Sharing-Mart
Online Auctions for Digital Content Trading and Content Incentivization

Christopher Leberknight, Ranjan Pal, Mung Chiang, and H. Vincent Poor
Email: rpal@usc.edu, {csl, chiangm, poor}@princeton.edu

***Abstract*—This paper introduces *Sharing-Mart* (S-Mart), an online digital trading platform developed at Princeton University to perform social file sharing experiments on top of technological networks as overlays. It describes the S-Mart system, the experiments conducted, and incentivization aspects associated with certain experiment types that might run on S-Mart. In the first part of the paper, the S-Mart system and the experiments conducted are explained, and the economic behaviors and dynamics of package auctions run on S-Mart are described. The major experimental observation that stands out here is that Internet users are not much incentivized to share content on competitive applications, whose success depends on the co-operation of other users in the system. To alleviate incentivization issues in these applications, in the second part of the paper a mathematical framework is proposed that derives user population threshold values, which hint at the necessity of a certain base population strength in S-Mart for co-operation to take place amongst all the users. An outline of two experiments to validate the theory is presented.**

***Keywords:* Sharing-Mart; Auctions; Incentives**

I. Introduction

The explosion of online content such as digital video, audio, and text, coupled with ubiquitous computing technologies has enabled new business opportunities and market places for digital content. For example, companies such as Netflix, Apple (iTunes), and Amazon (Kindle), all derive revenue based on business models involving the sale of digital content such as video, audio, and text, respectively. The pricing strategies for these digital goods are based on lower operational costs and other signals of demand from their tangible counterparts such as the sale of DVDs, music CDs, and books within traditional market places. The combination of existing business models based on digital content and the proliferation of Web 2.0 technologies such as blogs and social networking sites (i.e., YouTube, Vimeo, and Flickr) suggests the potential and feasibility of a market for user generated content. The value and demand for user generated content is much more complex to quantify and raises several challenges that are fundamental to classical economics. First, determining the value of user generated content is akin to monetizing information which is highly subjective and therefore cannot be accurately defined using a fixed or marked price. Second, classical economics emphasizes methods to efficiently allocate scarce resources primarily comprised of *private tangible goods*. However, user generated content within an electronic market place is more aptly characterized as a *public or intangible good* subject to *multiplicity* and *abundance*. To effectively address these two challenges, an open market pricing mechanism is required to enable consumers to provide the appropriate signaling function, specify preferences, and ration resources. We have developed a fully operational online virtual money file sharing system, known as **Sharing-Mart** (http://sharingmart.princeton.edu), which enables the transaction of digital goods using fixed price and multi-winner auctions. Sharing-Mart is the content equivalent of physical exchange markets such as eBay and Amazon and provides the opportunity to examine many different theoretical dimensions of auction theory, which can be substantiated or falsified through human subject experiments. In this paper, we[1] touch upon the following three important facets of the Sharing-Mart (S-Mart) system.

- We briefly describe the Sharing-Mart system, its architecture, and the principles behind its operation. (See Section II.)
- We give an overview of our experimental designs that were conducted on the S-Mart system, i.e., we describe the auction games played by users in each experiment and how subjects for the experiments were recruited via *Facebook*. (See Section III.)
- Through our experiments we study the economic behaviors and dynamics of package auctions for public goods in a virtual economy. We also observe through our preliminary experiments that cooperation amongst users (a desired property for the successful[2] working of the system) is not perceived as a dominant strategy, thereby indicating that our provided incentive to make users sell/contribute, is not powerful enough. (See Section IV.)
- In order to overcome the problem of strongly incentivizing users to contribute content, we propose voluntary as well as involuntary settings on the S-Mart framework, which are both realistic and practically implementable, and promote users to contribute content unselfishly. We

[1] R.Pal is with the Computer Science Department at the University of Southern California. C. Leberknight, M. Chiang, and H. V. Poor are with the Electrical Engineering Department at Princeton University.

[2] We emphasize here that for an application like *sharing course notes,* cooperation amongst the users is essential for increasing the overall knowledge of students, although the application is inherently competitive. Thus, in such an application, students need to contribute/sell their knowledge content for the better working of the system but may not, if strong incentives are not provided.

give an explanation of our settings, and describe the computation of population threshold parameters based on the settings. The population threshold parameters provide the S-Mart administrator with a valuable estimate of the mandatory presence of a certain number of users for an application in the S-Mart system so that the application is deemed effective in regard to unselfish content contribution. Our incentive settings and related analysis will also prove useful to knowledge management system (KMS) networks [4], [5], [6], in which a major challenge is to incentivize knowledge sharing amongst users by accounting for the dynamics of competition and cooperation at an organizational level. (See Section V.)

II. Sharing-Mart System

Sharing-Mart is a virtual money based file sharing system developed at Princeton University, which allows different digital rights (e.g., view only, download, and resell rights) of various file types (e.g., video, audio, graphics and documents) to be traded by means of different transaction styles (e.g., marked price transactions and multi-winner auctions). S-Mart has recently been integrated with Facebook to enable rapid and increased interaction and analysis of users within the S-Mart social network. Anyone with a Facebook account can use their Facebook identity to access S-Mart by going to http://sharingmart.princeton.edu. A screenshot of the methods to access the Sharing-Mart system is depicted in Figure 1.

The Sharing-Mart system can be accessed, either by visiting the main website http://sharingmart.princeton.edu and clicking on the Facebook icon, or by adding the application to a users application list on his/her main page within Facebook, by selecting the Sharing-Mart application from the Application Directory. The first time the system is accessed, a user will need to complete a brief registration form which will also initialize the user's personal homepage or directory within the system. Subsequently, each time a user accesses the application she will be presented with her personal login screen which will display information regarding past purchases and sales of digital content as well as other statistics regarding usage and popular content. Sharing-Mart has several features such as the *ability to view the most popular content*, *search for content*, *user statistics*, *and the ability to sell and purchase content using online auctions*. The full list of features as well as a user manual can be found at http://sharingmart.princeton.edu/HTML-User-Manual-v1.htm. The Sharing-Mart auction mechanism is implemented as a Vickrey-Clarke-Grove (VCG) auction, and at the same time enables the sale of multiple objects similar to package auctions [7].

III. Experimental Designs

In this section, we briefly describe two experimental designs that we conducted using the Sharing-Mart system. These experiments extend our previous research (open only to Princeton students) [1], in which agent-based competition was examined among seven graduate students at Princeton as a part of a homework problem set.

While auctions are generally competitive by nature, we further amplify the inherent competition through experiments which investigate the desire of students to cooperate with other students by contributing their knowledge in the form of digitized course notes. This allows us to not only examine different auction configurations such as package auctions [7] but also different patterns of human behavior.

A. Subject Recruitment

Subjects were recruited for two experiments from two undergraduate courses in social networking and systems analysis and design. There were 25 student subjects in each class. Additional subjects were also recruited using Facebook. A user account on Facebook, Share Mart, was used to recruit one hop neighbors from 5 students. A total of 64 members joined the Share Mart group and advertisements were broadcast to all Share Mart friends. A screen shot of the Share Mart account and advertisements are depicted in Figure 2.

Both experiments were based on a word game, and each subject was provided with details of the experiment and rules for winning the game. Winners in the first experiment would receive a monetary reward of up to $50.00 and winners of the second experiment would receive a monetary reward of up to $25.00.

B. Game Description

The objective of the word game used in each experiment is to obtain and correctly organize all the letters for a hidden word. Each subject is instructed to sell and buy letters using an auction in Sharing-Mart to collect all the letters. The winner of the game will receive a monetary prize if she/he is the first player to determine the word, and/or the player who maximizes his/her token balance by selling the letters. Therefore, players can win the game by selling and/or buying letters and there is no limit on the number of times each letter can be sold.

