

# The Social Psychological Effects of Feedback on the Production of Internet Information Pools

Coye Cheshire  
Judd Antin

UC Berkeley School of Information  
Berkeley

*A growing number of systems on the Internet create what we call information pools, or collections of online information goods for public, club or private consumption. Examples of information pools include collaborative editing websites (e.g. Wikipedia), peer-to-peer file sharing networks (e.g., Napster), multimedia contribution sites (e.g. YouTube), and amorphous collections of commentary (e.g., blogs). In this study, we specifically focus on information pools that create a public good. Following current theory and research, we argue that extremely low costs of contribution combined with very large networks of distribution facilitate the production of online information pools—despite an abundance of free-riding behavior. This paper presents results from a series of Internet field experiments that examine the effects of various feedback mechanisms on repeat contributions to an information pool. We demonstrate that the social psychological benefits from gratitude, historical reminders of past behavior, and ranking of one's contributions relative to those of others can significantly increase repeat contributions. In addition, the context in which individuals interact with the system may partially mitigate the positive effect of some types of feedback on contribution behavior.*

doi:10.1111/j.1083-6101.2008.00416.x

## Introduction

Systems that facilitate computer-mediated exchanges of digital information grew along with the Internet in the early 1990's, providing new ways to quickly and efficiently share text, music, movies, software and other digital goods. The purpose of these systems varies widely, from the production of an 'online encyclopedia that anyone can edit' (Wikipedia) to the distribution of digital media (e.g. peer-to-peer systems such as the original Napster). Though diverse in purpose, each system provides different types of incentives that can encourage individuals to contribute information. A growing body of research looks at how non-monetary incentives

harness social psychological processes to promote increased contributions (e.g., Ling et al., 2005; Rafaeli and Raban, 2005; Rafaeli et al., 2005; Rashid et al., 2006; Cheshire, 2007). This paper complements this line of research by focusing on the question of how non-monetary incentives encourage individuals to contribute small quantities of information in different contexts of interaction. We report the results of a field experiment which examines the effects of three different synchronous, direct feedback mechanisms on repeat contributions to an online system of information exchange.

When digital information goods from many different sources are collectively transmitted over a computer network so that they can be accessed by groups of individuals, they create an *information pool*. In these systems, individual contributions of digital information combine to produce information products for public, club, or private consumption.<sup>1</sup> Digital information goods may include (but are not limited to) software, photographs, art, music, speeches/lectures, videos, and general discourse (Kollock, 1999; Shapiro and Varian, 1998). Examples of information pools that are created through individual contributions of digital information goods include peer-to-peer exchange systems, collaborative editing systems that allow individuals to contribute text and multimedia content for a defined purpose, distributed work systems in which individuals contribute very small quantities of information to help complete much larger tasks (e.g., NASA Clickworkers, <http://clickworkers.arc.nasa.gov>), as well as amorphous collections of commentary (e.g., blogs).

Information pools exist in many different forms and occur through a variety of behaviors. However, in this paper we limit our focus to information systems that create public goods, which are defined by the situation in which one cannot prevent others from benefiting from the good and consumption by one does not affect consumption for others (Olson, 1965). As is the case in public good problems, individuals who participate in information pools must overcome the temptation to free-ride (consuming the public good without making a contribution to it). If everyone followed the strategy of free-riding, the public good would not be maintained or even produced at all.

Our notion of an information pool is similar to the concept of a *discretionary database*, which is a shared database structure within an organization that depends on contributions from a small group of individuals (Thorn and Connolly, 1987). Discretionary databases are exemplified by automatic meeting-scheduling systems (Rafaeli and LaRose, 1993: 279). Both information pools and discretionary databases acknowledge the collective action problem in shared information systems, and both focus on individual contribution levels as the primary dependent variable. However, an 'information pool' is a broad concept that includes systems of shared information goods that take many technological forms. While some information pools do create database structures (e.g. Wikipedia), many others do not (e.g. blogs). Similarly, the concept of 'discretion' is just one of many ways to describe the social production of information pools. Individuals who create information pools may or may not know how their contributions will be combined or distributed, and contributions may not

be discretionary in any given situation. Thus, we view discretionary databases as a subset the larger category of information pools.

At least two key features of information affect the way that it is transferred and collected on the Internet. First, information can be consumed by many individuals without losing much, or any, value (Shapiro and Varian, 1998; Rafaeli and Raban, 2005; Cheshire, 2007). In economic terms, this is called high jointness of supply or non-rival goods. The second key feature of information is that it can be transferred to another individual without the original owner losing her copy of the same information (Rafaeli and Raban, 2005; Cheshire, 2007). The property of replication allows the owner to retain the value of the information even when she shares it with others, while the property of pure (or very high) jointness of supply allows the value of the information to stay constant regardless of the number of individuals who consume it.<sup>2</sup>

If individuals can keep the information that they share while many others simultaneously benefit from it, then the overall cost-to-benefit ratio is greatly shifted in favor of sharing—especially when compared to the *de facto* standards of physical goods and services. Still, some costs do remain. These may take the form of time, computer equipment required to make copies, or the cost of the internet connection used by the contributor. However, these costs are arguably small compared to the value of the content in the information (Kollock, 1999; Cheshire, 2007).

The large disparity between the costs incurred by the contributor and the content value of the good is fundamental to understanding the creation of information pools. When a given individual shares information, she must contribute while knowing that her contributions are made at some small loss, and that they do not directly affect her current information gains. However, it is reasonable to assume that at least *some* people will still make the decision to contribute, especially if they have strong positive beliefs about the overall outcome. These individuals may be rare, and this type of behavior may or may not be typical. We only need to assume that the probability of some contribution in the population is greater than zero. As long as some individuals are willing to contribute in the face of small costs—perhaps altruistically—then a collective pool of information can be produced (Cheshire, 2007). The crucial point is that if the number of people in a given network is large enough, then it can dramatically change the potential for emergence because it becomes progressively more likely that altruistic actors will be present as group size increases (Oliver and Marwell, 1988). The Internet often creates a favorable situation for finding or attracting such individuals because it facilitates substantial network sizes that are improbable in face-to-face interactions.

Finally, an intriguing feature of many online information pools is that individuals can only observe cooperative behavior; they cannot necessarily see evidence of non-contributions (defection-like behavior in game-theoretical terms or free-riding in collective action). For example, when individuals visit online message boards or user-created content sites such as Wikipedia, contributions are visible, but users have no way to ‘see’ the vast majority of individuals who view, use, or derive value from the information without contributing. The same is true in online message boards,

where individuals see those who post and reply to messages, yet the number of 'lurkers' (individuals who read without posting) remains unknown. Though it may be possible for knowledgeable users to infer some knowledge about non-contribution, or 'lurking' behavior (see: Joyce & Kraut, 2006; Rafaeli et al., 2004), doing so may require expert knowledge of the system and impose a cost in time and effort.

