

Selective Incentives and Generalized Information Exchange

COYE CHESHIRE

University of California at Berkeley

The goal of this research is to understand how generalized exchange systems emerge when information, as the object of exchange, produces a collective good. When individuals contribute information for a collective benefit, it can create a group-generalized exchange system that involves a social dilemma. I argue that two properties of information, replication and high jointness of supply, are crucial for understanding the nature of the social dilemma in these exchange systems. Combined with low-cost contributions, these special features of information can allow social psychological selective incentives to significantly encourage cooperation. Experiments were conducted to examine the independent effects of two social psychological selective incentives (social approval and observational cooperation) on sharing behavior in a generalized information exchange system. The results indicate that observing high levels of cooperative behavior is beneficial in the short run, but ultimately it only leads to moderately higher levels of cooperation than when individuals cannot observe cooperative behavior. On the other hand, when individuals receive either high or low levels of social approval, it has a very positive, significant impact on cooperative behavior. This research has implications for real-world generalized information exchange systems such as those found on the Internet. In addition, the theory and results in this study can also be extended to public goods that share the features of low-costs contributions, replication, and high jointness of supply.

In the current age of global economies, digital Internet connectivity, and high mobility, information is an essential part of everyday exchanges between individuals and organizations. From the most basic exchanges of news between colleagues or neighbors to the complex exchanges of strategic information between armed forces and governments, the information we give and receive helps us carry out our objectives. With so much information being created, collected, and distributed on a daily basis, we face a growing challenge to understand who shares information and what types of incentives encourage such cooperative behavior.

Of all the new information that is regularly being created, the majority of it (as much as 92%) is stored on magnetic media used in

computers and computer networks (Lyman and Varian 2003). Currently, one of the most prevalent systems for sharing and distributing this type of digital information is through Internet peer-to-peer (P2P) file-sharing networks. In peer-to-peer systems, individuals often upload (share) and download (retrieve) digital files that store music, movies, software, images, and other types of information. In one of the most popular P2P systems (KaZaA), users collectively share as much as 5,000 terabytes of information; yet at any given time, only about 9% of the users are actually sharing files (Lyman and Varian 2003). In another study of the popular peer-to-peer software system, Gnutella, researchers found that about 70% of users do not share files, and those that do share do not necessarily share the most desirable files (Adar and Huberman 2000).

This type of behavior is not just unique to peer-to-peer systems. The development of open-source software such as Linux, Apache, and Sendmail are examples of collective information goods (software) that are created through the contributions of many different individuals (see Raymond 2001). Despite the millions of open-source software users, the

*Address correspondence concerning this paper to Coye Cheshire, School of Information, 102 South Hall, UC-Berkeley 94720-4600; coye@ischool.berkeley.edu. I would like to thank Karen S. Cook, Cecilia Ridgeway, Noah Mark, Morris Zelditch, Walter Powell, Spencer Cahill and several anonymous reviewers for all of their advice and feedback on this research. This project was funded through a National Science Foundation Dissertation Improvement Grant (Award ID: 0402614).

actual number of contributors is a relatively small percentage of the target population. Just as with peer-to-peer exchange systems, a majority of users retrieve and collect the information goods while a minority of users provides them. Given this observed discrepancy between those who share and those who only retrieve information, how can these types of information exchange networks emerge and persist?

To address this question, it is first useful to consider the examples of open-source software and Internet peer-to-peer file sharing networks in the broader context of social exchange systems, public goods, and social dilemmas. The above examples can be viewed as a type of social exchange system called *generalized exchange*. Generalized exchange occurs when, “the reward that an actor receives. . . (is) usually not directly contingent on the resources provided by that actor” (Yamagishi and Cook 1993: 235). Since no individual is in a position to make another’s receipt conditional on her giving, it becomes possible to free ride (i.e., to receive goods or services without giving anything back in return). Thus, generalized exchange networks contain inherent social dilemmas (Yamagishi and Cook 1993).

One of the major types of generalized exchange deals with individuals who contribute resources toward the production of a collective good, which Ekeh (1974) calls “group-focused generalized exchange” and Yamagishi and Cook (1993) call “group-generalized” exchange. Each individual may provide some input to the collective good, and any returned value comes from that collective good. In this paper, I focus only on this type of group-generalized exchange.¹

The nature of the object of exchange (i.e., the good) plays a key role in social exchange and collective goods. For example, a collective good might be a physical good, such as a small-town bridge, which has implications for how many people can use it at one time, and what kinds of costs are associated with pro-

ducing it. The relative value and cost associated with providing a collective good affects whether individuals will contribute at all (Olson 1965). Similarly, the benefit that an individual expects to receive as others collect from the public good also changes the likelihood of contribution (Oliver and Marwell 1988; Marwell and Oliver 1993).

Previous studies of generalized exchange and collective goods primarily focus on physical goods (see also Foa and Foa 1974) for an early attempt to describe the nature of different goods in social exchange). However, the exchange of information has remained a largely uncultivated area of research in social psychology and in social exchange. Information can be much like any other good, since it can be transferred and it has value. However, information also differs from most physical goods in at least two key ways. First, information has the quality of *replication*, which specifically means that information can be transferred to an individual without the original owner losing her copy of the same information (Shah and Levine 2003). As a result, the contributor often does not lose the information when she shares it. This is true when a copy of the information is assumed to have the same content as the original. For example, a recipe can be transferred from one person to the next and each person benefits from the same information without the owner losing any value from the good herself. No matter how many people the owner gives the recipe to, she still retains her copy.²

The second key feature of information is that it can be consumed by many individuals without losing much value, and in some cases, without losing any value at all. Thus, information can have very high *jointness of supply*.³ Jointness of supply refers to the way that consumption of the public good affects the avail-

² It is important to note that I am not referring to value that comes from rarity. Some information, such as a secret recipe, is valuable precisely because it is unique and the content of the information is not freely available. That is, its value is *directly tied* to its rarity.

³ Shah and Levine (2003) also call digital information non-rival goods. Jointness of supply and rival/non-rival goods are both economic terms that describe the same concept.

¹ Another common type of generalized exchange that is not covered here is chain-generalized (Ekeh 1974), or network-generalized exchange (Yamagishi and Cook 1993).

ability of that good for other individuals. As a result, jointness of supply ranges on a continuum from zero (i.e., the cost of providing a good increases proportionally as it is consumed) to one (i.e., the value of the good remains unchanged no matter how many individuals consume it). Pure jointness of supply occurs when everyone receives the content value from a good, regardless of the number of individuals that consume it (Marwell and Oliver 1993). Variations on the parameter of jointness of supply have been extensively explored theoretically (i.e. Olson 1965; Heckathorn 1991, 1996).

Information that is transmitted over a computer network provides an ideal example of how high jointness of supply and replication can operate in information exchange. These 'digital information goods' (Kollock 1999b) may include software, photographs, art, music, speeches/lectures, television shows, books, magazines, or movies. As with the example of a recipe, the exchange of digital goods is accomplished through replication. When an individual shares a digital good, it is always a copy of the original. Furthermore, no matter how many people copy the digital good, it does not reduce others' ability to enjoy the same item.