The game consists of two rounds with multiple auctions in each round. In **Round I**, the Sharing-Mart Operator (**id = smart**) creates l **sequential auctions**, where $l =$ **length of the word**. For example, if the *hidden word* is "cat," then $l = 3$. Therefore, the Sharing-Mart Operator would create 3 sequential auctions. That is, the auctions will not all occur at the same time. Each auction will sell 1 letter. The letters will not be visible to the players during the auctions. Players will bid on files containing a letter. Only the winner of the auction will see the contents of the file. For example, if the *hidden word* is "cat," 3 files labeled as 1.pdf, 2.pdf, 3.pdf will each be auctioned. Only one of the files will be sold in each auction. The files will be auctioned in increasing order. For example if the *hidden word* is "cat" the three auctions in Round I could be as follows:

1. Auction 1 sells 1.pdf
2. Auction 2 sells 2.pdf
3. Auction 3 sells 3.pdf

This will allow players to determine when Round I is over. For example, since the word contains three letters there will

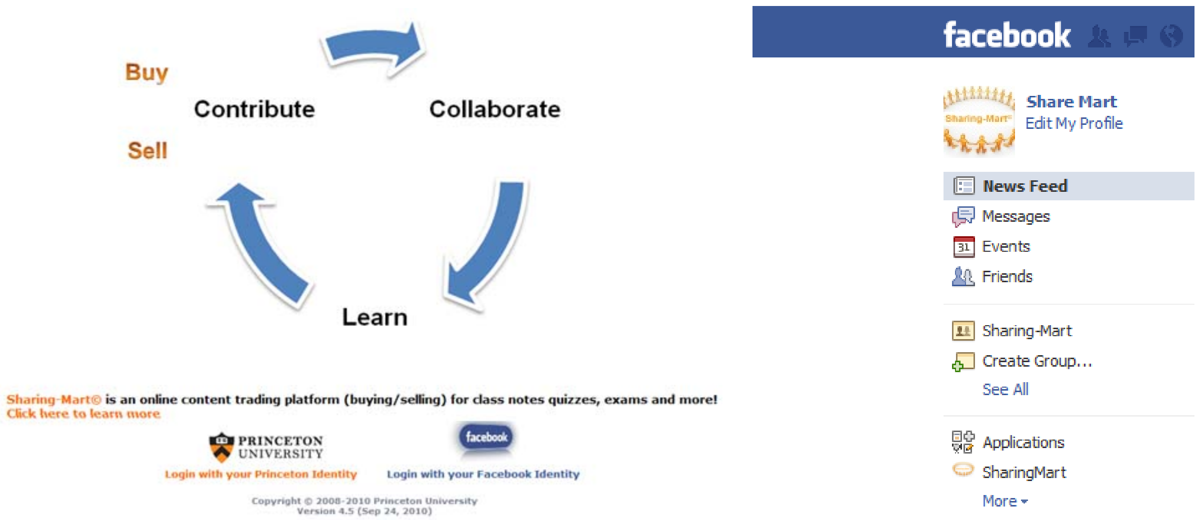

(a) Access via `http://sharingmart.princeton.edu` (b) Access within the Facebook application

Fig. 1 Accessing Sharing-Mart

Fig. 2 Facebook Share Mart Friends and Advertisements for Experiments

be three auctions. Therefore when you see the auction for 3.pdf you will know this is the last auction in Round I. At the end of Round I only the winners will know the contents of each file. Unless one of the players wins all three auctions the game continues to Round II. In Round II, each winner from Round I auctions his/her file(s) and everyone competes to purchase all the files/letters. This continues until a player has purchased all the letters and unscrambled the word. As soon as a player identifies the word she/he should send an email smart@princeton.edu with the solution to the hidden word. In addition to these details of the game subjects were also instructed to follow eight conditions, summarized in Table I, to ensure game accuracy and consistency among all players.

IV. Experimental Results

This section presents the results from two experiments in which subjects were required to purchase and sell four different files using the Sharing-Mart online auction module. Each experiment consisted of two rounds with multiple auctions in which subjects were required to purchase *all* the files. This type of auction configuration is most similar to a package auction [7]. The reason for designing the game as a package auction or auction of multiple objects with two rounds was to automate and simplify the process of generating items of interest for all subjects, and to discourage subjects from bidding their entire token balance in one auction. If the game consisted only of one auction, and the winner of the auction received the reward, then every player would be incentivized to bid his/her entire token balance to win the game and the demand function would simply be defined by the maximum token balance (all players have the same initial token balance). Since the tokens are virtual currency and do not represent any real value to the players, requiring players to purchase all files discourages players from bidding their entire token balance

1	The duration of each auction in Round I is six minutes.
2	Each player knows the length of the word, l. The length of the word for this experiment will be four letters.
3	The date and time of the experiment (Round 1 auctions) will be during one of your lectures determined by your Instructor.
4	The letters will not be visible in the auctions. The Sharing-Mart Operator will create l .pdf files. The .pdf files will be labeled 1.pdf, 2.pdf ... l.pdf. There are no duplicate letters. That is, all letters in the word will be unique.
5	Players will only be permitted to sell one instance of a letter during an auction in each auction, but they can sell the same letter in successive auctions.
6	Players should schedule their auctions in Round II to last between 5 and 30 minutes.
7	The game/experiment runs for one lecture (1 - 2 hours) or until one of the players determines the word.
8	All auctions in Round II must be scheduled during class lecture time. This is to ensure everyone is aware of the date and time of the auction which will help to increase the number of players during each auction.

TABLE I Game Conditions

Auction # (Filename)	**Letter**	**I-Price**	**F-Price**	**Letter**	**I-Price**	**F-Price**
1 (1.pdf)	N	16	450	F	2	325
2 (2.pdf)	S	14	550	C	3	500
3 (3.pdf)	K	11	600	A	8	500
4 (4.pdf)	I	19	850	L	4	490
Experiment	I Token Balance = 1200 Reward = $50.00			II Token Balance = 500 Reward = $25.00		
Round I	Keyword = "SINK"			Keywork = "CALF"		

TABLE II Initial and Final Prices for Experiments I and II

for one file and encourages them to sell the files they have won in previous auctions if their token balance drops below a certain threshold. This is the basis for designing the game as a package auction and motivating or incentivizing players to participate in the game. The two rounds are used to "seed" the players with items of interest to all players and then to let the dynamics of the game drive the activity in the system. In the first round the files were individually auctioned by the S-Mart Operator and subjects were required to bid on each file. This process "seeded" the group with digital content which would be in high demand by all players in the system. Since it is highly unlikely that one player would win all auctions in the first round, a player who won and paid too much for a file would have to sell her file to replenish her token balance in order to purchase the remaining files in the second round. This motivated players to contribute to the system by selling their files and let the dynamics of supply and demand drive the activity in the system. Also even more noteworthy is the implicit motivation for subjects to contribute content in the second round. This corresponds to encouraging cooperation among players in a competitive environment and is analogous to students contributing or selling their course notes to other students in the same class. Incentivizing cooperation is one of the main challenges in this research and also a main challenge for many other file-sharing systems and is typicallly known as the free rider problem [14], [15], [16], [17], [18], [19].

The hidden word for Experiment I was "sink" and the hidden word for Experiment II was "calf." In both experiments all subjects had the same initial token balance. In Experiment I, all subjects had an *initial balance of 1200 tokens* and in Experiment II, all subjects had an *initial balance of 500 tokens*. The letters were sold in four separate auctions in Round 1 with different initial prices (reserve prices). In addition, the reward for winning the game in Experiment I was equivalent to $50.00 and the reward for winning the game in Experiment II was $25.00. Details of initial price (I-Price) and final price (F-Price) for Round I of each experiment is provided in Table II.

We have developed several research questions to examine the issue of incentivizing content contribution in competitive environments. Three research questions form the basis for understanding economic behavior subject to the experimental configuration using the Sharing-Mart system. Our first research question investigates the final price subjects paid for content as a function of the initial price charged for the content and the subjects' token balances.

