. Benefit from reputation. People are also motivated to contribute to online communities by the reputation they gain by doing so. Many online communities play on this motivation by institutionalizing "leader boards" and other devices that show the most active contributor. Amazon.com, for instance, uses the "top reviewers list" to recognize people who have contributed many reviews. Even when official recognition is absent, active contributors often get recognized by other members as an expert in certain topics or an enthusiastic help-provider. In the model, agents who are among the top ten percent contributors receive reputation benefit. Sensitivity analyses indicate that the main results were robust when the proportion receiving reputation benefit varied between 5% and 15%.

2.3. Costs, Motivation, Member Entry and Exit

2.3.1. Motivation as a weighted sum of benefits. As mentioned earlier, agents' motivation to read and to post is calculated as a weighted sum of their benefits from reading and posting minus costs of reading and posting. These weights differ across communities (Ridings and Gefen 2004). In the model, we set the relative weights for information exchange, identity, bonds, and other benefits respectively at 0.5, 0.1, 0.3, and 0.1; these weights are consistent with Riding and Gefen's (2004) finding about interest communities[1]. Within a single community, members have various reasons for joining. Some people may go to a movie discussion site for information about which movies to watch; others for dates and companionship, and yet others because they identify with being a movie lover. In the model, weights for individual agents' were drawn from normal distributions around the community means.

2.3.2. Costs of participation and contribution. We model three types of cost associated with reading and posting messages. Access cost simulates the time and effort people spend logging in order to read and post messages. Posting cost simulates time and effort spent composing messages. Compared with reading a message, posting one is more time-consuming and, thus, incurs a higher cost. For simplicity, we assume that starting a new thread and replying to an existing thread incur equal cost. Reading and posting messages also incur opportunity cost, which is the time could have been spent on alternative activities, such as work, conversation with family members, or reading and posting in other communities. We assume that opportunity costs are constant across different online communities, but variable across individuals (e.g., opportunity cost is higher for mid-career wage earners than for teens or retires).

2.3.3. Member entry and exit. Members join and leave online communities. Because there is little prior research describing the rate with which newcomers enter online communities, we analyzed 100 Usenet groups to estimate some parameters relevant to entry. This analysis indicates that the number of newcomers joining a community is proportional to community size

[1]Note that the current model describes behaviors within an interest community, like a movie discussion group. Although we do not vary community type within the model, had we done so, it would be by varying these weights. For example, the relative weights in a technical support group, in which people typically care less about interpersonal bonds and more about information, identity and reputation would be 0.55, 0.25, 0.1 and 0.1.

(see Butler 2001 for similar results) and follows a truncated gamma distribution function. Larger communities attract more newcomers per day. Although a community will attract an average or smaller number of newcomers on most days, it will attract many newcomers on a small number of days. In our model, agents do not make conscious decisions to leave the community. If their experienced benefits are low, they simply stop coming back.

2.4. Model Implementation and Calibration

We implement the computational simulation using NetLogo, a cross-platform multi-agent modeling environment (Wilensky 1999). Within the agent-based model, agents take actions within a simulated day. All active agents in the simulated community are given the opportunity to make a reading and posting decision before anyone moves to the next day. Messages posted the previous day are distributed to all agents the next day and used to update member expectations of benefits. In the jargon of agent-based modeling, actions are organized in staged episodes, and time is simulated as forced parallel.

We took three steps to insure the external validity of the model. Whenever possible, we drew insights from existing theories to specify the key assumptions and relationships in the model. The prior sections described this rationale. When theory was insufficient, we mined a rich set of data of 100 Usenet groups to fix the important parameters such as the ratio of new threads to replies or the entry rate for newcomers. We also went through an iterative calibration process during which we systematically varied key parameters to replicate behavioral patterns that have been repeatedly discovered in empirical studies. We describe this calibration in more detail below. To assure the robustness of our results, we ran a series of sensitivity analyses by relaxing key assumptions and varying key parameters. Results do not differ substantially from those we report in the results section.

2.4.1. Model calibration and validation. Model calibration is the process of adjusting a computational model to produce results that match real data or stylized facts with reasonable tolerance (Carley 1996). We used pattern calibration to establish the reasonableness of the model and its potential for predictive accuracy (Carley 1996). Pattern calibration compares the pattern or distribution of results generated by the computational model with the pattern or distribution generated from real data. Previous studies show that three statistics describing online communities – posts per member, replies per post, and communication partners (out-degrees) per member – demonstrate power-law distributions (Fisher et al. 2006, Smith 1999). We use these three stylized facts to calibrate the model.

In the beginning of the calibration, we constructed and simulated twelve online group using statistics from twelve Usenet groups with various size and message volume. We then engaged in an iterative process in which we compared the distribution of all three statistics from simulation with data from the real groups. After each run, we examined mismatches between the simulated and the real data, reexamined assumptions, and made adjustments to the model in light of theoretical reasoning, empirical evidence, or our intuitions. After about ten iterations, the model replicated the power-law distribution for all three statistics. The iterative calibration process helped select parameters, variables, and relations that yield outcomes that correspondence to the real world (Burton and Obel 1995), which greatly increases the construct and external validity of our model.

We then constructed and simulated another 25 online group using statistics from real groups. The simulated statistics fit the real statistics for these 25 groups well and demonstrate the validity of the model. Figure 2 illustrates the real and simulated statistics (after log transformation) in one of these 25 groups. We also examined survival curves for members and messages during model calibration and validation and found reasonable similarity between simulation and real data. As shown in Figure 3, the survival curve from real data suggests that about 60% of new posters fail to return after their first post and on average about 10% to 20% posters stick around for over 100 days

Insert Figure 2 and Figure 3 about Here

3. Simulation Experiments and Results

3.1. Virtual Experimental Design

In this section, we describe a full-factorial simulation experiment examining the effects of conversation moderation on community performance when topical breadth and message volume vary. We simulated three types of moderation: no moderation, community-level moderation, under which messages whose topics did not conform to community purpose were removed, and personalized moderation, under which a personalized algorithm presents a subset of messages that match member interests. We simulated three levels of topical breadth, in which groups dealt with one, five, or nine topics, and three levels of message volume with on average about 10, 15, and 20 messages per day.