The lack of evidence about non-contribution behaviors is important for several reasons. First, the examples above illustrate that the design of a given system can partially or completely control the availability of information about contributors and free-riders. By choosing which information to provide and which to withhold, designers can often influence the future behavior of participants. Second, when information about others in the system is withheld, it can introduce additional barriers to participation. The critical mass hypothesis in computer-mediated communication highlights the importance of knowledge about the decisions of others when making one's own choice about using a given communication system. As Markus (1987) argues, computer-mediated systems can create a public good if a critical mass of users adopts a given system. However, the likelihood of reaching a critical mass is often tied to the availability of information about adoption rates in a given population (i.e., individuals are more likely to adopt if they believe that others are adopting the system, technology, etc).

Kollock (1999) argues that the unique properties of digital information can create public goods, even if there is only a single contributor. Because information goods can be non-rival and easily replicated, many individuals can benefit from a single contribution (i.e., information posted to a website, a song made available for download, etc). While the idea of the critical mass may at first seem to be in conflict with Kollock's (1999) notion of digital goods as public goods, the two arguments speak to different levels of analysis. A single digital good can create a public good on the Internet, yet a given *system* of distributing digital information goods can generate a public good as well. For example, an Internet bulletin board system (BBS) can become a valuable public good once enough individuals use the system (Rafaeli and LaRose, 1993). Critical mass explanations for adoption behavior in computer-mediated systems clearly apply to interactive communication systems such as email systems, BBS's, online newsgroups, and other virtual communities (see also: Jones and Rafaeli 2000). However, the critical mass arguments in computer-mediated systems are primarily concerned with *adoption* of technologies and information systems (Markus 1987), and are therefore less relevant for explaining individual contribution behaviors within a system.<sup>3</sup>

In sum, the unique nature of information helps us understand why the costs associated with contributing information goods are low compared to physical goods, and in turn how these low costs can influence the decision to contribute or not. Thus far, the differences between information and physical goods primarily inform the study of the factors that encourage initial contributions of information. We now turn to a second, related issue which is the primary focus of this study: how can we

encourage *repeat* offerings to an information pool once an individual has already made at least one contribution?

### **Social Psychological Motivations for Contributing to Information Pools**

In his landmark theory of collective action and public goods, Olson (1965) argued that one solution to the free-rider problem is to offer selective incentives to those who contribute to a public good. This solution leverages the excludability of some benefits, while still maintaining the non-excludability and non-rivalrous nature of the public good. Though effective at encouraging contributions, selective incentives are expensive to produce (Oliver, 1980). For example, handbags and flashlights might encourage contributions to public radio—but someone still has to pay for these gifts (or their costs must be subtracted from money made through donations).

One resolution to the production and cost issue associated with selective incentives is to focus on social psychological processes rather than monetary benefits. Social psychological incentives are *intrinsic* benefits that individuals experience when they make contributions to an information pool and to public goods more broadly. Although these social psychological benefits may be very small, the low-cost situation created in many information pools can allow these processes to have a relatively significant impact on behavior (Cheshire, 2007). Recent empirical work in this area indicates that at least four key processes can encourage contributions of information in online settings: uniqueness of contributions, goal-setting, social approval, and the observation of cooperative behavior.

Ling et al. (2005) used field experiments involving members of an online movie recommendation community, MovieLens, to study the effects of different system features on the outcome of information contributions. In this case, the contributions were written recommendations/reviews of movies. The researchers found that participants tended to contribute more when they were specifically reminded of the uniqueness of their contributions. In addition, individuals contributed more information when they were told that they were less similar to others (again, reflecting a tendency to contribute more when identified as unique). In a follow-up experiment, the researchers also found that email messages reminding participants to contribute produced more contributions when the email identified the uniqueness of the individual, as compared to when the email reminder did not contain any information about individual uniqueness. Overall, the results of this study provide support for the motivational benefits of the perception of uniqueness when contributing to an information pool.

Ling et al. (2005) also investigated the difference between individual goal-setting and group goal-setting. Using the same online movie recommender system as the earlier experiments, the researchers found that individuals who are given specific goals contribute more than those who are given 'do-your-best' type goals. However, they also found that group-oriented goals were actually more effective than individual-oriented goals at motivating contributions. Although this finding was counter to

their initial predictions, it may be explained by the in-group/out-group effect (i.e., Tajfel and Turner, 1986). Specifically, when individuals perceive an in-group distinction, it may lead to higher degrees of commitment to their own group, as well as an increased desire to reinforce the in-group identity through participation (or in this case, increased contributions) (see also: Ludford et al., 2004). Ling et al. (2005) use the notion of social facilitation (e.g., Zajonc, 1965) to help explain their findings. Social facilitation suggests that individuals who perceive that their contributions can be evaluated by others will be likely to increase their contributions towards the group task. Whether we interpret this finding as a result of in-group/out-group salience, or as a product of social facilitation, the evidence from Ling et al.'s (2005) research supports the idea that group-oriented goals can result in higher contribution rates than individual-oriented goals when producing an information pool.

Using a series of controlled laboratory experiments, Cheshire (2007) examined the effects of social approval and observation of cooperative behavior on contributions to a different type of information pool than the one used by Ling and her colleagues. The experiments in the Cheshire (2007) study were designed to create a system much like a peer-to-peer music swapping system on the Internet. Participants created a list of their favorite songs, books, and movies to use as their own 'digital goods' in the experiments. These digital goods had economic value to the participants, but there was a cost associated with sharing one's own information. Individuals chose whether or not to contribute one of their digital goods to the collective pool of information goods, creating a basic social dilemma game.

The results of the Cheshire (2007) study showed that when individuals were told that a high percentage of users liked their last contribution, it had a strong, significant impact on continued contributions. Furthermore, when individuals were told that a low percentage of users liked their last contribution, it also had a strong, significant impact on contributions. Post-questionnaire responses indicated that participants continued to contribute to the information pool even when they received low social approval because they wanted to keep trying to *raise* their social approval. In this case, the social approval rating was always dependent on only the last contribution and was only known to the individual contributor—meaning that it could not function as any kind of permanent, public reputation. Thus, as a completely internal and social psychological effect, social approval had a strong impact on contribution behavior simply by informing individuals about how much others liked their last contribution to the system.