In this paper, I propose a solution to the problem of the emergence of generalized information exchange systems through the use of social psychological processes as selective incentives. Social psychological processes have previously been identified as an important non-economic influence on behavior in collective-action situations (Klandermans 1984). However, social psychological selective incentives have been largely overlooked in research on generalized exchange and collective goods. I focus on two key social psychological selective incentives: social approval and observational cooperation. Social approval involves the ability to have one's contributions evaluated by others in the network and/or to evaluate the contributions of others. Observational cooperation concerns the ability to observe the total amount of sharing in the generalized information exchange system. I argue that the low cost of contributions, combined with the features of high jointness of

supply and replication, can allow small social psychological processes to have a relatively significant impact on cooperation in generalized information exchange. Although the scope of this study is limited to information exchange problems that produce a public good, the following theory and arguments can be extended to public goods that share the features of low-costs contributions, replication, and high jointness of supply.

A THEORY OF THE EMERGENCE OF GENERALIZED INFORMATION EXCHANGE

For generalized information exchange to emerge, individuals must overcome the temptation to receive without contributing and instead engage in sharing (cooperative) behavior. Once sharing begins, the overall value of the collective good can increase as more individuals contribute more goods (with high jointness of supply). In addition, individual information preferences are more likely to be met through diversity as the group size grows.

Even if the cost of contribution is extremely small, perhaps no more than the time it takes to make the contribution, it may appear that there is no rational incentive to contribute because a given individual's contributions will never benefit her. This is a crucial distinction between group-focused generalized information exchange and other common types of public goods problems in the real world and in much experimental work. In most public goods problems, individuals contribute to a public good and equally receive benefits from this collective good. For example, one of the most popular experimental representations of the public goods problem, the voluntary contribution mechanism, allows participants to choose whether to put money into a private account or into an equally distributed public account that earns interest (e.g., Isaac and Walker 1988a, 1988b; Dorsey 1992). Thus, individuals who invest in the public good can actually receive a portion of benefits from their own contributions. In generalized information exchange, however, all contributions to the public good are sunk costs because the contributor already has the information that she contributes.

Generalized information exchange can also be described as a social dilemma: a group structure that involves individually dominating strategies that are superior to all other strategies, no matter what other individuals do; all of the strategies converge on a deficient outcome that is less preferred by everyone (Dawes 1991; Yamagishi 1995). In the case of generalized information exchange, the individually dominating strategy is *not* to contribute to the pool of information, even though this leads to the deficient outcome of not producing a collective pool of information. The supply of the public resource (information pool) does not decrease regardless of the number of individuals that benefit from it, but there is still a cost associated with contributing to the public resource. This cost is small—and certainly much smaller than the actual value of the information goods. For example, it is fair to assume that there is an opportunity cost of time that could have been invested elsewhere (Heckathorn 1996).

Social Psychological Incentives in Generalized Information Exchange

Since individuals are better off when they receive information without giving it in return, it is necessary for individuals to make a seemingly altruistic contribution to a collective good for generalized information exchange to emerge. Empirical evidence suggests that altruistic behavior is a natural aspect of social interaction. For example, individuals donate blood and organs at some personal cost with no direct benefits (i.e., Healy 2001; Piliavin, Allyn, and Charng 1990; Simmons 1992). Altruistic behavior also exists in collective action problems where there appears to be no rational explanation why individuals would otherwise want to contribute. Coleman (1988) cites examples of rational zealotry, such as terrorists who engage in behavior on behalf of what they believe are public goods or hunger strikers who fast—in some cases until death.

Altruism and rational zealotry might explain why some individuals contribute to public goods even when doing so does not seem rational. However, if there is an added benefit that only occurs when an individual

contributes, then contributing (i.e., cooperation) becomes slightly more attractive regardless of the decisions of others. Selective incentives for contributors have been very successful as an economic solution to free-riding, because they offer additional motivations that make cooperation rational even when the initial conditions of the social dilemma make cooperation irrational (e.g., Olson 1965; Oliver 1980; Heckathorn 1996).

In generalized information exchange, removing some degree of anonymity between individuals can help to make social psychological processes act as selective incentives. When exchanges (or collective contributions) are no longer purely anonymous, individuals can have knowledge about other exchanges, as well as their own. For example, recent findings about reputation systems in other types of exchange (e.g., Takahashi 2000; Kollock 1999a) suggest that when individuals have some knowledge about others' contributions, it affects the way that the individuals consider their own contributions.

Social Approval. Individuals may gain some intrinsic satisfaction from the popularity of their own contributions in the form of psychological efficacy, leading to an increased desire to share more in the future. In this case, I define popularity as the frequency that a unit of an individual's information is considered desirable by other contributors. Specifically, as more individuals acknowledge recipe "x" as desirable, the popularity of recipe "x" increases.

In addition to receiving social approval, individuals may participate in giving social approval by rating the popularity of other's contributions. Thus, it is possible for giving and receiving social approval to affect behavior. Individuals may cooperate (or share) because they care about their popularity rating (receive social approval) and/or because they want to vote on the contributions of others (give social approval). In this study, social approval is a combination of both processes.

Assumption 1: *The popularity of one's information bestows some small, but significant, social psychological reward on the individual.*

Assumption 2: *The ability to rate the information of others bestows some small, but significant, social psychological reward on the individual.*

Assumption 3: *Individuals are forward-looking, anticipating rewards that could be bestowed after subsequent interactions.*

If individuals are forward-looking, and anticipate the actions of others, then assumptions 1, 2, and 3 lead to the prediction that social approval processes positively affect the desire to contribute.

Forward-looking Social Approval Reward

Hypothesis: *In generalized information exchange, when individuals are aware that they will be able to observe the popularity of their own contributions and rate the contributions of others, they will cooperate more than when they are not able to do so.*

Because individuals care about what others think about their information, high levels of positive feedback should have a greater positive impact on cooperation than low levels of positive feedback.

Social Approval Differential Hypothesis: *In generalized information exchange, high social approval will have a greater positive effect on the cooperation rate compared to low social approval.*

The social approval in either high or low conditions should help to keep cooperation rates from decreasing over time during the emergence of generalized information exchange. High social approval rewards individuals through positive feedback so that they continue to contribute over time. Individuals who receive consistently low social approval might continue to contribute goods in the hopes of receiving higher social approval over time. In addition, individuals may look forward to the ability to give social approval through rating the contributions of others.

Social Approval Rate Hypothesis: *In generalized information exchange, consistently high or low social approval will not have a negative effect on the cooperation rate over time.*

A scope condition for any situation where the social approval of one's contribution is an imperative is that an individual's information

must be important (i.e., have personal value) to the original owner. When this condition is met, the popularity of one's information is likely to be more rewarding because it is a reflection of the value(s) placed on it by the owner.