Under each experimental condition, we ran a 365-day simulation on five randomly constructed groups. All groups began with 30 seed members and 30 seed messages and evolved over time as newcomers joined and old-timers left. At each simulated day, each agent assessed prior benefits from having read and posted messages and decided whether to login to read and post messages. For purposes of the simulation, the precision of the personalized moderation was set to 80% of recommended messages matching a member's interests. Sensitivity analyses suggest that the main results remain robust when the precision of personalized moderation varies between 60% and 100%.

We examine the effects of conversation moderation on two community outcomes: (1) the number of new posts per day, which is an indictor of community activity and (2) the average number of login sessions per member, which is an indicator of member commitment. We also examined (3) information benefits and (4) bonds benefits members received at the 100th, 150th, 200th, 250th and 300th day of the experiment. These benefits serve as mediators, to understand the link between experimental conditions and community outcome measures. We ran two two-way ANOVA analyses to examine the main effects of moderation and its interactions with topical breadth and message volume.

3.2. Results

3.2.1. Effects of Moderation on Community Activity. Analyses of the impact of moderation on *community activity* (posts per day) in communities differing in *topical breadth* revealed a significant main effect of moderation. As shown in Figure 4, personalized moderation led to the highest level of community activity, significantly more posts per day than community-level moderation and no moderation ($p < .001$). There was no significant difference between community-level moderation and no moderation ($p = .24$). The analyses also revealed a significant interaction between moderation and topical breadth ($p = .05$). Personalized moderation lead to more posts than community-level moderation and no moderation in communities with moderate (five topics) and high topical breadth (nine topics; $p < .001$), but not in communities with a narrow range of topics (one topic; $p > .15$). Topical breadth had no effect on community activity (i.e., amount of posting; $p = .65$).

Insert Figure 4 about Here

Analyses of the impact of moderation on community activity in communities differing in *message volume* revealed two main effects and a significant interaction between moderation and message volume, as shown in Figure 5. By definition, communities with high message volume had significantly more posts per day than communities with low message volume ($p < .001$). Community-level moderation, by removing off-topic messages, led to fewer posts than no moderation ($p = .02$) and personalized moderation ($p < .001$). What is interesting is the effect of personalized moderation. Personalized moderation, compared with community-level or no moderation, led members to contribute more messages ($p < .001$), and the effect was much greater in communities with higher message volume ($p < .001$).

Insert Figure 5 about Here

3.2.2. Effects of Moderation on Member Commitment. Analyses of the impact of moderation on *member commitment* (login sessions) in communities differing in *topical breadth* revealed a significant main effect of moderation ($p < .001$) and a significant interaction between moderation and topical breadth ($p < .001$). As shown in Figure 6, personalized and community-level moderation both led to more login sessions than no moderation, yet under different circumstances. Communication-level moderation led to more logins in communities with a single topic; personalized moderation led to more logins in communities with more topics. Both differences were statistically significant at the $p < .001$ level. Topical breadth had no effect on member commitment ($p = .45$).

Insert Figure 6 about Here

Analyses of the impact of moderation on commitment in communities differing in *message volume* revealed two main effects and a significant interaction between moderation and message volume ($p < .001$). Higher message volume led to more frequent logins. Personalized moderation led to more frequently logins than community-level moderation, and community-level moderation led to more frequent logins than no moderation. Both differences are significant at the $p < .01$ level. The effect of personalized moderation, as shown in Figure 7, was greater in communities with medium and higher message volumes than in communities with low message volume ($p = .01$).

Insert Figure 7 about Here

3.3. Member Benefits from Information and Bonds

Posts and logins are observable behaviors. To better understand the route by which moderation interventions affect agents posting and login behaviors, we examined their impact on various benefits that agents received as intermediate variables. Below we present results on two benefits – benefit from accessing information (informational benefit) and benefit from interpersonal bonds (relational benefit) – to illustrate an important design trade-off involved in moderation decisions.

3.3.1. Member benefits in communities with different topical breadth[2]. Figure 8 shows the effects of moderation and topical breadth on the amount of informational and relational benefits that agents received, averaged across all active members and the five snapshots at which member benefits were recorded.

Members received greater *informational benefit* in the topically broad communities than the topically narrow communities ($p < .001$) and with either type of moderation than no moderation ($p < .001$). The significant interaction between moderation and topical breadth indicates that the effects of moderation on informational benefit varied across communities with different topical breadth ($p < .001$). Community-level moderation led to

[2] Because previous analyses revealed no significant difference between medium and broad topical breadth, we left out medium topical breadth in Figure 8 and medium message volume in Figure 9 to make the figures more readable.

greater informational benefit in communities with a narrow focus whereas personalized moderation led to greater benefit in communities with a broad focus.

Members received greater *relational benefit* in the topically narrow communities than the topically broad communities ($p < .001$). Personalized or no moderation led to greater relational benefit than community-level moderation ($p < .001$). The effects of moderation on relational benefit also varied across communities with different topical breadth ($p < .001$). No moderation led to the greatest relational benefit in communities with a narrow focus; whereas personalized moderation led to the greatest relational benefit in communities with a broad focus, followed by no and then community-level moderation.

Insert Figure 8 about Here

3.3.2. Member benefits in communities with different message volume. Figure 9 shows the effects of moderation and message volume on informational and relational benefits that agents received, averaged across all active members and the five snapshots.