> **Q1**: *Are subjects with greater purchasing power will ultimately more active in the system, do they perceive greater benefits, and are they more incentivized to engage in future contributions compared to subjects who have less purchasing power?*

Our corresponding hypothesis states:

> **H1**: *Subjects who pay more for files perceive greater benefit and therefore will contribute/participate more in the system.*

We conjecture that if winners in Experiment I were more active than the winners in Experiment II because they had a greater initial token balances, then it may suggest Experiment I winners will contribute to the system more in the future compared to other users. For Experiment I, subject 1 and subject 7 won the auctions in Round I and for Experiment II subject 7 and subject 12 won the auctions in Round I. Based on the results from Figure 7a and Figure 7b it can be observed that subjects who won auctions in the first round did contribute/participate more in the system compared to other subjects. Specifically, the results indicate winners in Experiment I Round I were more active throughout the entire game compared to the winners in Round I of Experiment II. This supports our first hypothesis and suggests individuals with a larger token balance may participate more in the system compared to individuals with smaller token balances. We may therefore expect that one parameter that may influence contributions to the system is an individual's token balance. Individuals with larger token balances may have a greater perceived benefit of using the system compared to individuals who have smaller token balances. To confirm this conjecture we might investigate varying token balances across all users in the system to test whether subjects are incentivized to participate based on their initial token balances.

A partial goal of the hypotheses in this research is to determine whether the patterns or economic behavior of players in the experiments follow real world patterns that are observed in real economies and auctions. For example, is the concept of "winner's curse" possible in the experimental settings? Do subjects bid higher than the item is actually worth? The latter issue is difficult to resolve since the transactions are done with virtual currency. Specifically, since the experiment uses tokens in lieu of real money it is not easy to determine if subjects are paying more than the "true value." However to address this challenge real monetary incentives in the form of rewards are attached to the tokens. For example, a reward of $50.00 is used in the first experiment and a reward of $25.00 is used in the second experiment. Therefore, bidding too high may result in depleting the initial token balance to a point that would forfeit or severely limit the likelihood of winning the rewards. In a real economy with real currency it is expected that people who have larger budgets will bid higher than those with a smaller budget if the "true value" of the item auctioned is public information. For example, if the value of a piece of fine art is known to be $1000.00 and there are two bidders with budgets of $300.00 and $800.00, it is highly likely that the individual with the higher budget, $800.00 will win the auction taking into account rationality and utility maximization principles. Therefore, in the experiments involving the transaction of digital goods we expect that experiment I, with a reward of $50.00, will have higher final prices than the experiment II with a reward of $25.00. That is, files will have a higher demand and consequently higher sale price in Experiment I compared to Experiment II. Based on the results in Table II, the average final price over all four files is 612.5 tokens in Experiment I compared to 453.75 tokens in Experiment II. Therefore, the higher token balance in Experiment I of 1200 tokens resulted in a higher average final price (612.5) for all items compared to Experiment II which had a lower token balance of 500 tokens and lower average final price (453.75) for all items. However, it is unclear thus far from the analysis whether the higher final prices were due to the higher reward, higher initial token balance or arrangement of initial prices for each file. In addition, since the tokens merely serve as a proxy for the reward, the analysis of the reward amount may prove more useful. As a result, in future experiments the reward amounts will be reversed keeping all other parameters (token balance, initial price sequence) constant. This will help to clarify whether the reward amount or token balance is the stronger predictor of final prices.

Another parameter that may have influenced the total potential revenue is the duration of the auction. The effects of the duration of the auctions on the final prices and demand will be studied in future research, but a time series analysis of the initial sale prices, final purchase prices and bid amounts is provided in Figure 3 to provide insight into the different dynamics observed in the system. Subjects were required to use the same name as the file names which were used in Round I to ensure subjects would not accidentally purchase the same file more than once in Round II. The results confirm that not only were the final prices for each file higher in Experiment I compared to Experiment II, but they also demonstrate that the auctions in Round I of Experiment I were longer (longer red horizontal lines for each file) than the auctions in Round I in Experiment II.

Thus far, the analysis has examined only the activity corresponding to Round I in which players were only allowed to purchase files. Round II accounts for situations in which buyers can be sellers and vice versa. That is, players are free to choose whether they would like to purchase or sell content. Even though the reward, token balance, and auction duration are greater in Experiment I, an opposite demand effect is observed between Experiments in Round I and II. Specifically, there is more activity or demand in Round II in Experiment II compared to round II in Experiment I which contradicts the previous results between experiments I and II in Round I. In the case of Round I, longer auction durations and higher initial token balances and reward amounts, are associated with higher final prices. Therefore, within group (Experiment I and II) demand was greater in Round I compared to Round II, but between group demand (i.e. between experiments) was greater in Round I for Experiment I compared to Experiment II Round I and less in Experiment I Round II compared to Experiment II Round II. The differences in demand within and between experiments for both rounds is presented in Figure 4. The total number of bids per round and experiment is depicted in Figure 4. Here it can be observed that a total of 66 bids were placed in Round II in Experiment II and only 46 bids were placed in Round II in Experiment I. In addition, the higher number of bids within each experiment between Round I and Round II, undescore the challenge of incentivizing users to contribute content. More people bid in Round I when there were only buy options. In Round II, players could buy and sell and there were fewer bids which suggest some players chose not to bid.

To gain deeper insight and a possible explanation for the difference in demands, the next step in this analysis examines the activity in Round II to explicitly understand why lower values (token balance, reward amounts and auction durations) generated more activity in Round II in Experiment II compared to Round II in Experiment I. An analysis of the number of bids per file per round and the number of times files were sold per round is provided in Figure 5. The results highlight which files were in greatest demand and the saturation of files in the system. While the results indicate there are fewer bids in Round II of Experiment I compared to Round II in Experiment II, Figures 5(a) and 5(c), also illustrate that there were more winners in Round II of Experiment I compared to Round II in Experiment II, Figures 5(b) and 5(d). The larger number of transactions or bids observed in Round II of Experiment II suggests that demand for the files was greater compared to the demand for files in Round II of Experiment I. This is

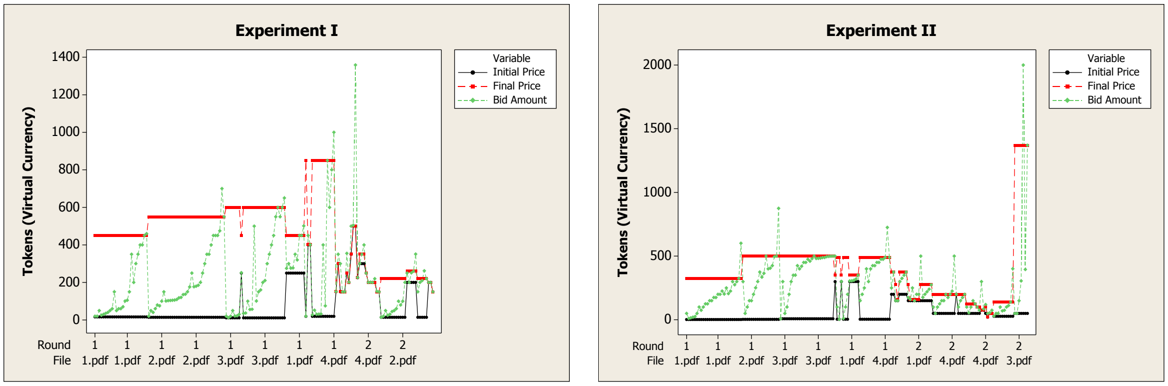

(a) (b)

Fig. 3 Time Series of Activity for Experiments I and II

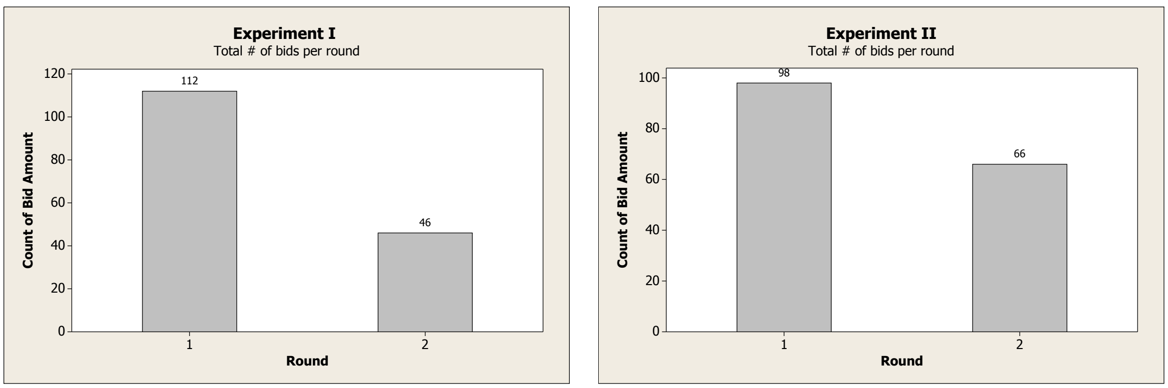

(a) (b)