Observational Cooperation. Previous research has shown that individuals will conditionally cooperate based on what they believe others are doing in a public goods situation (e.g., Fischbacher, Gächter, and Fehr 2001; Kurzban et al. 2001; Sell 1997; Keser and van Winden 2000). From a game-theoretical perspective, there is no strategic advantage to matching one's cooperation level to the rest of the group when others are already cooperating at a relatively high level (Sell 1997).

Keser and van Winden make a distinction between what they call "*future-oriented behavior* and simple *reactive behavior*" (2000: 32, emphasis in original) in conditional cooperation. Future-oriented behavior deals with the tendency for individuals to modify their behavior based on what they believe will happen in the future. This aspect of behavior has a similar logic to the game-theoretical approach to conditional cooperation. Behavior is explained as a set of rational strategies as individuals look forward to future interactions. Future-oriented behavior or game-theoretical reasoning can also be used to explain end-game effects and differential contributions at the beginning of some public goods games (Keser and van Winden 2000).

The other aspect of behavior that Keser and van Winden (2000) believe affects conditional cooperation is reactive behavior, wherein individuals tend to orient themselves towards the average behavior of other group members. When individuals can observe the degree to which other participants are cooperating, they can stimulate a normative response to reciprocate by cooperating as well. Thus, reactive behavior is closely tied to the principle of reciprocity (Keser and van Winden 2000). Furthermore, when individuals are able to make small contributions to a public good and then observe if others do the same, it allows the individual to signal their commitment to building the public good without tak-

ing too much risk at once (Kurzban et al. 2001: 1664).

The reactive (or reciprocal) perspective also has implications for how an individual will see herself in relation to others (See Yamagishi 1995 for a discussion of social orientation and motivation for contributing in social dilemma situations). When individuals observe cooperative behavior, it makes the decision to cooperate less impersonal. Furthermore, observing cooperative behavior can induce some level of obligation (Gouldner 1960; Granovetter and Soong 1983). Although an *obligation* can have a negative connotation in some circumstances, it can also make an otherwise uncongenial situation more personal. Individuals can experience at least a minimal amount of satisfaction from being a cooperator because they feel like they are part of the group (Keser and van Winden 2000).

I argue that individuals are likely to imitate the perceived aggregate actions of others, either through a sense of normative compliance or a sense of obligation to participate. Of course, this type of conditional cooperation has already been established through many studies, such as those described above. What is most important here is that when the costs for contributing information goods are small, observing cooperative behavior can have a relatively significant impact on overall cooperation. If individuals choose to cooperate based on what others are doing, then there must be a mechanism in place to allow for small (or low-cost) contributions so that individuals can observe whether or not others reciprocate (Kurzban et al. 2001).

Assumption 4: *Observing cooperative behavior induces a sense of normative compliance and/or obligation in individuals.*

Assumption 5: *Normative compliance and/or obligation create a small but significant desire to cooperate.*

Together, assumptions 4 and 5 imply that observing cooperative behavior in others stimulates normative compliance or obligation, thereby increasing an individual's rate of cooperation.

Observational Cooperation Hypothesis: *In generalized information exchange, when individ-*

uals are able to observe the amount of cooperative behavior in the exchange system, they will cooperate more than when they are not able to do so.

When individuals observe high amounts of cooperation, they know they are receiving benefits and that there is very little free-riding taking place. On the other hand, when individuals observe low amounts of cooperation, they know most individuals are free-riding on the contributions of a minority of participants (Sell 1997). For these reasons, it follows that individuals should cooperate more often when observational cooperation is high rather than low.

Observational Cooperation Differential Hypothesis: *In generalized information exchange, high observational cooperation will have a greater positive effect on the cooperation rate compared to low observational cooperation.*

Despite the positive effect that observational cooperation might initially have on an individual's own likelihood of cooperation, it is rational for individuals to decrease their own amount of cooperation over time in both the high and low observational cooperation situations. In a situation in which individuals always observe high amounts of cooperation, earnings are consistently being guaranteed. In a situation in which individuals always observe low amounts of cooperation, it becomes futile to keep contributing when the overall amount of cooperation is never changing for the better. A decrease over time is consistent with game-theoretical predictions about conditional cooperation, since individuals are expected to eventually modify their behavior over time to maximize their earnings. It is also consistent with previous theory and research showing that, over time, individuals tend to decrease their contributions unless additional motivations are introduced (e.g., Davis and Holt 1993).

Observational Cooperation Rate Hypothesis: *In generalized information exchange, consistently high or low observational cooperation has a negative effect on the cooperation rate over time.*

THE EXPERIMENTS

The experiments in this study use the concept of digital information goods (e.g., Kollock 1999b) as the object of exchange in a generalized exchange system. Digital information goods are items that can be converted to digital files and are easily exchanged across computer systems such as the Internet. As previously shown, digital information goods can have the features of replication and high jointness of supply when they are traded in real-world Internet exchanges. The experimental situation in this study is designed to approximate a simplified version of a real-world system of digital information exchange, like those found on the Internet.

There are five conditions in this study: four experimental conditions and one control condition. The control condition adheres only to the general procedures (described in detail below), and does not have any other experimental manipulations. The first two experimental conditions are designed to test the social approval process (high and low social approval). The two remaining experimental conditions are designed to test observational cooperation effects (high and low observational cooperation).

Participants

Potential participants were recruited by the Center for Social Research, located in the Department of Sociology at Stanford University. Recruitment slips were distributed in freshman-only courses, and four hundred completed slips were used to recruit potential participants. Emails were sent to each potential subject directing them to a sign-up webpage. The recruitment emails and webpage stated that the average payment for the current experiment was between "\$10–20." A total of 179 subjects signed up and participated. The first twenty subjects from this pool were used in a pre-test to refine the instructions and the procedures. Of the 159 non-pretest subjects, twenty-four subjects were eliminated due to high participant suspicion, as determined by the post-experimental questionnaire and debriefing. Thus, the suspicion rate was 15% (valid N = 135). There were between 20 and 29

subjects in each of the five experimental conditions,⁴ including 73 females and 62 males (54% and 46% of the valid sample, respectively). The mean age of the valid sample was 18.2 years.

Procedure Summary

The experiment consisted of ten rounds in which subjects made decisions about whether or not to contribute digital information goods to a collective information pool. The subjects did not know in advance how many rounds there would be during the experiment. The computer software for the experiment ran in an Internet browser and subjects were told that they would be participating with many other real people at several universities. However, the computer program was actually an elaborate script that used timed delays between display screens to convince subjects they were interacting with other people.