Formatted: Font: (Default) Times New Roman

Deleted: Figure 9

Members received greater *informational benefit* in communities with higher message volume ($p < .001$) and in communities with either type of moderation ($p < .001$). The interaction between moderation and message volume was not significant ($p = .10$), suggesting that the effect of moderation did not vary in communities with different levels of message volume.

Members experienced greater *relational benefit* in communities with higher message volume ($p < .001$). Personalized moderation led to the greatest relational benefit, followed by no moderation and community-level moderation ($p < .001$). The most striking result is the interaction between moderation and message volume ($p < .01$). As message volume increases, the effects of no and community-level moderation remained roughly the same whereas the positive effects of personalized moderation increased significantly and substantially.

Insert Figure 9 about Here

Overall, the examination of informational and relational benefits as intermediate variables linking moderation to behavioral outcomes highlights an interesting trade-off in moderating online discussion. While community-level moderation – using a single removal rule for off-topic messages for everyone – improved the benefits members receive from accessing valuable information, doing so lowers the benefits from interpersonal relationships. The effects are especially strong in communities covering a broad range of topics. It is for this reason that some communities, such as health support groups whose member seek both information and online relationships, can be difficult to design. To the extent that the simulation reflects the benefits members receive in real online communities, it suggests that community-level moderation, even though it is a widely adopted practice in managing online conversation, promotes informational benefit at the expense of relational benefit. In contrast, personalized moderation improves both informational and relational benefits, especially in communities that involve many topics and heavy traffic. It is a technology to handle the informational versus relational trade-off, and can be effective even if the algorithms to predict interests are only moderately accurate (e.g., 60 or 80 percent precision).

4. Discussion

In this study, we develop an agent-based model to express and synthesize social theories to understand trade-offs in online community design. By building upon and integrating insights from multiple theories, our model depicts a more complete picture, than any single theory has depicted on its own, of how individual motivation and interactions affect community dynamics. Our effort is successful in at least three regards. First, our application of the model to understand how moderation affects community activity and member commitment leads to plausible yet interesting predictions of the effectiveness of community-level and personalized moderation. Second, the availability of intermediate variables enables us to examine not only the effects of design interventions but also why and how these effects happen, which illuminates the critical trade-off between designing for informational benefit and designing for relational benefit. Third, in addition to providing a platform for examining the effectiveness of different types of conversational moderation, the model serves as a mid-level theory that can be further extended and applied to examine other design decisions such as leader board, community size, and newcomer socialization. Validity check demonstrates the potential value of our model as a test bed to inform and assist online community design.

4.1. Summary of Main Findings

Table 4 summarizes the main findings. The most interesting finding is that across multiple outcomes, community-level moderation is less successful than either its common use or experts' opinions would suggest. For instance, Preece (2000) notes a moderator's number one task is to "keep the group focused and on-topic" (p. 84) and the most important guideline for moderators is to "suspend personal opinion" (p. 293). In contrast, the simulation experiments presented here suggest that community-level moderation is only moderately effective. It did not encourage members to contribute, and its positive effect on member commitment occurred only in narrowly defined groups. As the analysis of benefits shows, one reason is that by removing off-topic messages community-level moderation increases the benefit that members receive from accessing information but reduces opportunities for developing online relationships.

Deleted: Table 4

Formatted: Font: (Default) Times New Roman, 12 pt, Not Bold

Comment: I don't understand the relevance of the suspend personal opinion quote.

Deleted: had no significant effect on

Deleted: ing

Deleted: ,

Deleted: ,

Insert Table 4 about Here

In comparison, the results illustrate the benefits of personalized moderation, an under-exploited yet promising approach for managing online discussion. Its positive effects are especially prominent in topically broad and high traffic communities when members are at risk of information overload. Compared with community moderation, in which moderation removes the same off-topic messages for everyone, personalized moderation resolves the trade-off between accessing information and establishing relationships by allowing community designers and managers to customize views and experiences to match different members' interests.

Deleted: developed

Deleted: ,

Deleted: the one-size-fits-all approach (i.e., filtering and

Deleted: ing

Deleted:)

The interactions between moderation styles and both topical breadth and message volume suggest that, as in the case with organizational design more generally (Galbraith 1973), online community design also involves many contingencies. There is no single best way to design and manage an online community. Rather, any of the three moderation styles can be a good solution, depending upon community characteristics (topical breadth, message volume, and developmental stage) and specific goals that designers wish to accomplish (to make members loyal or to increase their contribution). No moderation works well if the goal is to build relationship in a narrowly defined community. If the goal is to facilitate information exchanges about a single topic and to increase member commitment, then community-level moderation is best. If the goal is to encourage both contribution and commitment in an established community with diverse interests and high message volume, personalized moderation works best.

Deleted: moderating effects of

Deleted: similar to

Comment: We never talk about developmental stage. Do you need to add a hint in the results section?

The simulation produced two unexpected results that need further investigation. The first is that community-level moderation led to greater commitment but not contribution. One possible reason is that community-level moderation has different effects on posters versus lurkers. Due to the public goods nature of online conversations, posters and lurkers have equal access to the information provided by other members. By limiting and removing off-topic messages, lurkers, who are driven primarily by information benefit, receive disproportionate benefit and return more frequently whereas posters, who are driven by both information and social benefits, may be discouraged from coming back and posting more messages. It is also possible that community-level moderation encourages on-topic posts yet discourages off-topic posts. The two effects cancel each other out and lead to no overall increase in the total number of posts. Meanwhile, by removing off-topic messages, community-level moderation increases the signal to noise ratio and thus improves the overall experience of members who are interested in the community's topics. As a result, these members became more committed and visit the community more frequently.