In a different experimental condition, Cheshire (2007) also examined the effect of observing cooperative behavior on sharing. When participants were told that high percentages of others were currently contributing, these individuals contributed slightly more at the beginning of the study than those who were not told anything about the amount of cooperative behavior. The effect was short-lived, however, and contributions sank over time in the high observed cooperation condition. In a condition in which individuals were always told that low percentages of others were currently sharing, individual contributions were not significantly better than the

when no information about cooperative behavior was provided. Thus, observing high cooperative behavior only had an effect in the earliest stages of participation—when an individual is first deciding whether or not they should contribute. Over time, this incentive was not strong enough to maintain higher contributions compared to a situation where individuals had no information about the contributions of others. Thus, while there was evidence that individuals made decisions about contributing/not contributing based on the behavior of others (or their perception of such behavior), at the two extremes this effect was not significant beyond the initial decisions to contribute or not to the information pool.

In addition to the effects of specific incentives on contribution behavior, Sohn and Lekenby (2007) demonstrate that the *structure* of an information-sharing community has an independent effect on the quantity of contributions made by individuals in these systems. Sohn and Lekenby's (2007) approach is partially a response to the dominant utilitarian and normative explanations of the social dilemma in information sharing systems (such as those reviewed above). Sohn and Lekenby argue that most studies of online information systems and virtual communities assume a pooling structure (2007: 436), but computer-mediated systems can take many network forms that do not necessarily have to create a single information pool. Using an experimental design, they demonstrate that when individuals are able to exchange directly with one another, they make significantly more contributions compared to those who contribute to a pooled information resource. The researchers theorize that increased personal responsibility and contribution efficacy in the non-pooled structure explains the difference in contribution rates. Importantly, Sohn and Lekenby (2007) demonstrate that one's motivation to contribute information in a given system may be partially explained by the incentive structures in a given information exchange system.

### **Encouraging Contributions through the Social Psychological Effects of Interactive Feedback**

In this paper, we define feedback as the interactive process in which information is returned in response to a contributor's action.<sup>4</sup> Previous research has shown that providing some form of response when an individual contributes can motivate repeated contributions. For example, Joyce & Kraut (2006) found that when first time posters to online newsgroups received responses to their messages they were 12% more likely to post again. In such a case, the initial posting is the action that precedes feedback (e.g., responses), thereby affecting future behavior (increased postings by the original author). In many cases, simply receiving any response is enough to prompt additional action by the original contributor.

The studies reviewed above examine the social psychological effects of both synchronous and asynchronous feedback. Synchronous feedback immediately follows an action (i.e., an instant message thanking individuals for their current contribution), while asynchronous feedback occurs at a later time, perhaps as an email

or other message. In some cases, the feedback is direct, clearly and unambiguously coming from the same system to which the individual has made a contribution. In other cases, the feedback could be indirect, such as a note from a 3<sup>rd</sup> party encouraging contribution to one or more systems. In this study we examine the effects of direct, synchronous feedback on continued contributions to an information pool. Thus, we focus on immediate incentives which might encourage repeat contributions for those who have already made at least one contribution.

We examine three types of feedback: *Gratitude* for providing a contribution, a *Historical Reminder* of one's entire contribution record, and the *Relative Ranking* of one's contributions compared to others. We argue the three types of feedback provide intrinsic, social psychological benefits to the contributor. As current theory and research demonstrates, these effects may be very small—but the extremely low costs associated with contributing information in online systems can allow these benefits to influence behavior nonetheless (see: Kollock, 1999; Ling et al., 2005; Cheshire, 2007). Here, we briefly describe the types of feedback in this study and make several predictions about how they affect contribution behavior.

*Gratitude.* Positive emotional responses are an essential part of everyday social exchange processes. The psychological effect of a mild, positive emotion such as gratitude can produce moderate changes in behavior depending on the relative costs involved in a situation (Lawler, 2001: 322). Low-cost situations such as information pools should allow the small, positive effects of gratitude to have a relatively significant impact on contribution behavior. For example, Beenen et al. (2004) found that sending out a one-time 'thank-you' email to potential contributors in their MovieLens study created a short-term spike in contributions. In this study, we expand the notion of gratitude from a one-time response to a synchronous acknowledgement following every contribution made by an individual.

*Historical Reminder.* A historical reminder simply informs individuals about their prior contribution behavior. This type of feedback is dynamic since it gives updated information based on past behavior. An individual's own past behavior has been shown to be a predictor of future contributions in collective action problems (Gunnthorsdottir et al., 1999; Wilson and Sell, 1997). Unlike many controlled experimental settings, however, in real-world information pools the decisions to contribute or not may not be within a single defined time period (such as an experiment that lasts up to one hour). In the real world, an individual may contribute to an information pool contained within a website, only to visit the same website again many days, weeks, or months later. Thus, individuals could forget how much or how often they have contributed to a given system over time. We argue that a historical reminder encourages contributions because it prompts an individual to consider his or her own past contribution behaviors, making them salient in the current interaction.

*Relative Ranking.* As previous research has shown, aggregate knowledge about how much others are contributing can have a positive influence on an individual's likelihood of contributing to a public good (e.g., Sell, 1997; Cheshire, 2007). However, research has also shown that very specific knowledge of others' past contribution



behavior (e.g. reputations) can actually reduce contributions to a public good (Wilson and Sell, 1997). Thus, knowledge about cumulative group behavior can be beneficial to the production of a public good, while individual reputation information may actually reduce contributions in some situations.<sup>5</sup>

A relative ranking is aggregated information that combines one's own behavior with the average behavior of others. It is intended to stimulate the importance of one's behavior relative to all others who have contributed at least once to the system. The critical mass arguments in computer-mediated interaction could apply to the relative ranking since this form of feedback highlights an individual's actions in light of others' behavior. However, a relative ranking does not necessarily indicate actual numbers of contributors. Rankings that do not report absolute numbers of participants (i.e., percentages or proportions) make it difficult or impossible to know the number of other contributors.

### Context of Interaction

In addition to the independent effects of feedback, we argue that context of interaction is a key factor for understanding incentives in information pools. The context of interaction includes all of the characteristics of the situation in which an individual chooses whether to participate in an information pool or not. There are numerous ways to make distinctions between interaction situations in online settings, but we focus on a dichotomous division between an *internal website* which is directly associated with an information pool and an *external website* that is unrelated to the information pool but which provides opportunities for contributions.