Fig. 4 Total # of bids per round and experiment

because even though the incentive was lower in Round II (i.e. the rewards was \$25.00 compared to \$50.00 in Experiment I) it is likely that more bids were placed based on the observation that there were fewer winners. That is the supply of files in the system was lower for Experiment II Round II compared to Experiment I Round II. Fewer winners means the market was not as saturated with files since fewer subjects obtained the files which. According to the fundamental principles of supply and demand a lower supply is typically associated with a higher demand and higher price. An analysis of the ratio of the average initial prices to the token balance confirms that the behavior observed in these experiments corresponds to the expected behavior in a real economy. Specifically, since the ratio in Experiment II Round II is 16.8% compared to a ratio of 12.58% in Experiment I Round I, average initial prices for files relative to the respective initial token balances demonstrate that greater demand does indeed correspond to higher prices for the case of digital goods.

Our analysis now shifts focus from investigating the issues surrounding content contribution to examining the activity associated with winners of the game. We present the remaining two questions and corresponding hypotheses:

Q2: *Do players that bid the most number of times win the game?*
H2: *The higher the number of bids the higher the probability of winning the game.*

The intuition is that more aggressive or motivated players will bid more during the auctions compared to players who are less motivated, and are therefore more likely to win the game. However, we see that the three winners of the game in Experiment I, Figure 6(b) are subjects 8, 9 and 12 and only subject 8 is among the top three players who bid the most number of times (Figure 6(a) subjects who bid the most are subjects 7, 8 and 11). In Experiment II also showed similar results. In Experiment II there was only one winner, subject 7, who was not among the subjects who bid the most number of times (Figure 6(c) and 6(d)).

This suggests that other factors or behavior may have contributed to winning strategies and the results do not lend support for the third hypothesis which claims that subjects who bid the most number of times would most likely be among the winners of the game. Another factor that may be more likely to be correlated with the overall winning strategy is the final price paid for the files. This corresponds to the final question:

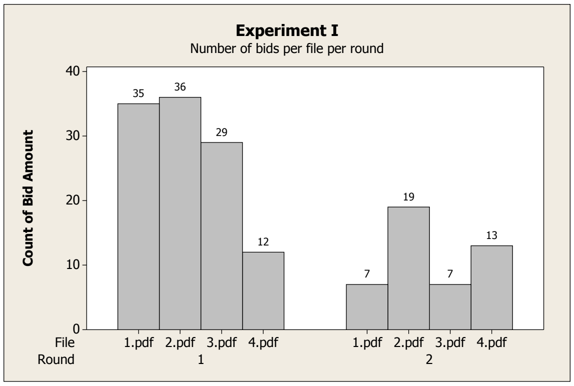

(a)

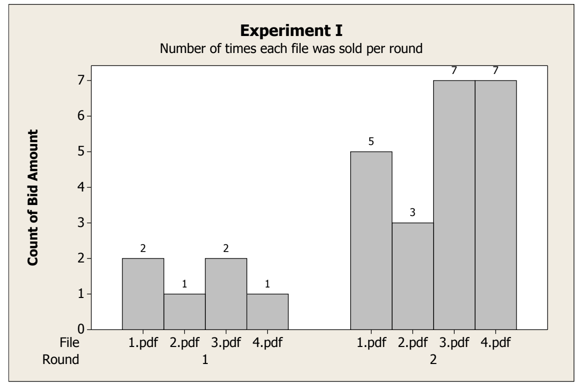

(b)

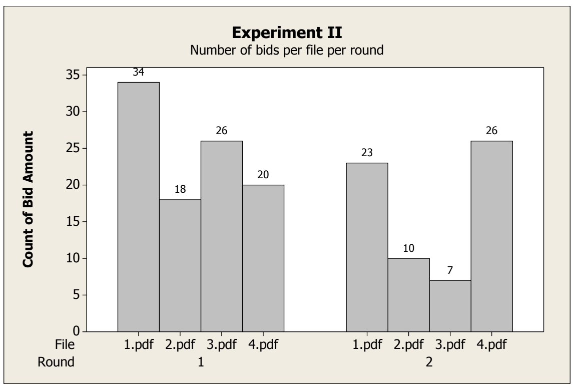

(c)

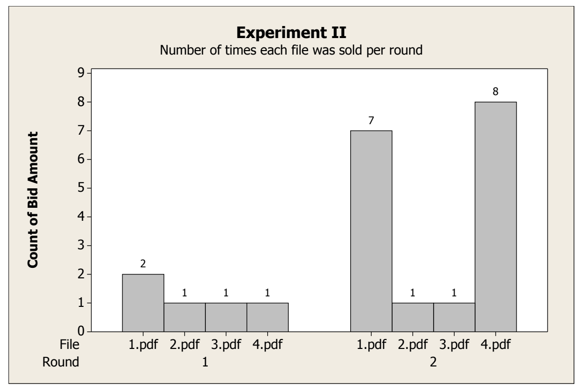

(d)

Fig. 5 Number of bids vs. files sold

Q3: *Do players who pay the highest amount on average win the game?*

In the experiments and auctions all players have common knowledge of the reward amounts. Therefore since the "true value" of the package auction is public information we might expect to observe that winners of the game would bid, on average, higher than other players and ultimately pay more for the files compared to other players. This is similar to the "winner's curse." That is, since players know the value the game and since it is played with virtual currency, players may be driven to bid high to win the auctions. The virtual currency is only a proxy for the actual reward and therefore may not discourage players from bidding too high. The following hypothesis examines the relation between the bid amount and winners of the game.

H3: *The higher the price paid for files the higher the probability of winning the game.*

Based on the results of the average prices paid by winners of the game, H3 is not supported.

The results in Figure 7 indicate that the highest average prices paid by winners of the auctions do not correlate with the winners of the game. For example, the winners for Experiment I were subjects 8, 9 and 12 and three other players paid more on average. Namely, subjects 1, 7 and 13. A similar observation is noted for Experiment II. There was only one winner in Experiment II, subject 7, and three other players paid more on average, subjects 8, 11 and 12. A possible explanation for the difference in the number of winners between the experiments is that Experiment I ran much longer than Experiment II which allowed more time for transactions.

V. Incentivizing User Cooperation

The major experimental observation that stands out from our experimental results in the previous section is that Internet users are not much incentivized to share content on competitive applications, whose success depends on the co-operation of other users in the system. In this section, we propose a mathematical framework that derives user population threshold values, which hint at the necessity of a certain base population strength in S-Mart for co-operation to take place amongst all the users. We compute (1) the *contributor threshold value* (CTV), which is defined as the minimum number of S-Mart users required to contribute valuable content (without any social influence) for all the S-Mart users to willfully contribute valuable content on a given topic, and (2) the *socially influenced population threshold* (SIPT), which we define to be the population of S-Mart users on a given topic required, in order to maximize the incentive of each S-Mart