Deleted: effect

Deleted: on

Deleted: new

Deleted: identify with

Second, the simulation revealed an unexpected, negative effect of topical breadth on relational benefit. Simulation results suggest that in communities dealing with many topics, members seem to be less likely to engage in repeated interactions to form strong relationships; whereas in communities with a single topic, members' shared interest in the common topic can serve as a powerful bonding mechanism that leads to interpersonal liking. This effect highlights the challenge in managing communities with mixed motives and diverse member interests, such as a health support group. In these communities, personalized moderation may help members to find and interact with other who share similar interests.

4.2. Limitations

This research is not without limitations. In constructing the model, we walked a fine line between transparency and accuracy. To make the model clear and interpretable, we made simplifying assumptions to capture the essence of member motivation. In this section, we acknowledge these limitations, speculate how altering these assumptions may change our results, and discuss ways to relax these assumptions and extend the model in future research.

We simulated one type of community, interest-based communities, such as a movie discussion group, organized around a set of shared interests or topics. There exist many other types of communities such as technical and health support groups, political discussion groups, online gaming communities, and social networking sites. These communities different in the amount that amount that members are motivated by information, social, reputation and recreational motivations. Caution needs to be taken to generalize our findings to other types of communities, and future research should examine the effects of conversation moderation and design trade-offs in all types of online community. For example, because members in technical support groups compared to hobby groups are generally less motivated by social benefits but more motivated by reputation, it is likely that different types of moderation will have differ effects in these groups.

We assumed that member preferences and interests are exogenous and static over the course of a year. In real communities, however, newcomers who originally join to talk about the nominal topic of the group may increasingly value friendship with one another after repeated encounters. Likewise, member interests or attitudes towards certain topics may shift over time in response to the message they are exposed to.

In developing the simulation, we replied upon our best judgment to estimate key parameters in the benefit functions or distributions, when theory or empirical evidence was absent (Sterman, 2002). We ran a series of sensitivity analyses to assure that our main findings are robust and not dependent upon these parameters taking certain values. Our main results, including the positive effects of personalized moderation, remained unchanged, which further enhanced our confidence in the validity and reliability of the model. We mentioned some of the parameters

Deleted: such as movie discussion

Comment: Mention how different communities will change the model & speculate on effects. I.e., t the weights on motivations are different. E.g, tech group little weight on social benefits while health support group high weight on both info & social.

Deleted: movies

Deleted: become

Deleted: friends

Comment: Drop. Endogenous motivations (i.e., changing the weights for the movations) may have large impact on the model.

Deleted: Given the one-year duration of our simulation experiment, we do not anticipate that our main findings will change substantially when these assumptions are relaxed

Deleted: There are several cases when theory or empirical evidence was absent, that we varied in the sensitivity analyses in model description, and a complete list is available upon request.

4.3. Concluding Remarks

Online communities are successful to the extent that members return repeatedly and contribute material that others value, and to the extent that members receive benefits when they visit. Because many design decisions are not motivated by a systematic understanding of member motivation and contribution in these communities and are designed through intuition and trial and error, many designs are less effective than they could be. In this study, we treat online communities as socio-technical systems that need to be carefully designed to fit their strategic goals and environments. In other words, we believe that online community design can go beyond intuition and trial and error and can benefit from the prescriptive power of social science theory. We believe that our model, incorporating multiple theoretical perspectives, has the potential to evolve into a multi-contingency tool for diagnosis and design (Burton and Obel 2004) of online communities. Theoretical knowledge and predictions embedded in the model can be combined with creative design intuition to generate effective design decisions. The model presented here was primarily based on social psychological theory, but the social science literature offers a wide range of theory explaining group and organization behavior that could be exploited for community design. We acknowledge that the design of online communities will rely for the foreseeable future on creative designers basing their designs upon observations of what has worked in the past and feedback from the field. We hope, however, that our model can serve as a test bed that help designers gain some preliminary knowledge of which features to experiment with. We also hope our research on the application of theory to the problem of online community design serves as a case study of how to to extract value from the social sciences for design.

Deleted: Lacking

Deleted: practices, driven

Deleted: by experiences and trials and errors, fail

Deleted: es

Deleted: creation

Deleted: psychological

Deleted: theori

Deleted: es

Deleted: tool

Deleted: to assist

Deleted: y design

Deleted: While community designers play the game of tweaking knobs to get their intended results, w

Deleted: community

Deleted: knobs to turn and in which direction.

Deleted: , which develops broad lessons on how

References

Amichai-Hamburger, Y. 2005. Internet minimal group paradigm. *CyberPsychology & Behavior*. 8(2)140-142.

Arguello, M., Butler, B., Joyce, E., Kraut, R., Ling, K.S., Wang, X. 2006. Talk to me: Foundations for successful individual-group interactions in online communities. *ACM Conference on Human-Factors in Computing Systems*, New York: ACM Press. 959-968.

Burton, R.M., Obel, B. 1995. The validity of computational models in organization science: From model realism to purpose of the model. *Computational and Mathematical Organization Theory*. 1(1) 57-71.

Burton, R.M., Obel, B. 2004. *Strategic Organizational Diagnosis and Design: The Dynamics of Fit, with OrgCon Software*. Kluwer Academic Publishers, Boston, Third Edition.

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

Butler, B., Sproull, L., Kiesler, S., Kraut, R. 2007. Community effort in online groups: Who does the work and why? Weisband, S.P., ed. *Leadership at a Distance: Research in Technologically-Supported Work*, Lawrence Erlbaum Associates, Mahwah, NJ. 171-194.

Butler, B.S. 2001. Membership size, communication activity, and sustainability: A resource-based model of online social structures. *Information Systems Research*. 12(4) 346-362.