To refine the difference between internal and external contexts of interaction, consider the following scenarios. Individuals who contribute to an information pool through an internal website are presumably visiting the site because they intend to examine the system, learn more about it, or just try it out. They may have learned about the site through a blog posting or a link on a colleague's website, for example. The Flickr<sup>6</sup> homepage is an 'internal' site by our definition—the website itself is the public representation of the information pool contained within (e.g., photograph sharing). It is reasonable to assume that those who visit the site are relatively focused on it and that they are there because of some level of pre-existing interest or curiosity. On the other hand, many otherwise unrelated websites allow individuals to contribute to, browse, or search Flickr through one of a variety of small, self-contained 'widgets' which are embedded in that website's content. The individuals who contribute through the external (unrelated) websites may have come across this opportunity to contribute information without knowing much, or anything, about Flickr, instead simply noticing the widget and casually engaging with it before returning to their original activities. In this case, it is reasonable to assume that an individual's focus is centered on the primary website that they are visiting, and the opportunity to interact with another system through this website is a secondary, peripheral activity.

## Hypotheses

Our first set of hypotheses simply states that each of the three types of feedback should lead to higher contribution rates compared to situations in which no feedback is provided. Each feedback incentive provides a positive social psychological benefit to the contributor, which should have a positive effect on repeat contribution behavior. These include positive emotional response (gratitude), salience of past behavior for current behavior (historical reminder), and salience of one's behavior compared to others (relative ranking). We argue that the positive effects should exist for those who contribute through external sites as well as those who contribute through internal sites. Thus, we make the same predictions for the external and internal contexts of interaction.

**H1a-1c:** For external contributors, providing feedback in the form of a Gratitude (1a), Historical Reminder (1b), and Relative Ranking (1c) will increase overall contributions compared to a situation in which no feedback is provided to the contributor.

**H2a-2c:** For internal contributors, providing feedback in the form of a Gratitude (2a), Historical Reminder (2b), and Relative Ranking (2c) will increase overall contributions compared to a situation in which no feedback is provided to the contributor.

In a direct comparison of internal versus external contributors, we expect internal contributors to exhibit a higher overall contribution rate compared to external users. Those who interact through an internal site may be more likely to identify with—and want to contribute to—a given information pool compared to those who contribute entirely through external sites. Individuals who are visiting an internal website are more likely to be focused on the information pool as their primary activity. On the other hand, individuals who contribute from an external site are presumably visiting that site for an unrelated purpose—making contributions to an information pool through the external site is likely to be a peripheral activity. Thus, while these individuals may make some contributions through the external site, they are likely to resume their interaction with the external website rather than continually contribute to the information pool.

**H3:** All things being equal, contributors who interact through the internal website will have higher repeat contributions compared to those who contribute through external websites.

## A Field Experiment Using the Mycroft System

For the experiments we report here, we worked with a small start-up company to use their custom-built Internet system called Mycroft. Mycroft is a web-based network which allows large tasks to be widely distributed and the results efficiently collected so that thousands of individual contributors can work on the same project at the same time. The system applies the peer production model (Benkler, 2006) to a wide variety of problems—creating information goods by combining and synthesizing the

efforts of many people who each contribute one small piece of information at a time. Mycroft accepts large jobs which cannot technically or efficiently be completed by computers and breaks them down into many constituent parts called ‘puzzles.’ The puzzles are distributed via banner ads on existing websites, in place of traditional advertising materials (see Figure 1). As each puzzle is answered, the results are combined with others at successively larger levels until the top-level job is complete.

Mycroft banners exist on a primary website and on many diverse websites. The banners are designed to facilitate entirely self-contained interactions. Users who visit a blog, shopping, or news site, for example, may find a Mycroft banner there and contribute to the information pool without leaving that page. Individuals can contribute anonymously, or they can choose to register and provide more information about themselves. This architecture takes advantage of ‘casual’ contributions, turning collective action into something that individuals might do while they are surfing from one website to the next. In addition, individuals could visit the primary Mycroft website to make contributions.

### Data Sources & Measurement

Between March 2006 and October 2006, thousands of people viewed Mycroft banners more than 100,000 times. The Mycroft banners were hosted on more than 20 diverse volunteer websites, which we will refer to as external sites. In addition, an internal home website hosted the banners. In both cases, the Mycroft banners appear near the top web pages in locations where commercial banner advertisements might normally exist.

Our sample of users comes from Internet users who either: 1) heard about and visited the Mycroft home website through email announcements, mailing lists, blog postings and word-of-mouth, or 2) interacted with Mycroft banners on one of the external host sites. Announcements and commentaries about Mycroft were posted by bloggers and journalists on a variety of websites, and usually provided a link to the Mycroft website that interested readers could follow. No additional incentives were provided for visiting the Mycroft website. Once directed to a page on the Mycroft home website, potential contributors could interact with Mycroft banners included on each page. Individuals who viewed Mycroft banners on an external site, on the other hand, did not visit the Mycroft website itself. All users, internal and external, could click on a, “What’s This?” link in the Mycroft banner interface to learn more about what Mycroft was and how it worked.



Figure 1 The Mycroft Banner Interface.

In our study, the main website and the external host sites serve as the two contexts of interaction. We do not examine the factors that influence an individual's decision to visit the main website, yet we acknowledge that internal contributors were notified about the Mycroft project in some way that led to their navigating to the home website while the external contributors were not. As these are essentially two different sample populations, we separate and compare these groups in all of our analyses. In addition, we discuss the implications of the two different samples in our results and discussion. Our final valid sample includes 791 individuals who contributed *at least once* to the system during the data collection period (467 internal contributors and 324 external contributors).

For the test period in this study, the Mycroft system distributed puzzles that asked individuals to type in small portions of text. Once an individual made a contribution, they were assigned to receive either one of the three feedback types or no feedback at all (the control condition). By structuring our intervention in the form of feedback provided immediately after a contribution, influence is limited to those who have already contributed. This allows us to test what we call *retention incentives*, or feedback mechanisms that encourage individuals to continue contributing over time. This is in contrast to *capture incentives*, which are aimed at collecting initial contributions. Since we do not provide any manipulations before a contribution is made, we are not concerned with capture incentives in this study.