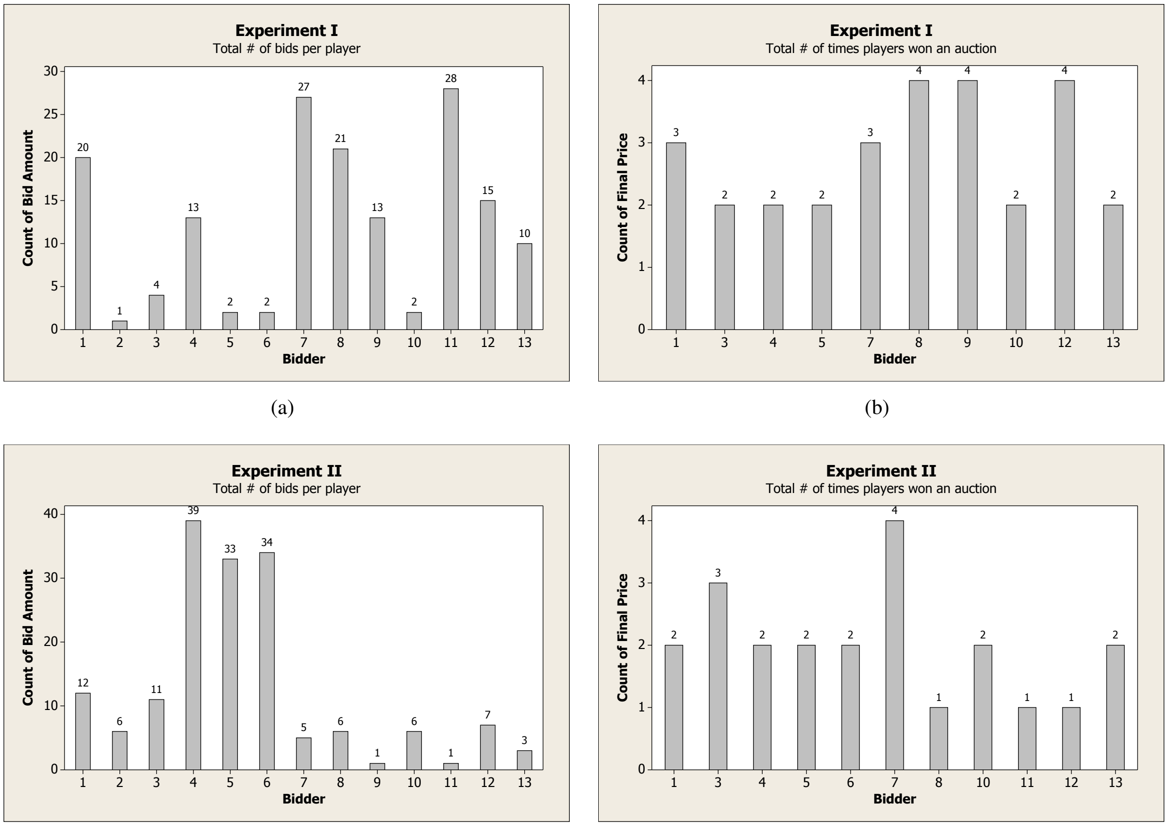

(a) (b) (c) (d)

Fig. 6 Number of bids vs. number of auctions won

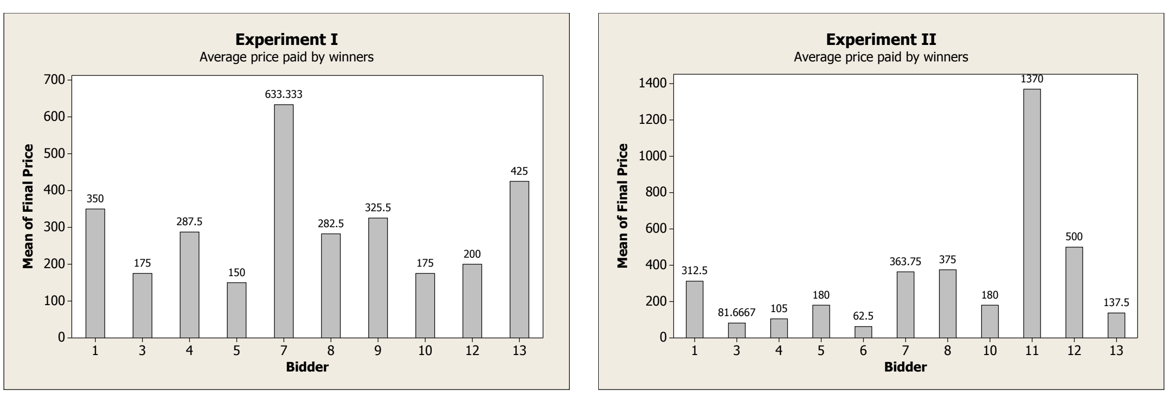

(a) (b)

Fig. 7 Average price paid by winners user to contribute valuable content on the given topic. We first provide an example application of a socially selfish file sharing system. Second, we describe our application setting as applied to the S-Mart system. Third, we define and briefly explain the notation is relevant to our mathematical framework. Finally, we outline two experiments that we wish to conduct in the near future, to validate the theory as proposed in the mathematical framework.

A. An Example of a Socially Selfish Application

The sharing of course notes and exercise hints/solutions amongst study friends is a common aspect of academic life. For example, students of MATH 101 (Calculus I) at a university may share notes about an exercise set on 'Lagrange Multipliers' with their friends in the same class, or elsewhere. The Internet today provides an ideal platform to share academic content of all types, viz., text, audio, video, and graphics. However, sharing academic content amongst those who are not study friends is not that common because of reasons attributed to student psychology. College and university academics are competitive and students want to gain a strategic advantage (ex., in terms of getting better grades) over their fellow students (students who are not study friends) by withholding information that they feel is valuable. Thus, knowledge flow, *if made possible* from a student/study group to other student/study groups would be deemed as valuable. In this regard, S-Mart provides a perfect platform for knowledge exchange to take place between individual students/study groups in return for virtual currency. However, in a socially selfish environment like in a school, college, or university, guaranteeing valuable information flow between competitors is a challenge. In this paper, we take steps to overcome this challenge.

B. Application Settings

We assume the following two types of settings.

1) S-Mart users share content on a particular topic without experiencing any regulation, social influence, quality demands, or central monitoring. This implies that S-Mart users may or may not contribute content, depending on their free will, and social friends of S-Mart users logged on to S-Mart *cannot* influence their friends to contribute. Content, if contributed by S-Mart users could be of any quality, and there is no central monitoring taking place to test S-Mart user misbehavior. By the term 'misbehavior', we mean either 'withholding' behavior, or 'cheating' behavior, i.e., contributing useless content, despite having good content. This setting is realistic of applications such as casual course notes sharing amongst students, where there is no pressure on anyone to contribute.

2) S-Mart users share content on a particular topic without any regulation, but there is a central monitoring system in place to detect[3] with a certain probability of success on whether S-Mart users withhold information or share low quality content despite having good quality content. Once user misbehavior above a certain level is detected, the S-Mart system can impose certain punishments on all the S-Mart users interested in a given topic. We will discuss punishments further in Section III. In this setting, S-Mart users may be influenced by social friends on not to withhold information so as to avoid global punishment laid down by S-Mart on all users relevant to a given topic. We assume here that social influence *always* motivates users to act altruistically. This influence can be exerted via *Facebook* like social sites, given that the S-Mart application is embedded in a social networking site. An example of an application fitting this setting is *faculty administered collaborative learning*. In this application, students of a class (ex., MATH 101) share course documents (related to homework sets) with other fellow students, and are awarded positive points for sharing valuable content, but the whole class gets negative points for exceeeding a certain degree of misbehavior (if detected)[4]. The points contribute to the final grades of the students in the class.

[3]The system can detect misbehavior based on information from other users, or by its own monitoring.

C. Notations

Throughout the rest of the paper, we use the term 'S-Mart user' and 'user' interchangeably. The following notation is also used.

n - fixed number of S-Mart users involved in a particular topic of interest (TOI). By the term 'involve', we imply that n users are logged on to Sharing Mart for learning about a specific topic. Ex., Lagrange Multipliers. A user could be both a consumer and a producer of content.

z_i - number of content units for a particular TOI of S-Mart user i. The content units imply the number of documents possessed by user i related to a particular TOI. The documents could be in the form of audio, video, text, or graphics.

v_{ij} - value of content unit j of user i. The value is a scalar number, which indicates the importance of unit j to user i.

c_{ij} - cost incurred by user i in contributing content unit j to the S-Mart repository. The cost is representative of a combination of 1) the effort of user i in accumulating j, 2) the importance of j to i, and 3) the strategic advantage i feels it would lose by contributing j. In this paper, we adopt a single scalar number for c_{ij} for the sake of simplicity. However, its an open question of how to appropriately quantify c_{ij} based on the three factors mentioned above, and this is a topic of our future work.