Carley, K. 1996. Validating computational models. Working Paper, Pittsburgh, PA.

Cartwright, D., Zander, A. 1953. Group cohesiveness: Introduction. Cartwright, D., Zander, A., eds. *Group Dynamics: Research and Theory*. Row Peterson, Evanston, IL.

Cotte, J., Chowdhury, T. G., Ratneshwar, S., Ricci, L. M. 2006. Pleasure or utility? Time planning style and Web usage behaviors. *Journal of Interactive Marketing*. 20 45-57.

Crowston, K., Howison, J. 2005. The social structure of free and open source software development. *First Monday*. 10(2) np.

Festinger, L. 1968. Informal Social Communication. Cartwright, D., Zander, A. eds. *Group Dynamics: Research and Theory*. Harper & Row, New York. 14 182-191.

Figallo, C. 1998. *Hosting Web Communities: Building Relationships, Increasing Customer Loyalty, and Maintaining A Competitive Edge*. John Wiley & Sons, New York.

Friedman, E., Resnick, P. 2001. The social cost of cheap pseudonyms. *Journal of Economics and Management Strategy*. 10(2) 173-199.

Galbraith, J. 1973. *Designing Complex Organizations*. Addison-Wesley, Reading, MA.

Gu, B., Konana, P., Rajagopalan, B., Chen, H. W. M. 2007. Competition among virtual communities and user valuation: The case of investing-related communities. *Information Systems Research,* 18(1) 68.

Harper, F.M., Frankowski, D., Drenner, S., Ren, Y., Kiesler, S., Terveen, L., Kraut, R.E., Reidl, J.T. 2007. Talk amongst yourselves: Inviting users to participate in online conversations. *12th International Conference on Intelligent User Interfaces*, Honolulu, Hawaii, 62-71.

Herlocker, J.L., Konstan, J.A., Terveen, L.G., Riedl, J.T. 2004. Evaluating collaborative filtering recommender systems. *ACM Transactions on Information Systems*. 22(1) 5-53.

Johnson, S.L., Faraj, S. 2005. Preferential attachment and mutuality in electronic knowledge networks. *26th International Conference on Information Systems*, Las Vegas, NV. 287-299.

Jones, Q., Ravid, G., Rafaeli, S. 2004. Information overload and the message dynamics of online interaction spaces: A theoretical model and empirical exploration. *Information Systems Research*. 15(2) 194-210.

Kalman, Y. M., Ravid, G., Raban, D. R., Rafaeli, S. 2006. Pauses and response latencies: A chronemic analysis of asynchronous cmc. *Journal of Computer-Mediated Communication*. 12(1) 1-23.

Karau, S.J., Williams, K.D. 1993. Social loafing: A meta-analytic review and theoretical integration. *Journal of Personality & Social Psychology*. 65(4) 681-706.

Kollock, P. 1999. The economies of online cooperation: Gifts and public goods in cyberspace. Smith, M.A., Kollock, P., eds. *Communities in Cyberspace*. Routledge, London.

Kraut, R. E., Wang, X., Butler, B., Joyce, E., Burke, M. 2007. Building commitment and contribution in online groups through interaction. Working Paper. Carnegie Mellon University.

Lampe, C., Johnston, E. 2005. Follow the (slash) dot: effects of feedback on new members in an online community. *ACM Conference on Supporting Group Work*. ACM Press. 11-20.

Ledyard, J.O. 1995. Public goods: A survey of experimental research. Kagel, J.H., Roth, A.E., eds. *The Handbook of Experimental Economics*. Princeton University Press, Princeton, NJ. 111-194.

Liang, T.-P., Lai, H.-J., Ku, Y.-C. 2007. Personalized content recommendation and user satisfaction: Theoretical synthesis and empirical findings. *Journal of Management Information Systems*. 23(3) 45-70.

Ling, K., Beenen, G., Ludford, P.J., Wang, X., Chang, K., Li, X., Cosley, D., Frankowski, D., Terveen, L., Rashid, A.M., Resnick, P., Kraut, R.E. 2005. Using social psychology to motivate contributions to online communities. *Journal of Computer Mediated Communication*. 10(4) article 10.

McKenna, K.Y.A., Green, A.S., Gleason, M.E.J. 2002. Relationship Formation on the Internet: What's the Big Attraction? *Journal of Social Issues*. 58(1) 9-31.

Ma, M., Agarwal, R. 2007. Through a glass darkly: Information technology design, identity verification, and knowledge contribution in online communities*.* *Information Systems Research,* 18(1) 42.

North, M.J., Macal, C.M. 2007. *Managing Business Complexity: Discovering Strategic Solutions with Agent-Based Modeling and Simulation*. Oxford University Press, London.

Pew Internet. 2007. Internet adoption. http://www.pewinternet.org/PPF/r/148/report_display.asp. accessed on May 1, 2007.

Pew Internet. 2001. Online communities: Networks that nurture long-distance relationships and local ties. http://www.pewinternet.org/PPF/r/47/report_display.asp, accessed on May 1, 2007.

Postmes, T., Spears, R. 2000. Refining the cognitive redefinition of the group: Deindividuation effects in common bond vs. common identity groups. Postmes, T., Spears, R., Lea, M., Reicher, S., eds. *SIDE effects centre stage: Recent developments in studies of de-individuation in groups*. KNAW, Amsterdam, the Netherlands 63-78.

Preece, J. 2000. *Online Communities: Designing Usability, Supporting Sociability*. Wiley, Chichester, England.

Preece, J., Maloney-Krichmar, D. 2003. Online communities. Jacko, J., Sears, A. A., eds. *Handbook of Human-Computer Interaction*. Lawrence Erlbaum Associates, Mahwah: NJ. 596-620.