The Experimental Task (All Conditions). When the researcher started the experiment, the subjects began by reading several pages of instructions about the experimental task. The participants then read and filled out an information goods creation form, which required the subject to list twenty of their favorite songs, movies, or books. Subjects were told that they would use these items in a series of Internet exchanges. Once they typed this list into the computer, each information good received a value of 10 points. Therefore, the participant started with twenty information items worth a total of 200 points. These initial points acted as the "bank account" to which all gains and losses would be added or removed during the experiment. The subjects were told that they would be given a choice on each round of the experiment to select one of their own information goods to contribute to the public information pool or not to contribute at all. Subjects were told that their information items were worth 10 points to everyone else in the exchange system. Similarly, subjects

⁴ To obtain the needed statistical power for the analyses, at least 20 subjects were needed in each experimental condition. After this goal was met, subjects were randomly assigned to the four non-control conditions until the subject pool was exhausted.

learned that they would receive 10 points for every contribution made by others. That is, the benefits from individual contributions were distributed to all participants except the original contributor. When an individual chose to share one of her items, however, a small transfer cost (5 points) was deducted from her virtual account. In sum, the only way to earn more points was if other individuals contribute, but all shared information goods cost each contributor a small number of points. Subjects were told that their total payment would be determined by how many points they earned during the experiment.⁵

All subjects in every condition were told that they would receive a complete list of the information goods that were contributed at the end of the experiment (regardless of their behavior), and they would not be told who contributed them. This small step helped to insure that subjects cared about the information goods in all of the experimental conditions (instead of thinking of the digital goods as only “points”).

Based on the cost/benefit structure, the earnings and rewards that could be earned would change with relation to the size of the exchange system. For this reason, subjects were not told the exact size of the system. Like many real-world systems of Internet exchange, it was not possible for the subject to determine the exact number of participants in the exchange system.

At the conclusion of the experiment, subjects were given a post-test questionnaire that asked them to describe why they chose to share, when and if they did so, and why they chose not to share (when applicable). Finally, the post-questionnaire asked the participant about task comprehension and believability. After the subjects completed the post-test questionnaire, the experimenter gave the subject a debriefing form that explained the true nature of the study.

Social Approval Conditions (High and Low Social Approval). The social approval

conditions included additional procedures in conjunction with the general procedures listed above. The instructions stated that a list of all contributions by others would be displayed to all participants who contributed on the previous round on the computer screen. Subjects who made contributions in the previous round would then get to vote for up to five items from the list of information goods as the “best items”. The list included all of the current contributions, except the contributor’s own contribution (so she could not vote for her own item). In reality, this list came from a pool of information items that were pre-programmed into the exchange script. This list was created from actual information goods that were shared by participants during the pre-test experiments. Subjects were also told that the collective votes created a popularity rating. After the voting ended, the participant saw a bar chart on the screen that prominently displayed the popularity rating of her contribution from the previous round. If the subject did not contribute on the previous round, then no popularity score was displayed and the message, “no current contribution”, appeared instead.

It is important to note that the popularity rating is tied to the ability to vote. Although it is not possible to tease out the independent effects of ‘the desire to vote’ and ‘the desire to view one’s own popularity’ within the experiments, the post-questionnaire responses suggest that both processes do affect behavior. Specifically, post-test questionnaire responses show that many subjects self-report a desire to vote or a desire to view their own popularity rating. However, there was no overlap between reporting ‘the desire to vote’ and ‘the desire to view one’s own popularity’. This indicates that both processes are at work, yet individuals do not self-report both at the same time. This issue is addressed further in the results.

Subjects were told that there was no monetary reward for having the most popular goods, but only participants who contributed would get to vote on the most popular good(s) at the end of each round. Subjects who did not contribute would not get to vote, and were told to “wait” until other subjects voted. These

⁵ Since the exchange system is actually a controlled script, all subjects are paid \$15 at the end of the experiment. Subjects are fully debriefed on the nature of the experiment at its conclusion.

wait-times were randomly generated as described before.

Popularity ratings were manipulated by having a bar chart on the subject's computer screen that ranged from 1–100%. In the high social approval condition, an individual always saw her contribution fall between 75–95% on the bar chart, and this range was labeled as 'high popularity' on the screen. In the low condition, an individual always saw her contribution fall between 5–25% on the bar chart, and this range was labeled as 'low popularity' on the screen. In both the high and low social approval conditions, the exact percentage was randomly determined by the computer within the specified range for that condition.

Observational Cooperation Conditions (High and Low Observational Cooperation). In addition to the general procedures, the subjects in the observational cooperation conditions were told that they would be able to view information regarding the amount of sharing on the previous round. Regardless of whether a participant decided to share or not, the participant was able to actively see how much contributive sharing was occurring in the network. Only information about the previous round was presented on each subsequent round throughout the experiment.

Observational cooperation was manipulated in almost exactly the same way as social approval. A bar chart that ranged from 1–100% was used to display the amount of observational cooperation in each round. In the high condition, an individual always saw that 75–95% of the subjects cooperated on the previous round (this amount was labeled 'high amount of sharing' on the bar chart). In the low condition, an individual always saw that 5–25% of the subjects cooperated on the previous round (this amount was labeled 'low amount of sharing' on the bar chart). In both the high and low observational cooperation conditions, the exact percentage was randomly determined by the computer within the specified range for that condition.

Dependent Variables. There are ten rounds in each of the five conditions in the experiment. The dependent variable for each round is the binary decision to share one 'information good' with the group or not (i.e., cooperation

versus defection). For the following analyses, the dependent variable is examined in two ways. First, an *average cooperation rate* is computed from the average choices for each subject across all ten rounds. Thus, the cooperation rate has decimal values that range from zero to one. Second, in analyses that account for the effect of decisions over time, five trial blocks were created, with two rounds per trial block.

RESULTS

Figure 1 shows the average cooperation rates for each experimental condition and how they compare to each other. Several *t*-tests were used to explore the various differences between average cooperation rates between conditions as well as between paired samples (beginning of experiment, end of experiment) within conditions. The first and fifth trial block measures the cooperation rate at the beginning and end of the experiment, respectively.

The average cooperation rates for high social approval (.87) and low social approval (.81) were significantly higher than the control condition (.55). As Table 1 shows, high and low social approval had a significant positive effect on the average cooperation rate, compared to the control condition. This result provides initial support for the Social Approval Reward Hypothesis.

Table 2 displays the overall results of the within-subject and between-subject effects of the repeated-measures ANOVA. The tests of the within-subject effects were significant for the main effect of trial block, $F(3.5, 259) = 4.45, p < .01$, and the trial block by experiment condition interaction, $F(6.9, 259) = 2.29, p < .05$. Thus, there were some significant differences across trial blocks and between conditions, and the changes over time varied between one or more conditions. Finally, the between-subjects effect of experiment condition was significant, $F(2, 75) = 12.73, p < .01$.