[4]Punishing everyone is a strategy to induce each user to contribute valuable content. Given that users are generally social and that Facebook like social websites can host an S-Mart application, it may not be that difficult to identify users who have cheated. Eventually the cheaters would end up losing their social value amongst friends due to the whole class suffering because of them.

TVC_i - total value of content information of user i on a particular TOI. It equals $\sum_{j=1}^{z_i} v_{ij}$.

cov_{ik} - degree of content overlap between users i and k. $cov_{ik} \,\epsilon\, [0,1]$. The degree of content overlap between two S-Mart users on a TOI is the measure of the 'sameness' of information possessed by i and k. For example, if two users A and B have the same content regarding a TOI then $cov_{AB} = 1$, and $cov_{AB} = 0$ if A and B have nothing in common with respect to TOI. cov_{AB} is between 0 and 1 if A and B have some overlap in their information content.

TVC_i^{ctrb} - amount of contribution to S-Mart by user i out of its total accumulation of a particular TOI, TVC_i. $TVC_i^{ctrb} = f_i \cdot TVC_i$, $f_i \,\epsilon\, [0,1]$.

TVC - total amount of information *open* to the n S-Mart users, who are interested in a particular TOI, and it equals $\sum_{i=1}^{n} TVC_i^{ctrb}$.

NUC_i - amount of new un-multiplicated content open to user i for potential use, i.e., this is the amount of information/knowledge that user i did not have with itself before w.r.t a particular TOI.

d_i - degree to which user i gains from NUC_i, $d_i \,\epsilon\, [0,1]$. It is the fraction of the new un-multiplicated content that is of interest to user i.

B_i - benefit to user i from NUC_i, which equals $d_i \cdot NUC_i$. The benefit here is in terms of new information gained.

NG_i - net gain of user i, which equals $B_i - f_i \sum_{j=1}^{z_i} c_{ij}$.

W_i - probability that user i withholds content or cheats.

D - probability that system detects misbehavior by users.

PM_i - punishment imposed upon user i when S-Mart detects misbehavior in the system.

C_i - social influence index of user i, where $C_i \,\epsilon\, [0,1]$.

CSI_i - cost to user i for not withholding valuable information/content due to social influence, when in fact user would have preferred selfish behavior without the social influence.

D. Computing CTV

Suppose there are n Sharing-Mart users comprising of content producers and consumers, on a certain topic of common interest. We assume that each producer in S-Mart has a certain initial amount of content with itself regarding the topic. Producers could be consumers and vice-versa. By the term 'topic', we refer to a subject, information about which is useful to the members of S-Mart interested in the subject, ex., the topic could be Lagrange Multipliers in a MATH 101 calculus course. Let $U_i(NC|\gamma)$ be the utility of a non-cooperative user, i, in S-Mart, when γ members in S-Mart decide to *contribute* content on a topic. Here, a non-cooperative user is an S-Mart user who either withholds information or provides low-quality content inspite of having better quality content. The *contributors* are assumed to be *altruistic* and share the best content they have with the S-Mart users. Similarly, we denote by $U_i(C|\gamma)$ the utility of the same non co-operative user, i, in S-Mart, when it turns co-operative (contributes), and γ members in S-Mart decide to contribute something on a topic. Throughout the rest of the paper, we use the terms 'co-operation' and 'contribution' interchangeably. We state the following relationship on an *individual* level:

$$U_i(NC|\gamma) > U_i(C|\gamma),\; 0 \leq \gamma < n. \tag{1}$$

The above inequality states that on an individual level, a non-cooperative S-Mart user is better off withholding content rather than sharing it with others, as it diminishes the user's strategic advantage. By withholding content, a non cooperative user enjoys all the benefits of other's contributions without giving anything away itself.

However, on the *group* level we derive the following relationship:

$$U_{grp}(0) < U_{grp}(k),\; k_t < k \leq n. \tag{2}$$

This equation implies that a group of size k greater than a threshold k_t, benefits in co-operation more that when no one in the group co-operates, because if everyone were to withhold content, there would be no benefit to the group, and in turn to any individual. We consider a group utility function to be the utility of S-Mart system. Thus, from equations (1) and (2), we observe that a user will not want to contribute individually, but might not benefit anything if all members in the group behave in the same manner. In this section, we propose a way to reverse the sense of inequality (1) such that S-Mart members are individually incentivized to contribute valuable content for the benefit of the system.

We formulate an individual user i's utility function when it contributes as

$$U_i(C) = d_i \cdot NUC_i - f_i \sum_{j=1}^{z_i} c_{ij} = B_i - f_i \sum_{j=1}^{z_i} c_{ij}. \tag{3}$$

It is evident that when user i decides to be non-cooperative, its utility function, $U_i(NC)$ equals $d_i \cdot NUC_i = B_i$. Thus $U_i(NC) > U_i(C)$. En-route to computing Contributor Threshold Value (CTV), we execute the following steps.

1) We derive k_t, the *minimum* number of S-Mart users amongst the n users, whose positive contribution results in the group utility being more than the utility when none of the users co-operate, i.e., $U_{grp}(0)$. k_t from equation 2 arises due to the fact that contributing content places a cost on users, and as a result the benefit due to co-operation amongst a certain number of users should exceed the cost of contribution before any group activity to take place.

2) Having executed the previous step, we ensure that contribution is efficient on the group level beyond a certain size. However, it does not help reverse the sense of inequality (1). In this step, we propose a system that provides bonuses to users who contribute, such that they can be compensated for their contribution costs. The system reverses the sense of inequality (1) and individually incentivizes S-Mart users to contribute content.

For purposes of getting a closed-form expression for CTV, we assume *homogeneity* of S-Mart users w.r.t the notations described in Section VC. Our simplified closed-form expression for the homogenous case also provides insights into the heterogenous case. Thus, in a system with n users, each user has z units of content about a particular TOI, each with identical value v. Each user incurs a cost c for contributing a unit of content and contributes a fraction, f of its total accumulation. The degree of content overlap between any two users on a particular TOI is cov, and the degree to which any user i gains from NUC_i is d. We also assume the events of content overlap between users i and j, and between users i and k are statistically independent. Therefore, NUC_i is evaluated in the following manner: starting with the content of an S-Mart user i, a second S-Mart user contributes $(1-cov)fzv$ new content, a third user contributes $(1-cov)(1-cov)fzv$ new content and so on. We thus have a geometric progression sum for NUC_i of the form

$$NUC_i = (1-cov)fzv + (1-cov)^2 fzv + ...(1-cov)^{n-1} fzv$$

or

$$NUC_i = \left(\frac{1-(1-cov)^n}{cov} - 1\right) fzv \tag{4}$$

The utility for an individual S-Mart contributing user i is then computed as

$$U_i(C) = d \cdot \left(\frac{1-(1-cov)^n}{cov} - 1\right) fzv - fzc \tag{5}$$

Summing up the utilities of individual contributing users, we get the utility for the whole group, i.e., S-Mart, w.r.t to a particular TOI as

$$U_{grp}(n) = n(d \cdot \left(\frac{1-(1-cov)^n}{cov} - 1\right) fzv - fzc) \tag{6}$$

The value of k_t is obtained by setting U_{grp} to 0, i.e., when on the group level the net benefit to S-Mart w.r.t a particular TOI is zero. Thus, we have

$$U_{grp}(k_t) = k_t(d \cdot \left(\frac{1-(1-cov)^{k_t}}{cov} - 1\right) fzv - fzc) = 0 \tag{7}$$

Solving for k_t, we get

$$k_t = \frac{\ln(1 - cov(1 + \frac{c}{vd}))}{\ln(1-cov)} \tag{8}$$

Any value of k above k_t results in a positive value of U_{grp}. Once we have the value of k_t, we know for sure that co-operative behavior by some threshold number of S-Mart users is better than joint non co-operation by all users.