Preece, J., Nonnecke, B., Andrews, D. 2004. The top 5 reasons for lurking: Improving community experiences for everyone. *Computers in Human Behavior*. 20(2) 201-223.

Prentice, D.A., Miller, D.T., Lightdale, J.R. 1994. Asymmetries in attachments to groups and to their members: Distinguishing between common-identity and common-bond groups. *Personality & Social Psychology Bulletin*. 20(5) 484-493.

Ren, Y., Kraut, R.E., Kiesler, S. 2007. Applying common identity and bond theory to design of online communities. *Organization Studies*. 28(3) 377-408.

Ridings, C.M., Gefen, D. 2004. Virtual community attraction: Why people hang out online. *Journal of Computer Mediated Communication*. 10(1) np.

Rogers, E.M., and Agarwala-Rogers, R. 1975. Organizational communication. Communication Behavior. Hanneman, G.L., McEwen, W.J., eds. Addison-Wesley, Reading, MA. 218-236.

Sassenberg, K. 2002. Common bond and common identity groups on the Internet: Attachment and normative behavior in on-topic and off-topic chats. *Group Dynamics*. 6(1) 27-37.

Shchafer, J., Konstan, J., Riedl, J. 2001. E-commerce recommendation application. *Data Mining and Knowledge Discovery*, Kluwer Academic Publishers.

Smith, M.A. 1999. Invisible crowds in cyberspace: Mapping the social structure of the Usenet. Smith, M. A., Kollock, P., eds. *Communities in Cyberspace*. Routledge, London.

Tam, K. Y., Ho, S.Y. 2005. Web personalization as a persuasion strategy: An elaboration likelihood model perspective. *Information Systems Research*. 16(3) 271-291.

Wasko, M., Faraj, S. 2005. Why should I share? Examining social capital and knowledge contribution in electronic networks of practice. *MIS Quarterly*. 29(1) 35-58.

Wilensky, U. 1999. NetLogo. Center for Connected Learning and Computer-Based Modeling, Northwestern University, Evanston, IL.

Williams, K., Harkins, S. G., Latane, B. 1981. Identifiability as a deterrent to social loafing: Two cheering experiments. *Journal of Personality & Social Psychology*. 40(2), 303-311.

Wright, K. 2000. The communication of social support within an on-line community for older adults: A qualitative analysis of the SeniorNet community. *Qualitative Research Reports in Communication*. 1(2) 33-43.

Table 1. Definition and Rules for Member Decisions

Decisions	Definitions	Rules
Participation	Reading messages	If expected benefit from reading exceeds expected cost of reading
Contribution	Posting messages	If expected benefit from posting exceeds expected cost of posting
Message selection	Which messages to read?	Latest messages followed by recent messages, proportional to expected benefit from

		reading
Topic selection	What is the message topic?	Jointly determined by topics of recently viewed messages, personal interest, (and topic of original message when posting a reply message)
Conversation selection	Which message to respond to?	Jointly determined by preferential attachment, reciprocity, and match of personal interest

Table 2. Rationale, Rules, and Functions to Calculate Member Benefits

	Rationale	Rules	Function and Parameters
Information Benefit			
From accessing information (*$InfoB_{access}$*)	Information overload	Only messages matching an agent's interest provide $InfoB_{access}$, and $InfoB_{access}$ is a marginally decreasing function of the number of messages read	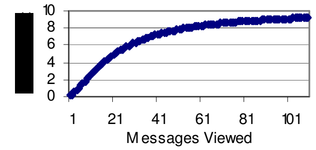
From sharing information (*$InfoB_{share}$*)	Collective effort model Public goods theory	$InfoB_{share}$ is conditional on liking task or group, is greater when others under-contribute, and is greater when group size is smaller	$\begin{cases} 0, & \text{if } \max(IntrB_{rec}, SocB_{iden}, SocB_{bond}) < 3 \\ f(avg\ msg, group\ size), & otherwise \end{cases}$
Social Benefit			
From attachment to group ($SocB_{iden}$)	Group identity	$SocB_{iden}$ is greater when agent's interests are similar to group interests	$f\left(\frac{count\,(viewed\ messages\ that\ match)}{count\,(viewed\ messages)}\right)$
From attachment to members ($SocB_{bond}$)	Interpersonal bonds Empirical studies of Usenet groups	$SocB_{bond}$ is greater with repeated, mutual interactions, with immediate responses from other community members, and is a marginally decreasing function of the number of relationships an agent forms	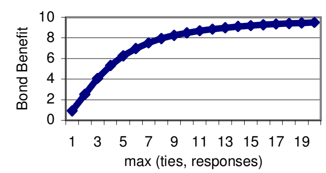
Other Benefit			
From recreation (*$IntrB_{rec}$*)	Intrinsic motivation Empirical studies of online behaviors	$IntrB_{rec}$ is a function of individual differences, distributed normally at the agent level, and distributed as a gamma distribution at the community level	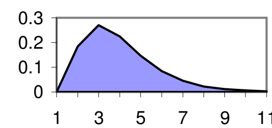
From reputation (*$IntrB_{rep}$*)	Incentive mechanisms	When an individual member contributes more than10% of the max contribution	$f\left(\frac{self\ contribution}{\max contribution / 10}\right)$

Table 3. Pseudo-code for calculating benefits from sharing information

```
Initialize information sharing benefit to zero
/* only contribute when task valence or group valence is high */
IF any of identity benefit, bonds benefit, or recreation benefit > 3 THEN
    /* more likely to contribute when group is at stake*/
    IF total messages < 100 THEN
        Increase information sharing benefit by 5 times (100 – total messages) / 100
    /*more likely to contribute when perceiving others as under-contribution*/
    IF average other contribution < 10% of self contribution THEN
        Increase information sharing benefit by 3 times self / other contribution
ENDIF
/* less likely to contribute in groups larger than 15*/
IF group size > 15 THEN
    Multiply information sharing benefit by (1 – (group size – 15) / (group size + 15))
ENDIF
```