To measure retention, our standard dependent variable is the repeat contribution rate (RCR). The RCR is the average number of contributions that a given individual makes over all contribution sessions. For our analyses, a session begins the first time an individual contributes to a Mycroft banner, and ends either when the individual leaves the website or after a continuous hour of inactivity. Similar to most online contribution systems, a graph of contributors to Mycroft follows a power law function – a few individuals contribute a large amount and most individuals contribute very little or not at all.<sup>7</sup>

In order to account for multiple sessions by the same user, Mycroft banners use an industry-standard mechanism to identify users called 'cookies.' When an individual's web browser loads a banner, Mycroft writes a small amount of text to the user's computer. This identification system does not collect or return any additional information about the user to Mycroft.<sup>8</sup> The sole purpose of the 'cookie' is to uniquely identify those who contribute to Mycroft over multiple sessions. We obtained aggregate counts of contributions by randomly-generated ID's. If an individual contributes to a Mycroft banner at any time in the future using the same web browser, we recorded this as a separate session for the same user.

### Experimental Manipulations

Individuals who made at least one contribution were randomly placed into one of four feedback conditions (three experimental, one control). In the experimental conditions, individuals consistently received one of the three types of feedback: gratitude, historical reminder, or relative ranking. We also included a control


condition in which no feedback was given to the contributor. We ran each of these four conditions in the two different contexts of interaction (internal website and external websites), producing a total of eight conditions.

The gratitude feedback took the form of a simple static ‘thank you’ presented to the individual immediately after a contribution. The message was always the same and it was given after every contribution. The historical reminder feedback presented a simple count of the number of times an individual had contributed (e.g. ‘You have contributed 12 times’). The number was dynamic, counting upwards each time a contribution is made. Finally, the relative ranking feedback was operationalized as a percentage ranking of one’s contributions compared to all other current contributors. For example, a first-time contributor might see that they are in the top 99% percent of all contributors (i.e., most people have contributed more than they have), while a frequent contributor might find herself in the top 5 or 10 percent of all contributors. See Figure 2 for examples of the feedback mechanisms as they appeared to the participants.

## Results

Table 1 displays the descriptive statistics for key variables by context of interaction (internal versus external sites). Generally, these results already begin to show a pattern of higher contributions among internal contributors compared to the external contributors. The average number of sessions, however, is quite similar between the internal and external sites (1.13 and 1.08, respectively).

Our first set of hypotheses (1a-1c) predicts that, for individuals who contribute through the external websites, the three forms of feedback should produce higher



**Figure 2** Examples of the Three Types of Feedback (Gratitude, Historical Reminder, and Relative Ranking).

**Table 1** Descriptive Statistics by Internal and External Sites

Variable	Site	Mean (S.D.)	Min	Max
Contributions Per Session	Internal	4.25 (5.0)	1	44
	External	1.99 (2.68)	1	21
Contributions Over All Sessions	Internal	4.97 (6.84)	1	58
	External	2.41 (6.02)	1	97
Number of Sessions	Internal	1.13 (.60)	1	8
	External	1.08 (.55)	1	8

$N = 791$ ; 467 Internal, 324 External.

levels of contributions than when no feedback is present. Table 2 shows the ANOVA results for types of feedback on the repeat contribution rate among external contributors. The main effect of feedback is significant,  $F(3, 320) = 3.24, p < .05$ . Furthermore, post-hoc comparisons indicate that the repeat contribution rates are significantly higher for all three of the feedback conditions compared to the control condition ( $p < .05$ ). Thus, Hypotheses 1a, 1b and 1c each receive support.

Our second set of hypotheses also predicts a positive effect for the three feedback mechanisms among those who contribute through the internal website. Table 3 shows the ANOVA results for feedback type on repeat contributions for the internal contributors. The overall effect of feedback is borderline significant,  $F(3, 463) = 2.06$ ,

**Table 2** Analysis of Variance of Type of Incentive on Repeat Contribution Rate (External Contributors)

<i>Tests of Between-Subjects Effects</i>				
Source	Type III Sum of Squares	Df	Mean Square	F
Feedback Incentive	68.5	3	22.8	3.24*
Intercept	1261	1	1261	179.2**
Error(Trial Block)	2252	320	7	

$N = 324$ .

\*\* $p < .01$ . \* $p < .05$ .

**Table 3** Analysis of Variance of Type of Incentive on Repeat Contribution Rate (Internal Contributors)


<i>Tests of Between-Subjects Effects</i>				
Source	Type III Sum of Squares	Df	Mean Square	F
Feedback Incentive	154.4	3	51.5	2.06 <sup>+</sup>
Intercept	8463.6	1	8463.6	338.52**
Error(Trial Block)	11541.9	463	24.9	

$N = 467$ .

\*\* $p < .01$ . <sup>+</sup> $p = .10$ .

$p = .10$ , and post-hoc tests show a moderately significant difference between the relative ranking and control condition (mean difference =  $.96$ ,  $p = .10$ ). None of the other feedback types are significantly higher than the control condition. As Figure 3 shows, the gratitude feedback actually appears to be *lower* than the control condition (though the difference is not statistically significant). Interestingly, the relative ranking condition produces significantly higher contribution rates compared to the gratitude condition (mean difference =  $1.70$ ,  $p < .01$ ). Thus, Hypothesis 2c receives modest support and Hypotheses 2a and 2b are not supported.

Finally, Hypothesis 3 predicts that those who interact through internal webpages will have higher repeat contributions compared to external contributors. Table 4 shows the ANOVA of the feedback types and context of interaction on repeat contribution rates. The main effect of being an internal contributor is highly significant,  $F(1, 782) = 35.47$ ,  $p < .001$ ). Indeed, the average repeat contribution rates in all conditions are higher for internal contributors than for those who interact through external sites (See Figure 3). Thus, Hypothesis 3 receives strong support.



**Figure 3** Mean Contributions Per Session by Feedback and Context of Interaction.

**Table 4** Analysis of Variance of Type of Incentive on Repeat Contribution Rate (Internal and External Contributors)

<i>Tests of Between-Subjects Effects</i>				
Source	Type III Sum of Squares	Df	Mean Square	F
Feedback Incentive	151.5	3	50.5	1.78 <sup>+</sup>
Internal	1009	1	1009	35.47***
Incentive * Internal	128.8	3	42.9	1.51
Number of Visits	10961.7	1	10961.7	385.38***
Intercept	1977.9	1	1977.9	69.54***
Error(Trial Block)	22243.3	782	28.4	

$N = 791$ .

\*\*\* $p < .001$ . \*\* $p < .01$ . + $p = .10$ .

## Discussion

The field experiments in this study produce several clear findings. First, each of the three feedback mechanisms (gratitude, historical reminder, and relative ranking) have a significant impact on repeat contributions among those who interact with the Mycroft banners through an external website. This type of situation (interacting with a banner on an unrelated website) is representative of much real-world activity on the Internet where various forms of content are displayed in or around the periphery of unrelated websites. Thus, these results may generalize to similar situations in which individuals contribute via small areas of larger websites (perhaps through banners, frames, or embedded widgets). Given the growing interest in user-created content<sup>9</sup> on the Internet, even a modest increase in repeat contributions associated with these types of feedback mechanisms could have a substantial effect on the production of information goods in systems with large numbers of users.