However, the issue remains of how to make sure that each user could have a greater utility when choosing to contribute rather than when choosing not to contribute. To address this question, we propose the following simple *bonus* structure that provides S-Mart users with an incentive to contribute: for any unit of content contribution, the S-Mart system offers a bonus of c units to the contributing user. This is the amount which is equal to the cost incurred by a user in contributing a unit of content on a particular TOI. The intuition is that once we compensate users for their costs of contribution, they will be willing to contribute. S-Mart provides the bonus from its own net benefit, $U_{grp}(n)$. Thus, the overall benefit, $U1_{grp}(n)$, to S-Mart after providing total bonuses amounting to fzc to contributing users is $U_{grp}(n) - fzc$, which is equivalent to

$$U1_{grp}(n) = n(d \cdot \left(\frac{1-(1-cov)^n}{cov} - 1\right) fzv - fzc - fzc) \tag{9}$$

Therefore, the contributor threshold value, CTV, is computed by equating $U1_{grp}$ to zero. Thus, we have the following equation.

$$U1_{grp}(CTV) = CTV(d \cdot \left(\frac{1-(1-cov)^{CTV}}{cov} - 1\right) fzv - 2fzc) = 0 \tag{10}$$

Solving for CTV from (10), we get the value of the minimum number of S-Mart users required to contribute valuable content such that all the S-Mart users contribute valuable content a particular TOI. The closed form expression for CTV is given as

$$CTV = \frac{\ln(1 - cov(1 + \frac{2c}{vd}))}{\ln(1-cov)} \tag{11}$$

Sensitivity Analysis Based on the CTV expression, we observe that CTV values increase with increasing $\frac{c}{v}$ values. This is intuitive as the cost incurred by a user for sharing topic information increases w.r.t. the benefits obtained, and as a result users are less incentivized to contribute and the critical number increases. We also observe that CTV values increases with increasing d values. This result is intuitive as well because higher values of d imply that a user benefits more from the shared information pool and this happens only when the critical number increases.

E. Computing SIPT

In this section, we study the role of *social influence* and S-Mart punishments in ensuing co-operative behavior amongst S-Mart users in socially selfish applications (e.g., sharing course notes/lectures). It is evident that if *every* S-Mart user w.r.t to a TOI contributes, we are guaranteed a successful operating S-Mart system with every user doing its best to help the other users gain knowledge. However, in reality this is hardly the case. Users are non-cooperative by nature and do not want to share valuable content with others. In such situations, social influence from friends, or imposing punishments upon detecting selfish behavior could change user mindset in favor of contributing valuable content. Given the tremendous popularity of social networking websites, its not difficult to embed and administer educational S-Mart applications on a site like Facebook (refer to application setting 2 in Section I). In such cases, it is important that *each* user is incentivized to contribute for the benefit of the whole system. The entire system could represent a course in an university, in which one of the main goals of the instructor is to facilitate collaborative learning amongst students for altruistic knowledge dissemination. In this section, we compute the socially influenced population threshold (SIPT), which we define to be the number of members(users) needed in a system functioning on the S-Mart framework such that each user in the system is maximally incentivized to co-operate.

Let W_i be the probability that user i withholds or cheats on valuable information. Let D_i denote the probability that the S-Mart system detects this misbehavior. We define P_i to be the probability that user i withholds information and the system (we use the term 'system' and 'S-Mart system' interchangeably) detects it. We assume independence of the events that users cheat and the system detects, and denote P_i to be the product of W_i and D_i. We also assume that the S-Mart system imposes a punishment if it detects any user misbehaving. Thus,the probability P that at least one user withholds content is given as

$$P = 1 - (1 - P_1)(1 - P_2)........(1 - P_n) \tag{12}$$

or

$$P = 1 - \prod_{i=1}^{n}(1 - P_i) \tag{13}$$

or

$$P = 1 - \prod_{i=1}^{n}(1 - W_i \cdot D_i). \tag{14}$$

Given that the system punishes all the users once it detects any misbehavior, a user could either be 1) insensitive to any punishment, or 2) concerned about the punishment. Let U_i^{nc} denote the utility of a user not concerned with punishments imparted by S-Mart. We formulate U_i^{nc} as

$$U_i^{nc} = (U_i + PM_i)P + U_i(1 - P), \tag{15}$$

where U_i denotes the individual utility of an S-Mart user when it decides to co-operate, i.e., $U_i(C)$, or the utility when it chooses not to co-operate, i.e., $U_i(NC)$, and $PM_i < 0$ is the punishment imparted to user i by S-Mart. A user concerned with punishments would try its best to avoid it. Individually, he would not want to cheat, and would also want others not to cheat. One way a user could prevent others from misbehaving is by influencing its friends, who in turn influence their friends, and so on. Given that the S-Mart application is embedded in a social networking site friend influence should be possible. Let U_i^c denote the utility of a user concerned with punishments imparted by S-Mart. We formulate U_i^c as

$$U_i^c = (U_i + PM_i - CSI_i)P' + (U_i - CSI_i)(1 - P'), \tag{16}$$

where C_i is the social influence index of user i, where $C_i \,\epsilon\, [0, 1]$. This quantity indicates the degree to which a user is influenced by his friends to not withhold valuable content for the benefit of S-Mart. CSI_i denotes the cost to user i for not withholding valuable information/content due to social influence, when in fact he would have preferred selfish behavior without the social influence. P' is the probability that at least one user, after being socially influenced, is caught misbehaving in the system. We denote P' as

$$P' = 1 - \prod_{i=1}^{n} W_i \cdot (1 - C_i) \cdot D_i. \tag{17}$$

The difference in utility, U_i^{diff}, between user i's mindset of being concerned and unconcerned about S-Mart punishments is given as

$$U_i^{diff} = U_i^c - U_i^{nc}. \tag{18}$$

A user prefers being concerned about punishments to being unconcerned if $U_i^{diff} \geq 0$. This implies the following relationship after some algebraic computations

$$U_i^{diff} = PM_i(P' - P) - CSI_i \geq 0. \tag{19}$$

Considering a *homogenous* user system w.r.t. parameters, the above equation can be expressed as

$$U_i^{diff} = PM(P' - P) - CSI \geq 0, \tag{20}$$

where P' is given as

$$P' = [1 - (W(1 - C)D)^n]. \tag{21}$$

In order to obtain the number of members(users) needed in a system functioning on the S-Mart framework such that each user in the system is *maximally* incentivized to co-operate, we first take the derivative of U_i^{diff} w.r.t. n. The first-order derivative for the homogenous user case is

$$\frac{dU_i^{diff}}{dn} = PM \cdot \{(1 - W \cdot D)^n \cdot A - [1 - W \cdot (1 - C) \cdot D]^n \cdot B\}, \tag{22}$$

where

$$A = ln(1 - W \cdot D) \tag{23}$$

and

$$B = ln(1 - W \cdot (1 - C) \cdot D) \tag{24}$$

Once we have the first-order derivative, we equate it to zero to find the value of n for which U_i^{diff} reaches a maximum, i.e., the value of n for which each user in the S-Mart system interested in a particular TOI is maximally incentivized to cooperate. This value of n is the SIPT. Equating the first-derivative to zero, we get the value of SIPT as

$$SIPT = \frac{\ln\{\frac{\ln(1 - W \cdot D)}{\ln[1 - W \cdot (1 - C) \cdot D]}\}}{\ln\{\frac{[1 - W \cdot (1 - C) \cdot D]}{1 - W \cdot D}\}} \tag{25}$$

Sensitivity Analysis Based on the SIPT expression, we observe that SIPT values decrease with increase in the values of D. The intuition behind this result is the fact that with increasing values of D - the detection probability, the users would willingly contribute for the fear of punishments, even in the case of a low number of users present in the system. We also observe that the SIPT values decrease with increase in the C values. This result is also intuitive as with increasing social influence, it requires fewer of users to be present so as to maximize user willingness to contribute content. However, we see that the SIPT values increase with increase in W because an increase in the withholding probability of users implies the requirement of greater number of users in the system to maximize user willingness to contribute content.