Table 3 displays the post-hoc analysis for the difference scores between the cooperation rates in each experimental condition. The

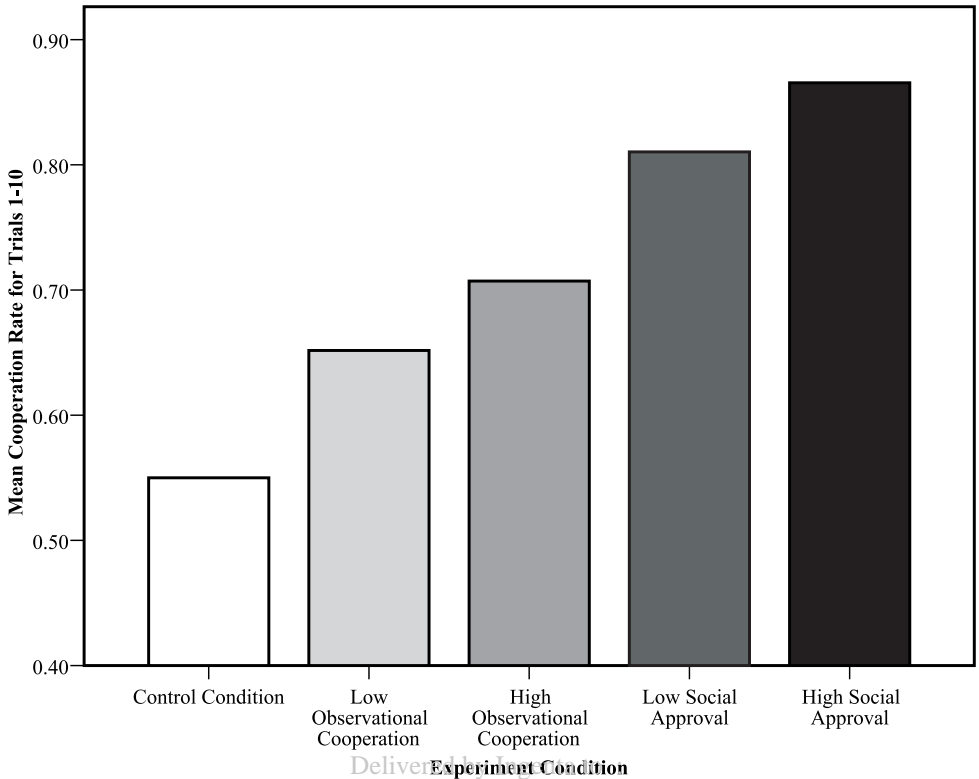


Figure 1. Average Cooperation Rates for Experimental Conditions

Table 1. Mean Comparisons for Cumulative and Trial Block Cooperation Rates in Each Experimental Condition.

Experiment Condition	Means			Mean Difference: First Trial Block–Last Trial Block§
	Total‡	First Trial Block†	Last Trial Block†	
Control	.55 (.27)	.73 (.30)	.45 (.39)	.28* (.47)
Low Observational Cooperation	.65 (.19)	.71** (.31)	.59 (.40)	.12 (.42)
High Observational Cooperation	.71* (.27)	.89** (.21)	.63 (.44)	.27** (.40)
Low Social Approval	.81** (.19)	.83 (.28)	.78 (.31)	.05 (.39)
High Social Approval	.87** (.22)	.90 (.25)	.84 (.33)	.05 (.34)

N = 135

(Standard Deviations in Parentheses)

* $p < .05$; ** $p < .01$ (two-tailed tests)‡ Analyses in this column are Independent-Sample t -tests between each experimental condition and the control condition.† Analyses in these columns are independent-Sample t -tests between the High and Low situations within each experimental condition (e.g., high observational cooperation versus low observational cooperation; high social approval versus low social approval).§ Analyses in this column are paired-sample t -test between the first trial block and the last trial block in each experimental condition.

Table 2. Repeated Measures ANOVA for Cooperation Rates across Trial Blocks and between Social Approval Conditions.

Tests of Within-Subjects Effects†				
Source	Type III Sum of Squares	df	Mean Square	F
Trial Block	0.91	3.50	0.26	4.45**
Trial Block × Condition	0.94	6.90	0.14	2.29*
Error (Trial Block)	15.43	259	0.06	
Tests of Between-Subjects Effects				
Intercept	208.23	1.0	208.23	830.18**
Experimental Condition	6.39	2.0	3.19	12.73**
Error	13.81	75	.25	

N = 135

* $p < .05$, ** $p < .01$ (Pillai's Trace Multivariate Test)† Sphericity cannot be assumed ($p < .01$). As a result, the Greenhouse-Geisser adjustment is applied to the degrees of freedom.

mean differences between high social approval compared to the control ($-.32$, $p < .01$) and low social approval compared to the control ($-.26$, $p < .01$) were both statistically significant. Thus, the Forward-looking Social Approval Reward Hypothesis received strong support across the various statistical tests.

Using a t -test, the difference between the cumulative cooperation rate in the high social approval condition (.87) and the low social approval condition (.81) was not statistically significant. In addition, t -tests were used to examine the difference between conditions for the first trial block only, and the last trial block only. As Table 1 shows, in each case the differ-

ence between the two social approval conditions was not statistically significant. The Bonferroni post-hoc comparison tests in Table 3 also show that the difference is not significant. Although the difference was in the hypothesized direction, the Social Approval Differential Hypothesis was not supported. However, the lack of support for the Social Approval Differential Hypothesis may not speak so much to the weakness of high social approval (which has a very high average cooperation rate across all trial blocks), but rather it demonstrates the strength of low social approval. I explore this intriguing result further in the discussion.

Table 3. Post-Hoc Analysis of Mean Difference Scores for Cumulative Cooperation Rates between All Experiment Conditions (Bonferroni Comparisons)

(I) Experiment Condition	(J) Experiment Condition	Cooperation Rate Mean Difference (I-J)
		Total
Control Condition	Low Observational Cooperation	-.10 (.07)
	High Observational Cooperation	-.16† (.07)
	Low Social Approval	-.26** (.07)
	High Social Approval	-.32** (.07)
Low Observational Cooperation	High Observational Cooperation	-.06 (.06)
Low Social Approval	High Social Approval	-.06

N = 135

(Standard Errors in Parentheses)

† $p < .10$; ** $p < .01$ (two-tailed tests)




Figure 2. Estimated Marginal Means of Cooperation Rate over Trial Blocks for Social Approval and Control Conditions

As Figure 2 shows, the cooperation rates for both social approval conditions remained very high across all trial blocks. Although there were some small fluctuations across trial blocks (more so for low than high social approval), the cooperation rates for both conditions were relatively stable. In fact, the mean difference for both conditions between the first and last trial block was only .05 (a non-statistically significant difference). Since the cooperation rates did not decrease for either of the two social approval conditions, the Social Approval Rate Hypothesis was supported.

Observational Cooperation

The average cooperation rate in the low observational cooperation condition (.65) was not significantly larger than the average cooperation rate in the control condition (.55). However, the average cooperation rate in the high observational cooperation condition (.71) was significantly higher than that of the control condition ($t = -1.9, p = .05$). As Figure 3 shows, the cooperation rate in the low observational cooperation condition started at about the same point as the control condition, but did

not decrease as drastically as the control. On the other hand, the high observational cooperation condition started much higher than the low observational cooperation condition through the first few trial blocks, yet the cooperation rate fell to the same level as that of the low observational cooperation condition by trial block 3 and remained at about the same level from that point forward.