The second major finding from this study is that the context in which an individual interacts with an online system has an important influence on the effectiveness of some types of feedback. While we found significant, positive effects for all three types of feedback in the external website conditions, we found only a modest effect for one type of feedback (relative ranking) in the internal website conditions. The internal contributors clearly contributed more overall than the external contributors across all conditions—so what might account for this discrepancy in the effectiveness of feedback incentives?

We believe that the non-finding among the internal website contributors may point to a far more interesting set of issues between retention incentives and the context of interaction. In some collective action problems, individuals create *collective identities* which help explain the mobilization of participation in the absence of other incentives (Polletta and Jasper, 2001). The collective identities are built from shared beliefs and goals, even if these beliefs and goals are not explicitly known. Collective identities tend to build affective connections—promoting norms of



reciprocity and perhaps even obligation—to those who disseminate information about a collective action effort (Polletta and Jasper, 2001). Although the notion of the collective identity has often been employed in socio-political movements (e.g., Snow et al., 1980), the concept is relevant to many public goods problems in which individual motivations may not be entirely explained by explicit costs and benefits associated with contributions. In our study, individuals who interacted through the internal site were *directly recruited* through emails, solicitations, and other online announcements that were disseminated from the researchers' network of relationships (e.g. personal blog sites, mailing lists, and links on acquaintances' websites). Thus, one possibility is that individuals who chose to come to the internal websites may have been demonstrating a type of obligation or other normative response from our requests for participation. External contributors, on the other hand, were never solicited in advance for their participation—they simply saw the Mycroft interface on various websites and made a decision to contribute or not. Thus, the external contributors are far less likely to have any sense of normative obligation compared to those who visited the internal site.

Given the situation described above, if individuals were already contributing (or planning to contribute) through obligation, pre-existing motivations or normative expectations, then any *additional* inducements may conflict with the intrinsic motivation that brought them to the site in the first place. In situations where intrinsic motivation is high, applying extrinsic motivations (such as our feedback mechanisms) can counteract the effects of the intrinsic motivation and lower overall contribution behavior. Psychologists have pointed out that intrinsic and extrinsic motivations are sometimes not additive but substitutive – a trade-off that can reduce motivation especially in cases where intrinsic motivations are high. This effect has been termed the 'hidden cost of reward' (Lepper and Greene, 1978) or the 'corruption effect of extrinsic motivation' (Deci, 1971, adapted from Osterloh & Frey, 2000). Some economists have also labeled this phenomenon 'crowding out' (Frey, 1997; Osterloh and Frey, 2000). Ironically, the same incentives that help to promote contributions from individuals with little or no pre-existing motivation can undermine contributions for those who may have been planning to contribute anyway.

The differential effect of feedback by context of interaction could be particularly salient for the design of Internet systems that depend on user-generated content that comes from a variety of sources. Feedback systems may be useful for encouraging repeat contributions among users who interact with a system through mechanisms on dispersed websites (i.e. external contributors in our study). However, for individuals who are directly recruited and encouraged to participate, or more generally for individuals who have a strong sense of collective identity, providing some types of feedback mechanisms may be counter-productive. The additional feedback incentives could undermine the user's existing motivations to contribute more, producing nearly the same number of repeat contributions than if no feedback mechanisms were in place. In some cases, these feedback incentives could actually reduce repeat contributions, as demonstrated through our surprising (though non-statistically

significant) result of slightly lower repeat contributions for those who received the gratitude feedback versus those who received no feedback at all in the internal sample.

It might be misleading to assume that all forms of feedback are ineffective for those who only contribute through an internal website. Of course, those who regularly contribute to information pools are among the most interesting (and perhaps arcane) users for designers of information pools—precisely because they often provide the vast majority of contributions. For example, in many information pools, a minority of individuals produce most of the desirable content (Adar and Huberman, 2000; Lyman and Varian, 2003; Swartz, 2006). Thus, while the feedback mechanisms in this study may not make much of an impact on internal participants, what is needed is an understanding of the unique motivations and needs of this type of contributor. As current research is beginning to show, the behavior of individuals who often make the most contributions seems to be tied to pre-existing characteristics and role differentiation (Welser et al., 2007). Future research should continue to examine and differentiate these users from the larger pool of casual contributors, while recognizing that incentives for ongoing contributions may work differently between these user types.

### **Limitations and Directions for Future Research**

Some important factors should be considered regarding these field experiments. First, sampling and sample characteristics are difficult or impossible to capture, which is a common issue for online research. (See: Bainbridge, 2002). Compared to traditional telephone and in-person surveys, it is more difficult to define a bounded sample and achieve high response rates in an open Internet study such as this one. Indeed, we did not restrict who viewed or interacted with Mycroft banners in any way. We chose not to collect any survey data before a contribution was made because it would interfere with the real-world nature of the experiment (e.g. a survey initiated by clicking on the banners would be suspicious, intrusive, and would create additional barriers to contributing). Furthermore, we did not want to collect survey data once an individual made one or more contributions because this might confound the experimental manipulations (e.g., asking demographic or attitudinal questions might interact with the likelihood of future interactions with the Mycroft system). Although any user could register with Mycroft and provide simple demographic information, very few participants chose to do this (< 2% of all users). For these reasons, we know very little about the characteristics of the sample population because no additional information was collected aside from contribution behavior.

Mycroft also incorporated a robust mechanism for tracking users across many interactions on multiple websites regardless of whether they chose to register or participated anonymously. However, no such mechanism, outside of a controlled laboratory setting, is completely free of potential problems. Technical limitations

associated with an open Internet system mean that, in rare cases, a participant might not have been persistently assigned to an experimental condition, or that their contributions might be counted as a new participant.<sup>10</sup> However, our logging information allowed us to correct and eliminate data in cases where problems could be identified. In addition, we have taken every possible precaution when analyzing the data to account for and eliminate redundant and erroneous information.

This study shows clear differences between different contexts of interactions and how feedback mechanisms can encourage repeat contributions in one type of information pool. We believe that our findings generalize to other types of information pools that are highly ordered, coordinated, and produce a public good.<sup>11</sup> Although we are ultimately interested in building theory and empirical evidence on the effects of feedback incentives and social psychological processes on contributions to many different types of information pools, it is important to note that the findings from this study may not necessarily generalize to *all other types* of information pools. Not all information pools produce public goods, and the structural affordances and limitations of different systems vary widely (see: Sohn and Leckenby, 2007). Diverse structural conditions, as well as the emergent (and perhaps unintended) behaviors of users, may influence the appropriateness and effectiveness of various feedback incentives on different outcomes (Antin and Cheshire, 2007). Thus, we must identify and acknowledge such differences as we analyze and compare the wide range of Internet systems that create information pools.