Table 4. Summary of the Effects of Community-Level and Personalization Moderation

	Community-level moderation	**Personalized moderation**
Member commitment (logins)	• positive, significant • a greater effect in narrow-focus or high-traffic groups	• positive significant • a greater effect in broad-focus or high-traffic groups
Member contribution (posts)	• negative, not significant	• positive, significant • a greater effect in broad-focus or high-traffic groups
Benefit from information access	• positive, significant	• positive, significant • a greater effect in broad-focus or high-traffic groups
Benefit from interpersonal bonds	• negative, significant	• positive, significant • a greater effect in broad-focus or high-traffic groups

Figure 1. The Conceptual Framework for the Agent-Based Model

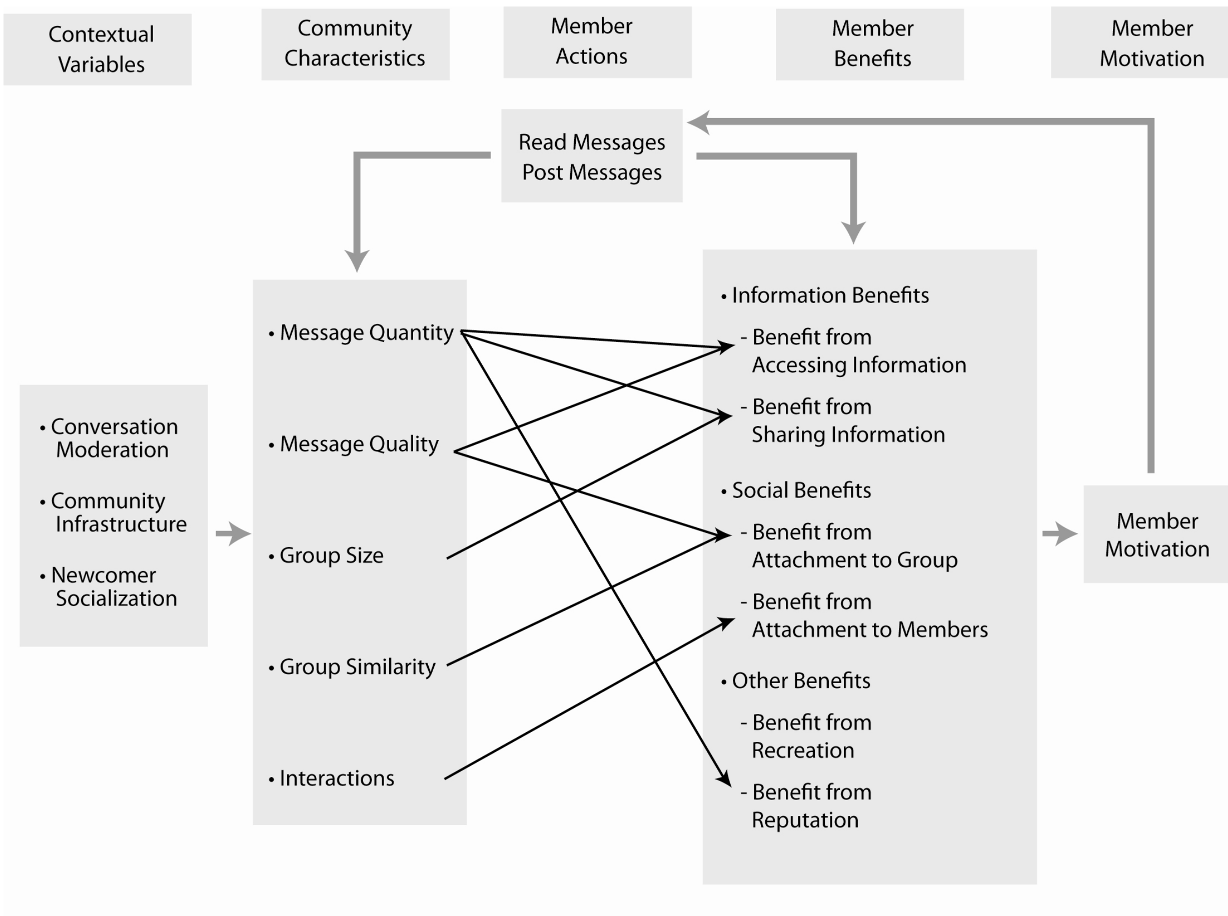

Figure 2. Comparison of Real (left) and Simulated (right) Data in Community Statistics

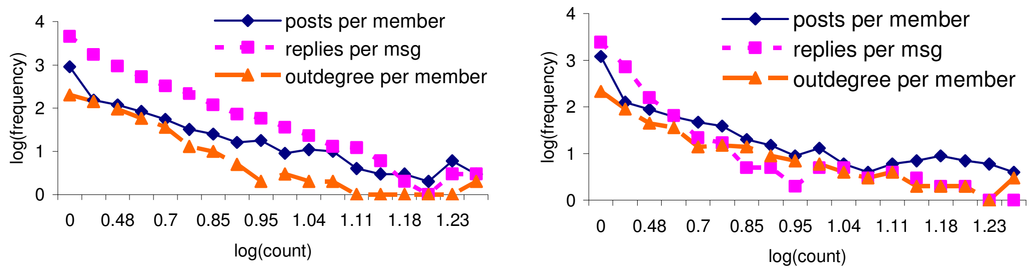

Figure 3. Comparison of Real (left) and Simulated (right) Data in Member Survival

