Efficiency and Nash Equilibria in a Scrip System for P2P Networks

Eric J. Friedman
School of Operations Research and Industrial Engineering
Cornell University
ejf27@cornell.edu

Joseph Y. Halpern
Computer Science Dept.
Cornell University
halpern@cs.cornell.edu

Ian Kash
Computer Science Dept.
Cornell University
kash@cs.cornell.edu

ABSTRACT

A model of providing service in a P2P network is analyzed. It is shown that by adding a scrip system, a mechanism that admits a reasonable Nash equilibrium that reduces free riding can be obtained. The effect of varying the total amount of money (scrip) in the system on efficiency (i.e., social welfare) is analyzed, and it is shown that by maintaining the appropriate ratio between the total amount of money and the number of agents, efficiency is maximized. The work has implications for many online systems, not only P2P networks but also a wide variety of online forums for which scrip systems are popular, but formal analyses have been lacking.

1. INTRODUCTION

A common feature of many online distributed systems is that individuals provide services for each other. Peer-to-peer (P2P) networks (such as Kazaa [30] or BitTorrent [3]) have proved popular as mechanisms for file sharing, and applications such as distributed computation and file storage are on the horizon; systems such as Seti@home [29] provide computational assistance; systems such as Slashdot [24] provide content, evaluations, and advice forums in which people answer each other's questions. Having individuals provide each other with service typically increases the social welfare: the individual utilizing the resources of the system derives a greater benefit from it than the cost to the individual providing it. However, the cost of providing service can still be nontrivial. For example, users of Kazaa and BitTorrent may be charged for bandwidth usage; in addition, in some file-sharing systems, there is the possibility of being sued, which can be viewed as part of the cost. Thus, in many systems there is a strong incentive to become a *free rider* and benefit from the system without contributing to it. This is not merely a theoretical problem; a 2000 study of the Gnutella [25] network showed that almost 70 percent of users share no files and nearly 50 percent of responses are from the top 1 percent of sharing hosts [1]. A followup study has shown that over time the disparity has only gotten worse [17]

EC'06, June 11–15, 2006, Ann Arbor, Michigan.

Having relatively few users provide most of the service creates a point of centralization; the disappearance of a small percentage of users can greatly impair the functionality of the system. Moreover, current trends seem to be leading to the elimination of the "altruistic" users on which these systems rely. These heavy users are some of the most expensive customers ISPs have. Thus, as the amount of traffic has grown, ISPs have begun to seek ways to reduce this traffic. Some universities have started charging students for excessive bandwidth usage; others revoke network access for it [5]. A number of companies have also formed whose service is to detect excessive bandwidth usage [22].

These trends make developing a system that encourages a more equal distribution of the work critical for the continued viability of P2P networks and other distributed online systems. A significant amount of research has gone into designing reputation systems to give preferential treatment to users who are sharing files. Some of the P2P networks currently in use have implemented versions of these techniques. However, these approaches tend to fall into one of two categories: either they are "barter-like" or reputational. By barter-like, we mean that each agent bases its decisions only on information it has derived from its own interactions. Perhaps the best-known example of a barter-like system is BitTorrent, where clients downloading a file try to find other clients with parts they are missing so that they can trade, thus creating a roughly equal amount of work. Since the barter is restricted to users currently interested in a single file, this works well for popular files, but tends to have problems maintaining availability of less popular ones. An example of a barter-like system built on top of a more traditional file-sharing system is the credit system used by eMule [8]. Each user tracks his history of interactions with other users and gives priority to those he has downloaded from in the past. However, in a large system, the probability that a pair of randomly-chosen users will have interacted before is quite small, so this interaction history will not be terribly helpful. Anagnostakis and Greenwald [2] present a more sophisticated version of this approach, but it still seems to suffer from similar problems.

A number of attempts have been made at providing general reputation systems (e.g. [14, 15, 20, 32]). The basic idea is to aggregate each user's experience into a global number for each individual that intuitively represents the system's view of that individual's reputation. However, these attempts tend to suffer from practical problems because they implicitly view users as either "good" or "bad", assume that the "good" users will act according to the specified protocol, and that there are relatively few "bad" users. Unfortunately, if there are easy ways to game the system, once this information becomes widely available, rational users are likely to make use of it. We cannot count on only a few users being "bad" (in the sense of not following the prescribed protocol). For example, Kazaa uses a measure of the ratio of the number of uploads to the number of downloads to identify good and bad users. However, to avoid penalizing new users, they gave new users an average rating. Users discovered that they could use this relatively good rating to free ride for a while and, once it started to get bad, they could delete their stored information and effectively come back as a "new" user, thus circumventing the system (see [2] for a discussion and [11] for a formal analysis of this "whitewashing"). Thus Kazaa's reputation system is ineffective.