The combined results of this study demonstrate that the context of interaction matters for the effectiveness of social incentives provided through feedback mechanisms in computer-mediated exchange systems. In addition, retention incentives such as those used in this study are useful for encouraging repeat contributions from 'casual' types of contributors. Much theoretical and empirical work will be necessary in order to specify and test the effects of feedback mechanisms in different contexts of interaction. Field experiments are an essential part of this goal, combining theoretically-motivated experimental manipulations with the ecological validity afforded by a real-world setting. One of our current directions is to investigate other factors such as *interpersonal trust networks* (Cheshire and Cook, 2004), especially as they relate to user-generated content systems. In addition, we aim to study motivations in information pools that have emergent structures of order and coordination (Cheshire and Antin, 2008).

For each advantage provided by using field experiments, there are also a number of questions that can only be addressed with tight experimental controls and other forms of standardization. One of the next steps in this line of research is to combine laboratory experiments with field studies such as this one, addressing concerns about uniformity, control, and generalizability. In addition, an important limitation that we aim to address in future studies is the interaction between multiple forms of feedback. Combinations of feedback incentives, the addition of asynchronous mechanisms, and personalized feedback messages<sup>12</sup> may also provide crucial insight into contribution behaviors. Finally, as current research indicates (e.g., Sohn and Leckenby,

2007), future studies should consider different structural arrangements and their effect on contributions of information in online settings. Our understanding of motivations in information-sharing systems depends on a more nuanced understanding of different network structures (i.e., pooled versus direct-exchange networks), types of interactive feedback and different contexts of interaction.

## Notes

- 1 In collective action problems, benefits are often divided into types that include individual consumption (private goods), closed-group consumption (club goods) and unrestricted public consumption (public goods). We primarily focus on Internet information pools that produce public goods, but all three types of goods are possible.
- 2 Though related, non-rivalry and replication are two distinct properties of some public goods. It is possible to have non-rival public goods that are not necessarily replicated (e.g., clean air). The opposite may or may not always be true, as pure replication arguably tends towards non-rivalry. One of the key distinctions between digital information and most physical goods is that, as a public good, information is often non-rival and replicable at the same time.
- 3 We draw a distinction here between arguments of critical mass in information technology adoption (i.e., Markus 1987) and a related literature on critical mass of collective action in sociology (i.e., Granovetter 1978, Oliver and Marwell 1988). The latter tradition is more closely aligned with our arguments regarding ongoing contribution behaviors in collective action systems, rather than the former's concentration on technology adoption.
- 4 As one anonymous reviewer notes, the term 'feedback' implies an underlying interactivity between at least two actors. In our study, the 'actor' that provides feedback is actually a computer system that is designed to give responses based on user input.
- 5 As Wilson and Sell (1997) acknowledge, more empirical research is needed to tease out the discrepancy between personal reputation information and knowledge of aggregate cooperation behavior. Still, their (1997) research shows a clear, negative impact of 'too much information' about the behavior of others on the production of a public good.
- 6 Flickr (<http://flickr.com>) is a website that allows individuals to upload and share digital photos. Users can also tag photos with keywords, search and discuss them, find other photographers, and form social groups.
- 7 As we are interested in repeat contributions, visitors who never made any contribution are eliminated from our analyses.
- 8 Individuals who visited the main Mycroft website were able to optionally create an account, but very few users chose to do so.
- 9 This is often referred to in the popular press as 'Web 2.0' (see: O'Reilly 2005).
- 10 Although we used Internet browser 'cookies' to track repeat users, this method is subject to a few limitations. These cookies identify browsers, not individuals, so if several individuals use the same computer and web browser without logging into Mycroft their results would be combined. Note also that cookies are tied to individual browsers, so if a contributor uses several different browsers or computers, or uses any number of programs that delete cookies, our ability to track contributors is lost.

- 11 See Cheshire & Antin (2008) for a classification system of different types of information pools.
- 12 We thank our anonymous reviewers for providing some of these suggestions.

## Acknowledgements

We would like to thank Benjamin Hill for his work with the Mycroft system and several anonymous reviewers for their feedback and suggestions.

## References

- Adar, E. & Huberman, B. (2000). Free riding on Gnutella. Xerox PARC, Palo Alto, CA.
- Antin, J. and Cheshire, C. (2007) Designing Social Psychological Incentives for Online Collective Action. Conference Paper. General Online Research (GOR), Leipzig, Germany.
- Bainbridge W. S. (2002). Validity of web-based surveys. In: Orville Vernon Burton. ed. *Computing in the Social Sciences and Humanities*. University of Illinois Press, Urbana, pp. 51–66.
- Beenen, G. (2004). Social facilitation and assigned goals in an online community. Tepper School of Business, Carnegie Mellon University, Pittsburg, PA.
- Benkler, Y. (2006). *The Wealth of Networks*. New Haven, CT: Yale University Press.
- Cheshire, C. (2007). Social Psychological Selective Incentives and the Emergence of Generalized Information Exchange. *Social Psychology Quarterly*, 70, 82–100.
- Cheshire, C., and Antin, J. (2008). Order, Coordination and Uncertainty: A Field Study of Contributions to an Online Information Exchange System. Working Paper, UC-Berkeley School of Information.
- Cheshire, C. and Cook, K. S. (2004). The Emergence of Trust Networks: Implications for Online Interaction. *Analyse and Kritik*, 28, 220–240.
- Deci, E. L. (1971). Effects of Externally Mediated Rewards on Intrinsic Motivation. *Journal of Personality and Social Psychology*, 18, 105–115.
- Frey, B. S. (1997). *Not Just for the Money: An Economic Theory of Personal Motivation*. Cheltenham, UK and Brookfield, USA: Edward Elgar.
- Granovetter, M. (1978). Threshold Models of Collective Behavior. *American Journal of Sociology*, 83(6), 1420–1443.
- Gunnthorsdottir, A., Houser, D., McCabe, K., & Ameden, H. (1999). Excluding free-riders improves reciprocity and promotes the private provision of public goods. Working Paper, University of Arizona Department of Economics.
- Joyce, E., and Kraut, R. E. (2006). Predicting continued participation in newsgroups. *Journal of Computer-Mediated Communication*, 11(3), 723–747.
- Jones, Q., and Rafaei, S. (2000). Time to Split, Virtually: ‘Discourse Architecture’ and ‘Community Building’ as means to Creating Vibrant Virtual Publics. *Electronic Markets: The International Journal of Electronic Commerce and Business Media*. 10(4), 214–223.
- Kollock, P. (1999). The economies of online cooperation: Gifts and public goods in cyberspace. In M. Smith and P. Kollock. (Eds.). *Communities in cyberspace*, (pp. 220–239). London: Routledge.