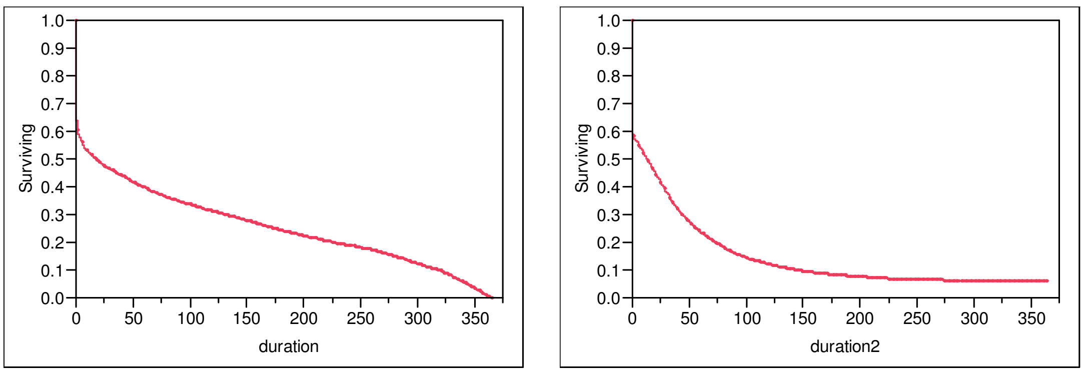

Figure 4. Effects of Moderation on Community Activity When Topical Breadth Varies

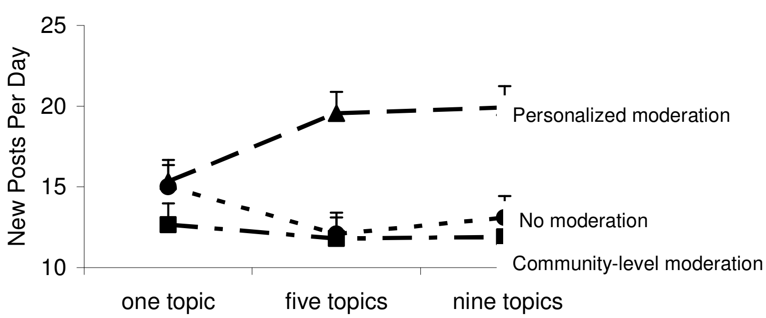

Figure 5. Effects of Moderation on Community Activity When Message Volume Varies

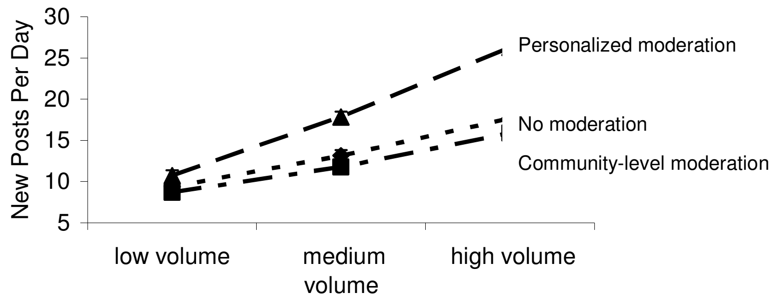

Figure 6. Effects of Moderation on Member Commitment When Topical Breadth Varies

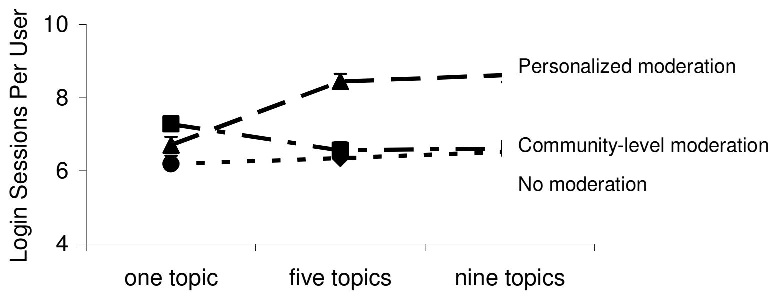

Figure 7. Effects of Moderation on Member Commitment When Message Volume Varies

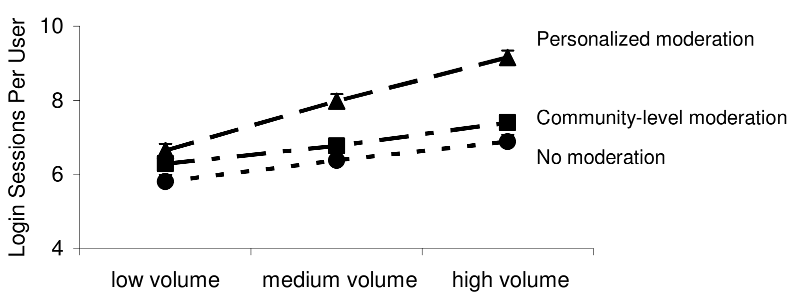

Figure 8. Effects of Moderation on Benefits with Narrow and Broad Topic Breadths

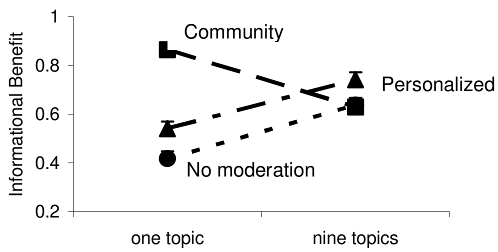

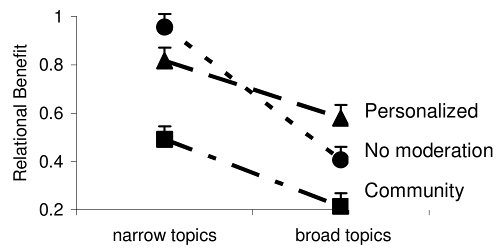

Figure 9. Effects of Moderation on Benefits with Low and High Message Volumes