F. Future Experiments

In this section we give an experimental outline of how to go about determining CTV and SIPT values empirically. Our goal of conducting the experiments is to validate the theory proposed in the mathematical framework. All the experiments are designed to facilitate *collaborative learning.*

Experiment outline to measure CTV: Assume a class assignment, for which students are required to write a class report regarding a given topic. (E.g., a survey paper on routing protocols in wireless networks.) The students are evaluated based on the quality of the report, which is determined by the number of salient points in the report. It is obvious that most students would search the Internet for related literature on the topic. However, apart from some very common information, students would vary w.r.t. one another in terms of topic points. We assume that there is a class organizer such as the Professor or the teaching assistant (TA). We plan to conduct the experiment in two rounds. In the first round, the Professor/TA gathers topic points separately from each student as part of 'ongoing progress'. In the second round, the Professor/TA opens up an online discussion board, in which students could share their topic points. The sharing is not made compulsory; however, if students share their knowledge, they are awarded a certain number of points for their contributions. The flip side to this benefit is that students might lose a competitive advantage to other competitors. We measure this loss of advantage in terms of a cost. The CTV value could be estimated by the Professor/TA in the second round by observing the rate at which students upload content points. We expect a sudden surge of content uploads over time. We need to keep track of the time when the surge occurs, and identify the number of users just before the surge occurs. This number will give us an estimate of the CTV.

Experiment outline to measure SIPT: We design an experiment along similar lines as the one to measure CTV. The only difference is that we incorporate social influence, misbehavior detection, and punishments. Social influence is a natural property and is not within the control of the Professor/TA. We assume here that a student may be positively influenced through chat or Facebook like mechanisms to willfully contribute content. However, it is clear that greater the number of friends of a person, the higher the chances of it being positively influenced. We capture misbehavior detection via student complaints to the Professor/TA about someone having information and not sharing it, or someone willfully sharing wrong information. We emphasize here that there is an inherent challenge in detecting user misbehavior as 1) its difficult to prove that someone is unwilling to share information, and 2) even if students share information and the information is wrong, we cannot infer for sure that the students wanted to harm other students. It could well be the case that their information sources were non-credible, and it is infeasible for anyone to check all Internet sources. We need to address these issues in our experiment. Having detected user misbehavior, punishments are in the form of points deducted from every student in the class. Similar to the experiment to measure CTV, the experiment here will have two rounds and the Professor/TA can estimate the SIPT from the second round by observing the contributing population count at the time when a surge of content contribution occurs.

VI. Conclusion and Future Work

In this paper, we have described Sharing-Mart, a virtual file sharing platform, and have investigated whether similar economic behavior is observed in the virtual economy as in a real economy. Three research questions and hypotheses (H1, H2, and H3) have been presented to understand the economic behaviors and dynamics of the package auction for public goods using the Sharing-Mart system. H1, which states subjects who initally pay more for content are more likley to particiapte more in the system compared to subjects who pay less, appears to be true. However, analysis of the results for H2 and H3 indicate neither hypothesis is supported. Therefore, while higher bid amounts may correpsond to high activity or participation the most active bidders and users who paid the most for content do not necessarily win the game. While the limited sample size does not allow for more rigorous statistical analysis and formalized hypothesis testing, the results from the analysis have provided insight into future directions and merit further analysis. The experimental observations also highlight the fact that users are not willing to co-operate on competitive applications, whose success depends on mutual cooperation amongst the users. To alleviate this problem we have proposed a mathematical framework that derives user population threshold values, which hint at the necessity of a certain base population strength in S-Mart for co-operation to take place amongst all the users.

As part of future research, we will run the experiments, whose outlines we have provided in the previous section. We will examine how the value or demand for public goods varies over time and the effects of the auction duration on the final prices paid for the content. In addition, based on the analysis of empirical results, more detailed analysis of the reward amount will be conducted in two different ways. First, by reversing the reward amounts between experiments and keeping all other parameters constant we will be able to confirm whether the reward amount is ultimately related to the demand for files. Second, we plan on randomizing reward amounts from a list of rewards to create an auction, which more accurately represents a real auction in which players have different values for the same item. The experimental configuration in this research afforded all players common knowledge of the reward amount. That is, the game design was based on public information. In our future experiments, players will be provided with different values regarding the reward amount which will more accurately simulate a real world auction scenario. Therefore, players will have private information regarding the value of the package auction.

VII. Acknowledgement

This was was supported in part by the National Science Foundation under Grant CNS-09-05086. We would also like to thank Felix Ming-Fai Wong for helping us with the editing of the paper.

References

[1] Inaltekin, H., Ogzencil, H., Xu, H., Chiang, M., and Poor, H. V. (2010) “Online content-trading: Sharing Mart system and auction experiments,” Proc. IEEE Workshop on Social Networking, Cape Town, South Africa, May 2010.

[2] Alavi, M., and Leidner, D. E. Review: Knowledge Management and Knowledge Management Systems: Conceptual Foundations and Research Issues, MIS Quarterly (25:1), 2001, pp. 107-136.

[3] Alavi, M., and Leidner, D. E. Knowledge management systems: Issues, challenges, and benefits. Communications of the AIS, 1, 7 (February, 1999)

[4] King, W. R., and Marks, Jr., P. V. Motivating knowledge sharing through a knowledge management system, Omega 36 (2008) 131146.

[5] Lin, T.-C., and Huang, C.-C. (2008) Understanding knowledge management system usage antecedents: an integration of social cognitive theory and task technology fit. Information & Management 45(6), 410417.

[6] Kankanhalli, A., Tan, B. C. Y., and Wei, K.-K. “Contributing Knowledge to Electronic Knowledge Repositories: An Empirical Investigation,” MIS Quarterly (29:1), March 2005, pp. 113-143.

[7] Milgrom, P. “Putting Auction Theory to Work: The Simultaneous Ascending Auction.” Journal of Political Economy, 2000, 108, pp. 24572.

[8] Krishna, V. (2002) Auction Theory Academic Press, San Diego, USA.

[9] Milgrom, P. (2004) Putting Auction to Work Cambridge University Press.

[10] Bichler, M. (2000). “Trading Financial Derivatives on the Web. An Approach Towards Automating Negotiations on the OTC Markets.” Information Systems Frontiers 1(4).

[11] Klemperer, Paul D. ed. 2000a The Economic Theory of Auctions. Cheltenham, U.K.: Edward Elgar.

[12] Klemperer, Paul D. 1999a. “Auction Theory: A Guide to the Literature.” Journal of Economic Surveys, 13:3, pp. 227-86.

[13] Edelman, B., Ostrovsky, M., and Schwarz, M. (2007) “Internet Advertising and the Generalized Second-Price Auction: Selling Billions of Dollars Worth of Keywords,” American Economic Review 97(1), 242-259.

[14] Adar, E., and Huberman, B. A. “Free riding on Gnutella,” Technical report, Xerox PARC, August 2000.

[15] Hughes, D., Coulson, G., and Walkerdine, J. “Free riding on Gnutella revisited: The bell tolls?,” IEEE Dist. Systems Online, vol. 6, no. 6, 2005.

[16] Jun, S., and Ahamad, M. “Incentives in BitTorrent induce free riding,” Proceedings of the 2005 ACM SIGCOMM Workshop on Economics of Peer-to-Peer Systems, pp. 116-121, Philadelphia, PA, 2005.

[17] Liogkas, N., Nelson, R., Kohler, E., and Zhang, L. “Exploiting BitTorrent for fun (but not profit),” Proceedings of the International Workshop on Peer-to-Peer Systems (IPTPS), Santa Barbara, CA, 2006.

[18] Piatek, M., Isdal, T., Anderson, T., Krishnamurthy, A., and Venkataramani, A. “Do incentives build robustness in BitTorrent?,” Proceedings of 4th USENIX Symposium on Networked Systems Design & Implementation (NSDI 2007), Cambridge, MA, April 2007.

[19] Sirivianos, M., Park, J. H., Chen, R., and Yang, X. “Free-riding in BitTorrent networks with the large view exploit,” Proceedings of the 6th International Workshop on Peer-to-Peer Systems (IPTPS), Bellevue, WA, 2007.