The repeated-measures ANOVA of the observational cooperation conditions and the control condition display similar results, with some important differences. As Table 4 shows, the within-subjects effect of trial block was significant, $F(4, 296) = 9.81, p < .01$. However, the trial block by experiment condition interaction was not significant using Pillai's Trace multivariate test. However, Roy's Largest Root multivariate test did indicate a small, significant effect for the trial block by experiment condition interaction ($p = .05$).⁶

⁶ Roy's Largest Root is a less robust statistical test than Pillai's Trace (since it compares the largest difference between the conditions), so this result should be considered in light of this.

Table 4. Repeated Measures ANOVA for Cooperation Rates across Trial Blocks and between Observational Cooperation Conditions.

Tests of Within-Subjects Effects§				
Source	Type III Sum of Squares	df	Mean Square	F
Trial Block	3.03	4.0	0.78	9.81**
Trial Block × Condition	0.70	8.0	0.09	1.13‡
Error(Trial Block)	22.85	296	0.08	
Tests of Between-Subjects Effects				
Intercept	151.58	1.0	151.58	500.63**
Experimental Condition	1.45	2.0	0.73	2.39†
Error	22.41	74	0.3	

N = 135

† $p < .10$, * $p < .05$, ** $p < .01$ (Pillai's Trace Multivariate Test)

‡ $p = .05$ for Roy's Largest Root, a less robust multivariate test based on the largest eigenvalue. Pillai's Trace test is not significant ($p = .26$).

§ Test of Sphericity was rejected ($p = .22$). No correction was applied to the degrees of freedom.

The between-subjects effect of experiment condition was borderline significant, $F(2, 74) = 2.39$, $p < .10$. The post-hoc tests in Table 3 show that the mean cooperation rate in the high observation cooperation condition differed from the control condition, though the difference was only borderline significant (mean difference = $-.16$, $p < .10$). There was no significant difference between the low observational cooperation condition and the control condition (mean difference = $-.10$, $p = n.s.$).

Based on these results, the Observational Cooperation Hypothesis received only partial support. When individuals believed that there was a high amount of cooperation, they cooperated at a higher rate than when they were unable to observe cooperative behavior. However, when individuals believed that there was a low amount of cooperation, they did not cooperate more than when they were unable to observe cooperative behavior.

The average mean cooperation rate in the high observational cooperation condition (.71) and the low observational cooperation condition (.65) was not significantly different. However, the mean comparisons in Table 1 show that the difference between the high (.89) and low (.71) observational cooperation conditions in the *first trial block* was significantly different ($t = -2.6$, $p < .01$). By the last trial block, the difference was not statistically sig-

nificant. Thus, the Observational Cooperation Differential Hypothesis was partially supported, but the caveat is that the significant difference between high and low observational cooperation only exists during the initial exchanges.

Figure 3 shows the means plots across the five trial blocks for the observational cooperation conditions and the control condition. The cooperation rate in the low observational cooperation condition did not significantly decrease (mean difference .12, $t = 1.6$, $p = .12$). However, the high observational condition did significantly decrease over time (difference .27, $t = 3.6$, $p < .01$). As previously shown, the repeated-measures ANOVA results indicate that the interaction between trial block and experimental condition was only partially supported when using a less robust multivariate test (Roy's Largest Root). While there was a significant cooperation rate decrease over time in the high observational cooperation condition and the control condition (as shown in Table 1), it is not clear that these changes were significantly different between the three conditions. Still, since the high observational cooperation condition did show a significant decrease in the t -test comparisons while the low observational cooperation condition did not (and there was not a very clear, significant interaction between the conditions and trial block), the results provide




Figure 3. Estimated Marginal Means of Cooperation Rate over Trial Blocks for Observational Cooperation and Control Conditions

limited support for the Observational Cooperation Rate Hypothesis. (Since the percentage of those sharing was always low in the

Post-experimental Questionnaire Results: Subject Rationale for Sharing or Not Sharing

At the conclusion of the experiment, each subject filled out a short post-experimental questionnaire to measure the subjects' rationale for sharing or not during the experiment. A qualitative software text analysis program (SPSS Text Analysis for Surveys) was used to analyze the answers and pool common terms and categories from the responses. In addition, manual coding and error-checking of the responses assured they were appropriately categorized.

Many subjects chose to share because they thought that it might help them earn more money in the long run (between 40–50% in all conditions); however for each condition, between 19–59% of subjects also claimed that making more money was their motivation for not sharing. Among those in the low observation cooperation condition, 31% said that they shared in order to build cooperation in the network (compared to 7% in the high observa-

tional cooperation condition). Since the percentage of those sharing was always low in the low observational cooperation condition, it follows that subjects must have felt as if their efforts had no effect on others. In fact, 21% of the subjects in the low observational cooperation condition said that they chose not to share specifically because of the percentage value displayed on the bar chart. This helps to explain the low amount of overall sharing in this condition. In the high observational cooperation condition, the situation is very different. Thirty percent of the subjects in the high observational cooperation condition said that they cared about the percentage value of the bar chart. Since the percentage value was always high in this condition, it is reasonable to assume that many individuals may have stopped sharing because they simply knew that they were already going to make a lot of points. Indeed, 41% of subjects in the high observational cooperation condition said that they did not share because of the percentage value on the bar chart.

In the social approval conditions, 19% of the participants stated that one of the reasons they contributed was their ability to view their

popularity rating. On the other hand, 15% of these subjects stated that they contributed because they wanted to vote. As previously mentioned, these two groups do not overlap. Individuals who self-report that they have a desire to vote do not report a desire to view their popularity ratings, and vice-versa. However, this finding becomes even more fascinating when the two social approval conditions are compared to each other. Of the subjects in the low social approval condition, 38% indicated that they chose to share because they cared about the popularity of their contributions. What makes this so intriguing is that none of the subjects in the high social approval condition self-reported that they cared about their popularity rating. Thus, not only is self-reporting 'the desire to vote' mutually exclusive from 'the desire to view one's popularity rating', it also appears that only individuals in the low social approval condition self-report a desire to view their own popularity. These findings highlight the potential for future research into the differential effects of these two aspects of social approval (i.e. desire to vote or comment on others and the desire to receive feedback from others).

Given that low social approval leads to a very high average cooperation rate (.81), it appears that many subjects continued to share in this condition because they wanted to raise their popularity rating. If individuals were trying to increase their popularity rating, it makes sense that this is one of the most salient issues self-reported on the post questionnaire. However, the subjects in the high social approval condition may not have indicated that the popularity rating was important to them because they became used to receiving consistently high ratings. Once high popularity ratings became the norm for these individuals, the actual 'popularity rating' may not have been as salient when they were asked to self-report their motivations for contributing. These individuals tended to give other reasons for why they shared so much. Fifteen percent of the subjects in this condition stated that they shared for altruistic reasons, while 11% said that they shared because of a sense of fairness, and 11% said that they shared because they liked to vote on other contributions. Given the

overwhelmingly high cooperation rate in the high social approval condition (.87), it appears that consistently high social approval did influence continued cooperation—even if subjects did not realize this fact (or want to admit it). Indeed, just because a social psychological incentive is working does not mean that individuals are cognizant of it.