This is a simple case of a more general vulnerability of such systems to *sybil* attacks [6], where a single user maintains multiple identities and uses them in a coordinated fashion to get better service than he otherwise would. Recent work has shown that most common reputation systems are vulnerable to such attacks [4]; however, the degree of this vulnerability is still unclear. The analyses of the practical vulnerabilities and the existence of such systems that are immune to such attacks remains an area of active research (e.g., [4, 16, 33]).

Simple economic systems based on a scrip or money seem to avoid many of these problems, are easy to implement, and are quite popular (see, e.g., [15, 18, 31]). However, they have a different set of problems. Perhaps the most common involve determining the amount of money in the system. Roughly speaking, if there is too little money in the system relative to the number of agents, then relatively few users can afford to make request. On the other hand, if there is too much money, then users will not feel the need to respond to a request; they have enough money already. A related problem involves handling newcomers. If newcomers are each given a positive amount of money, then the system is open to sybil attacks. Perhaps not surprisingly, scrip systems end up having to deal with standard economic woes such as inflation, bubbles, and crashes [31].

In this paper, we provide a formal model in which to analyze scrip systems. We describe a simple scrip system and show that, under reasonable assumptions, for each fixed amount of money there is a nontrivial Nash equilibrium involving *threshold strategies*, where an agent accepts a request if he has less than $\$k$ for some threshold k.[1] An interesting aspect of our analysis is that, in equilibrium, the distribution of users with each amount of money is the distribution that maximizes entropy (subject to the money supply constraint). This allows us to compute the money supply that maximizes *efficiency* (social welfare), given the number of agents. It also leads to a solution for the problem of dealing with newcomers: we simply assume that new users come in with no money, and adjust the price of service (which is equivalent to adjusting the money supply) to maintain the ratio that maximizes efficiency. While assuming that new users come in with no money will not work in all settings, we believe the approach will be widely applicable. In systems where the goal is to do work, new users can acquire money by performing work. It should also work in Kazaa-like system where a user can come in with some resources (e.g., a private collection of MP3s).

The rest of the paper is organized as follows. In Section 2, we present our formal model and observe that it can be used to understand the effect of altruists. In Section 3, we examine what happens in the game under nonstrategic play, if all agents use the same threshold strategy. We show that, in this case, the system quickly converges to a situation where the distribution of money is characterized by maximum entropy. Using this analysis, we show in Section 4 that, under minimal assumptions, there is a nontrivial Nash equilibrium in the game where all agents use some threshold strategy. Moreover, we show in Section 5 that the analysis leads to an understanding of how to choose the amount of money in the system (or, equivalently, the cost to fulfill a request) so as to maximize efficiency, and also shows how to handle new users. In Section 6, we discuss the extent to which our approach can handle sybils and collusion. We conclude in Section 7.

2. THE MODEL

To begin, we formalize providing service in a P2P network as a non-cooperative game. Unlike much of the modeling in this area, our model will model the asymmetric interactions in a file sharing system in which the matching of players (those requesting a file with those who have that particular file) is a key part of the system. This is in contrast with much previous work which uses random matching in a prisoner's dilemma. Such models were studied in the economics literature [7, 21] and first applied to online reputations in [11]; an application to P2P is found in [9]. A variant of this model that uses the multiperson prisoner's dilemma rather than random matching can be found in [27]. Some work on micropayments is similar in spirit, but it assumes random matchings and only allows agents to decide on a rough sharing level which prevents the use of more nuanced strategies such as threshold strategies [12]

This random-matching model fails to capture some salient aspects of a number of important settings. When a request is made, there are typically many people in the network who can potentially satisfy it (especially in a large P2P network), but not all can. For example, some people may not have the time or resources to satisfy the request. The random-matching process ignores the fact that some people may not be able to satisfy the request. Presumably, if the person matched with the requester could not satisfy the match, he would have to defect. Moreover, it does not capture the fact that the decision as to whether to "volunteer" to satisfy the request should be made before the matching process, not after. That is, the matching process does not capture the fact that if someone is unwilling to satisfy the request, there will doubtless be others who can satisfy it. Finally, the actions and payoffs in the prisoner's dilemma game do not obviously correspond to actual choices that can be made. For example, it is not clear what defection on the part of the requester means. In our model we try to deal with all these issues.