- Lawler, E. J. (2001). An Affect Theory of Social Exchange, *American Journal of Sociology*, 107(2), 321–52.
- Lepper, M. R. & Greene, D. (1978). *The Hidden Costs of Reward: New Perspectives on the Psychology of Human Motivation*. Hillsdale, NY: Erlbaum Publishers.
- Ling, K., Beenen, G., Wang, X., Chang, K., Frankowski, D., Resnick, P., & Kraut, R. E. (2005). Using social psychology to motivate contributions to online communities. *Journal of Computer Mediated Communication*, 10.
- Lyman, P. & Varian, H. R. (2003). *How Much Information?* Retrieved November 7, 2006 URL: <http://www.sims.berkeley.edu/how-much-info>
- Ludford, P. J., Cosley, D., Frankowski, D., & Terveen. L. (2004). Think Different: Increasing Online Community Participation Using Uniqueness and Group Dissimilarity. (pp. 631–638) in *Proceedings of CHI 2004*. Vienna, Austria.
- Markus, M. L. (1987). Toward a “Critical Mass” Theory of Interactive Media: Universal Access, Interdependence and Diffusion. *Communication Research*, 14, 491–511.
- Oliver, P. E. (1980). Rewards and Punishments as Selective Incentives for Collective Action: Theoretical Investigations. *American Journal of Sociology* 85, 1356–75.
- Oliver, P. E. & Marwell G. (1988). The Paradox of Group Size in Collective Action: A Theory of the Critical Mass. II. *American Sociological Review*, 53, 1–8.
- Olson, M. (1965). *The logic of collective action; public goods and the theory of groups*. Cambridge, Mass.: Harvard University Press.
- O’Reilly, T. (2005). *What is Web 2.0*, Retrieved November 22, 2006 (<http://www.oreillynet.com/pub/a/oreilly/tim/news/2005/09/30/what-is-web-20.html>)
- Osterloh, M. & Frey, B. S. (2000). Motivation, Knowledge Transfer, and Organizational Forms. *Organization Science*, 11, 538–550.
- Polletta, F. & Jasper, J. M. (2001). Collective Identity and Social Movements. *Annual Review of Sociology* 27, 238–305.
- Rafaeli S., LaRose R. J. (1993). Electronic Bulletin Boards and ‘Public Goods’ Explanations of Collective Mass Media. *Communication Research*, 20, 277–97.
- Rafaeli, Sheizaf, Gliad Ravid, and V. Soroka. (2004). De-lurking in virtual communities: a social communication network approach to measuring the effects of social and cultural capital. in *Hawaii International Conference on System Sciences (HICSS 37) Collaboration Systems and Technology Track*. Big Island, Hawaii.
- Rafaeli, S. & Raban, D. R. (2005). Information sharing online: a research challenge. *International Journal of Knowledge and Learning*, 1, 62–79.
- Rafaeli, S., Raban, D. R., & Ravid G. (2005). Social and Economic Incentives in Google Answers. in *GROUP05-Sustaining Community: The Role and Design of Incentive Mechanisms in Online Systems Workshop*. Sanibel Island, FL.
- Rashid, A. M., Ling, K., Tassone, R. D., Resnick, P., Kraut, R. E., & Riedl J. (2006). Motivating Participation by Displaying the Value of Contribution. in *ACM Conference on Computer Human Interaction (CHI)*. Montreal, Canada.
- Sell, J. (1997). Gender, Strategies, and Contributions to Public Goods. *Social Psychology Quarterly*, 60, 252–265.
- Shapiro, C., Varian H. R. (1998). *Information rules: a strategic guide to the network economy*. NetLibrary Inc. Retrieved August, 2007.
- Sohn D., Leckenby J. D. (2007). A Structural Solution to the Communication Dilemmas in a Virtual Community. *Journal of Communication*, 57, 435–449

- Snow, D. A., Zurcher L. A., Eklund-Olson, S. R. (1980). Social networks and social movements: a microstructural approach to differential recruitment. *American Sociological Review*, 45, 787–801.
- Swartz, A. (2006). *Who Writes Wikipedia?* Retrieved December 8th, 2006 (<http://www.aaronsw.com/weblog/whowriteswikipedia>)
- Tajfel, H., and Turner. J. C. (1986). "The social identity theory of intergroup behavior." pp. 7–24 In S. Worchel and W.G. Austin (Eds.). *Psychology of intergroup relations* Chicago: Nelson-Hall.
- Thorn B. K., Connolly T. (1987). Discretionary Data Bases: A Theory and Some Experimental Findings. *Communication Research*, 14, 512–28.
- Welser, H. T., Gleave, E., Fisher, D., & Smith, M. (2007). Visualizing the Signatures of Social Roles in Online Discussion Groups. *The Journal of Social Structure*, 8(2). <http://www.cmu.edu/joss/content/articles/volume8/Welser/>
- Wikipedia. (2006). *Wikipedia Statistics*, Retrieved December 8th, 2006 (<http://stats.wikimedia.org>)
- Wilson, R., & Sell, J. (1997). Liar, liar...: Cheap talk and reputation in repeated public goods settings. *Journal of Conflict Resolution* 41(5), 695–717.

### About the Authors

Coye Cheshire ([coye@ischool.berkeley.edu](mailto:coye@ischool.berkeley.edu)) is an Assistant Professor at the UC Berkeley School of Information. His research focuses on how various forms of exchange are produced and maintained, especially in computer-mediated environments such as the Internet. His current projects investigate shifts in modes of social exchange, interpersonal trust-building, online relationship formation, and the application of social psychological selective incentives to collective action problems.

**Address:** UC Berkeley School of Information, 102 South Hall, Berkeley CA 94720-4600

Judd Antin ([jantin@ischool.berkeley.edu](mailto:jantin@ischool.berkeley.edu)) is a Doctoral Student at the UC Berkeley School of Information. His research focuses on understanding and testing social psychological incentives for online collaboration.

**Address:** UC Berkeley School of Information, 102 South Hall, Berkeley CA 94720-4600