DISCUSSION AND CONCLUSION

This study demonstrates that social psychological processes can act as selective incentives that help to encourage cooperative behavior in generalized information exchange. The theory and results are particularly relevant to generalized information exchange systems such as those found on the Internet, as well as other collective-action problems in which the costs of contribution are low, individual contributions have inherent non-economic value to the contributor, and the public good has a relatively high degree of jointness of supply.

Selective incentives have long been considered a potential solution to the social dilemma in collective-action problems (Yamagishi 1995). Olson (1965) realized that selective incentives can change the situation in a collective-action problem so that individuals will have additional interests beyond the original costs and benefits when deciding whether to contribute or not to the public good. As Oliver states, "selective incentives can turn a collective-action situation in which cooperation is *irrational* into one in which collective action is *rational*" (Oliver 1980: 1359, emphasis in original). The reasoning is that a rational, self-interested individual may be initially inclined to free-ride in a public goods problem, but the addition of selective incentives can change the value functions and relative costs in favor of cooperation.

At first glance, it might seem as if selective incentives are the obvious solution to any collective-action problem, since they can potentially lead to cooperation among rational, economically self-interested actors. However, the production of selective incentives is itself a collective-action problem (Oliver 1980: 1361). For example, public radio fundraising efforts often use gifts and public recognition (such as

having one's name displayed in the newspaper) as selective incentives. The selective incentives have their own costs associated with them, so a new problem emerges: who will provide the selective incentives? In such cases, the production of selective incentives may constitute a second-order social dilemma (Yamagishi 1995).

Some social psychological selective incentives can potentially reduce or eliminate this second-order social dilemma because they are internal to the contributor. If mechanisms are in place to encourage individuals to draw on their own intrinsic feelings and motivations, then these incentives do not necessarily have to depend on additional maintenance costs. Indeed, the results of this study illustrate how the social psychological processes of observational cooperation and (especially) social approval can, in some cases, dramatically encourage cooperation when the means to produce them are in place. Social psychological selective incentives are easily formed because, like information, they have high jointness of supply (they do not necessarily cost more to produce no matter how many individuals contribute to the collective good). For example, everyone can potentially feel solidarity or experience the positive effects of social approval from their peers if they contribute to the public good.

Klandermans (1984) indicates that social psychological processes in collective action may come from an individual's expectations about the behavior of others. These may include expectations about the number of participants, expectations about one's own contribution to the probability of success, and expectations about the probability of success if many people participate. Using Klandermans' three expectations, it is quite possible that the subjects in the high observational cooperation condition believed: (1) there were many participants, (2) their own contribution would not make a big difference for the group, and (3) the probability of success was guaranteed to be high (because of the graph that displayed the percentage of others who were cooperating). To put it differently, "If an individual takes other people's contributions as given, she will contribute *less* as other people

contribute *more* (Sugden 1984: 773, emphasis in original). It seems that the rational tendency to share less when the contributions of others are guaranteed mostly offsets the social psychological benefits of solidarity and positive feelings. Still, there is evidence that individuals may contribute more to a public good when they see other people contributing (e.g., Bryan and Test 1967). For this reason, one implication from this study is that observational cooperation might be much stronger when the level of cooperation is average or largely fluctuating (since it is not guaranteed to be high or low).

The other social psychological selective incentive in this study is the social approval that comes from giving and receiving popularity votes. In many collective-action problems, social approval can act as a selective incentive only when a relatively small number of individuals can earn it by cooperating (Oliver 1980). This is because many forms of social approval may require an additional contribution cost, such as the cost associated with calling a friend to encourage participation in a collective-action effort. In this study, social approval comes from individual popularity ratings derived from a voting procedure that is part of the process of sharing information goods. Indeed, subjects report that they actually like to vote on the contributions of others. One implication here is that individuals may view the process of giving social approval as a kind of reward in and of itself, so long as the costs (e.g., time) are very small.

Hollander (1990) uses an economic mathematical model to demonstrate how the expectation of social approval can motivate cooperative behavior. In this model, the production of a collective good depends on the size of the approval incentive. Hollander finds that in large settings in particular, "social approval alone may be sufficient to support significant cooperation" (1990: 1166). The results of this experimental study support this finding. In this study, social approval is not diminished by the number of individuals who receive it. Again, this is an important distinction because many arguments and theories about collective action assume that social approval will inevitably be limited to only a small number of contributors.

It is possible that in a real-world system of generalized information exchange such as an Internet peer-to-peer file-sharing system, every contribution might not receive many votes—leading to low popularity ratings (i.e., low social approval). However, the results show that even low social approval induces cooperation in the early stages of generalized information exchange. This is perhaps one of the most unexpected and exciting findings in this study. In the formative stages of generalized information exchange, high or low social approval can be a powerful motivator for cooperation. This is intriguing because this type of social approval is completely restricted to the individual—no one else knows anyone else's popularity rating. Again, the kind of social approval used in this study is very different than public reputations which may motivate behavior for different reasons. In other types of social approval, there may be a social desirability component (like a public reputation). However, the popularity ratings that produce feelings of social approval in this study are only known to the individual who makes a contribution. It is encouraging to know that the so-called “pat on the back” or simple “thank you” may have more than just a little impact on behavior in some cases.

The results from this study demonstrate that simply observing high amounts of cooperative behavior has a positive effect on the cooperation rates of participants in generalized information exchange—especially in the initial stages. On the other hand, observing low amounts of cooperative behavior does not necessarily lead to significantly lower cooperation rates than the control condition (where individuals have no information about how much sharing/cooperation is taking place). Overall, this implies that the ability to observe the amount of cooperation may be useful for fostering cooperation in real-world peer-to-peer systems, though the effects are likely short-lived. Specifically, normative compliance may not have a particularly strong impact on cooperation, but it does seem clear that individuals cooperate more when they have some information about the rest of the network than when they have no information at all.

High and low social approval has a strong effect on the cooperation rate in generalized information exchange. The implications of these results are very clear for peer-to-peer systems: participants share significantly more digital goods, despite the small costs associated with a contribution, when they are able to give and get feedback on contributions. In the formative stages of a generalized information exchange system, even low social approval has a positive effect on cooperation. This implies that individuals really care about their own social approval. A peer-to-peer network that allows individuals to give some type of feedback to the participants creates a system where the participants feel like their contributions (and perhaps by association, their own selves) are being evaluated by a community of peers. Since peer-to-peer networks can (and often do) have anonymous participants, participation and sharing in these systems can seem isolating and inconsequential to the users. In some cases, individuals in anonymous Internet exchange may choose not to share (or defect) because they perceive the uncertainty associated with anonymity as a sign of greater risk (Cheshire and Cook 2004). The use of popularity ratings or other social approval measures can reduce this perception of anonymity and, as a result, encourage cooperation.

Understanding how generalized exchange can emerge and persist is an important problem for several disciplines, including sociology, economics, anthropology, communications and business. As much as we rely on information in our daily lives, the theoretical and practical problems surrounding information and information exchange continue to grow. Understanding how different types of information exchange systems emerge is not just an intellectual concern; it is a central problem for our increasingly interconnected societies.

REFERENCES

- Adar, Eytan and Bernardo A. Huberman. 2000. “Free Riding on Gnutella.” *First Monday* 5:10 (October 2000), URL: http://firstmonday.org/issues/issue5_10/adar/index.html
- Bryan, J. H., & Test, M. A. (1967). “Models and helping: Naturalistic studies in aiding behavior.”

- Journal of Personality and Social Psychology* 6:400–407.
- Cheshire, Coye and Karen S. Cook. 2004. “The Emergence of Trust Networks: Implications for Online Interaction.” *Analyse and Kritik* 26:220–240.
- Coleman, James S. 1988. “Free Riders and Zealots: The Role of Social Networks.” *Sociological Theory* 6:52–57.
- Davis, D. D., and Holt, C.A. 1993. *Experimental Economics*. Princeton, NJ: Princeton University Press
- Dawes, R.M. 1991. “Social Dilemmas, Economic Self-interest and Evolutionary Theory”. Pp 53–79 in *Recent Research in Psychology: Frontiers of Mathematical Psychology: Essays in Honor of Clyde Coombs*, edited by D.R. Brown and J.E.K. Smith. New York: Springer-Verlag.
- Dorsey, R. E. 1992. “The Voluntary Contributions Mechanism with Real Time Revisions.” *Public Choice* 72:261–282.
- Ekeh, Peter. 1974. *Social Exchange Theory: The Two Traditions*. Cambridge: Harvard University Press.
- Fischbacher, Urs, Simon Gächter, and Ernst Fehr. 2001. “Are People Conditionally Cooperative? Evidence From a Public Goods Experiment,” *Economics Letters* 71:397–404.
- Foa, Uriel G. and Edna B. Foa. 1974. *Societal structures of the mind*. Springfield, IL: Thomas.
- Gouldner, Alvin. 1960. “A Norm of Reciprocity: A Preliminary Statement.” *American Sociological Review* 25:161–178.
- Granovetter, Mark and Roland Soong. 1983. “Threshold Models of Diffusion and Collective Behavior.” *The Journal of Mathematical Sociology* 9:165–179.
- Healy, Kieran. 2001. “Exchange in Blood and Organs.” Ph.D. dissertation, Princeton University.
- Heckathorn, Douglas D. 1991. “Extensions of the Prisoner’s Dilemma Paradigm: The Altruist’s Dilemma and Group Solidarity.” *Sociological Theory* 9:34–52.
- . 1996. “The Dynamics and Dilemmas of Collective Action.” *American Sociological Review* 61:250–277.
- Hollander, Heinz. 1990. “A Social Exchange Approach to Voluntary Cooperation.” *The American Economic Review* 80:1157–1167.
- Isaac, R. M., and J.M.Walker. 1988a. “Communication and Free-Riding behavior: The Voluntary Contribution Mechanism.” *Economic Inquiry* 26:585–608.
- . 1988b. “Group Size Effects in Public Goods Provisions: The Voluntary Contributions Mechanism.” *Quarterly Journal of Economics* 103:179–200.
- Keser, C. and F. van Winden. 2000. “Conditional Cooperation and Voluntary Contributions to Public Goods.” *Scandinavian Journal of Economics* 102,1:23–39.
- Klendermans, Bert. 1984. “Mobilization and Participation: Social-Psychological Expansions of Resource Mobilization Theory.” *American Sociological Review* 49:583–600.
- Kollock, Peter. 1999a. “The Production of Trust in Online Markets.” *Advances in Group Processes*: 99–123.
- . 1999b. “The Economies of Online Cooperation: Gifts and Public Goods in Cyberspace.” Pp 220–239 in *Communities in cyberspace*, edited by Peter Kollock and Marc Smith. London: Routledge.
- Kurzban, R., K. McCabe, V. L. Smith, and B. J. Wilson. 2001. “Incremental Commitment and Reciprocity in a Real-Time Public Goods Game.” *Personality and Social Psychology Bulletin* 27:1662–1673.
- Lyman, Peter and Hal Varian. 2003. “How Much Information 2003?” Executive Summary, October 27, 2003. URL: http://www.sims.berkeley.edu/research/projects/how-much-info-2003/printable_execsum.pdf.
- Marwell, Gerald, and Pamela Oliver. 1993. *The Critical Mass in Collective Action: a Micro-social theory*. Cambridge ; New York: Cambridge University Press.
- Oliver, Pamela. 1980. “Rewards and Punishments as Selective Incentives for Collective Action: Theoretical Investigations.” *American Journal of Sociology* 85: 1356–75
- Oliver, Pamela E. and Gerald Marwell. 1988. “The Paradox of Group Size in Collective Action: A Theory of the Critical Mass. II.” *American Sociological Review* 53:1–8.
- Olson, Mancur. 1965. *The Logic of Collective Action: Public Goods and the Theory of Groups*. Cambridge, MA: Harvard University Press.
- Piliavin, Jane Allyn and Hong-Wen Charng. 1990. “Altruism: A Review of Recent Theory and Research.” *Annual Review of Sociology* 16:27–65.
- Raymond, Eric S. 1999. *The Cathedral & the Bazaar*. Sebastopol, California: O’Reilly.
- Sell, Jane. 1997. “Gender, Strategies, and Contributions to Public Goods.” *Social Psychology Quarterly* 3:252–265.
- Shah, S. K., and Sheen S. Levine. 2003. “Towards a Theory of Large-scale Generalized Exchange.” Conference Paper. American Sociological Association, Theoretical Issues in Economic Sociology Session. Atlanta, GA. August 2003.
- Simmons, Roberta G. 1992. “Giving Blood: The Development of an Altruistic Identity.” *Contemporary Sociology* 21:519–520.
- Sugden, Robert. 1984. “Reciprocity: The Supply of Public Goods through Voluntary Contributions.” *The Economic Journal* 94:772–787.
- Takahashi, Nobuyuki. 2000. “The Emergence of

- Generalized Exchange.” *American Journal of Sociology* 105:1105–1134.
- Yamagishi, Toshio. 1995. “Social Dilemmas.” Pp. 311–354 in *Sociological Perspectives on Social Psychology*, edited by Karen S. Cook, Gary A. Fine and James S. House. Needham Heights, NY: Allyn and Bacon.
- Yamagishi, Toshio, and Karen S. Cook. 1993. “Generalized Exchange and Social Dilemmas.” *Social Psychology Quarterly* 56: 235–248.

Coye Cheshire 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 include a series of theory-driven experiments examining shifts in modes of social exchange, a large-scale study of online dating and relationship formation, experimental studies of trust and trustworthiness in online settings, and the application of social psychological selective incentives to real-world Internet collective action problems.