Suppose that there are n agents. At each round, an agent is picked uniformly at random to make a request. Each other

[1] Although we refer to our unit of scrip as the dollar, these are not real dollars nor do we view them as convertible to dollars.

agent is able to satisfy this request with probability $\beta > 0$ at all times, independent of previous behavior. The term β is intended to capture the probability that an agent is busy, or does not have the resources to fulfill the request. Assuming that β is time-independent does not capture the intution that being an unable to fulfill a request at time t may well be correlated with being unable to fulfill it at time $t+1$. We believe that, in large systems, we should be able to drop the independence assumption, but we leave this for future work. In any case, those agents that are able to satisfy the request must choose whether or not to volunteer to satisfy it. If at least one agent volunteers, the requester gets a benefit of 1 util (the job is done) and one of volunteers is chosen at random to fulfill the request. The agent that fulfills the request pays a cost of $\alpha < 1$. As is standard in the literature, we assume that agents discount future payoffs by a factor of δ per time unit. This captures the intuition that a util now is worth more than a util tomorrow, and allows us to compute the total utility derived by an agent in an infinite game. Lastly, we assume that with more players requests come more often. Thus we assume that the time between rounds is $1/n$. This captures the fact that the systems we want to model are really processing many requests in parallel, so we would expect the number of concurrent requests to be proportional to the number of users.[2]

Let $G(n, \delta, \alpha, \beta)$ denote this game with n agents, a discount factor of δ, a cost of α to satisfy requests, and a probability β of being able to satisfy requests. When the latter two parameters are not relevant, we sometimes write $G(n, \delta)$.

We use the following notation throughout the paper:

- p^t denotes the agent chosen in round t.
- $B_i^t \in \{0, 1\}$ denotes whether agent i can satisfy the request in round t. $B_i^t = 1$ with probability $\beta > 0$ and B_i^t is independent of $B_i^{t'}$ for all $t' \neq t$.
- $V_i^t \in \{0, 1\}$ denotes agent i's decision about whether to volunteer in round t; 1 indicates volunteering. V_i^t is determined by agent i's strategy.
- $v^t \in \{j \mid V_j^t B_j^t = 1\}$ denotes the agent chosen to satisfy the request. This agent is chosen uniformly at random from those who are willing ($V_j^t = 1$) and able ($B_j^t = 1$) to satisfy the request.
- u_i^t denotes agent i's utility in round t.

  A *standard* agent is one whose utility is determined as discussed in the introduction; namely, the agent gets a utility of 1 for a fulfilled request and utility $-\alpha$ for fulfilling a request. Thus, if i is a standard agent, then

$$u_i^t = \begin{cases} 1 & \text{if } i = p_t \text{ and } \sum_{j \neq i} V_j^t B_j^t > 0 \\ -\alpha & \text{if } i = v^t \\ 0 & \text{otherwise.} \end{cases}$$

- $U_i = \sum_{t=0}^{\infty} \delta^{t/n} u_i^t$ denotes the total utility for agent i. It is the discounted total of agent i's utility in each round. Note that the effective discount factor is $\delta^{1/n}$ since an increase in n leads to a shortening of the time between rounds.

Now that we have a model of making and satisfying requests, we use it to analyze free riding. Take an altruist to be someone who always fulfills requests. Agent i might rationally behave altruistically if agent i's utility function has the following form, for some $\alpha' > 0$:

$$u_i^t = \begin{cases} 1 & \text{if } i = p_t \text{ and } \sum_{j \neq i} V_j^t B_j^t > 0 \\ \alpha' & \text{if } i = v^t \\ 0 & \text{otherwise.} \end{cases}$$

Thus, rather than suffering a loss of utility when satisfying a request, an agent derives positive utility from satisfying it. Such a utility function is a reasonable representation of the pleasure that some people get from the sense that they provide the music that everyone is playing. For such altruistic agents, playing the strategy that sets $V_i^t = 1$ for all t is dominant. While having a nonstandard utility function might be one reason that a rational agent might use this strategy, there are certainly others. For example a naive user of filesharing software with a good connection might well follow this strategy. All that matters for the following discussion is that there are some agents that use this strategy, for whatever reason.

As we have observed, such users seem to exist in some large systems. Suppose that our system has a altruists. Intuitively, if a is moderately large, they will manage to satisfy most of the requests in the system even if other agents do no work. Thus, there is little incentive for any other agent to volunteer, because he is already getting full advantage of participating in the system. Based on this intuition, it is a relatively straightforward calculation to determine a value of a that depends only on α, β, and δ, but not the number n of players in the system, such that the dominant strategy for all standard agents i is to never volunteer to satisfy any requests (i.e., $V_i^t = 0$ for all t).

PROPOSITION 2.1. *There exists an a that depends only on α, β, and δ such that, in $G(n, \delta, \alpha, \beta)$ with at least a altruists, not volunteering in every round is a dominant strategy for all standard agents.*

PROOF. Consider the strategy for a standard player j in the presence of a altruists. Even with no money, player j will get a request satisfied with probability $1 - (1 - \beta)^a$ just through the actions of these altruists. Thus, even if j is chosen to make a request in every round, the most additional expected utility he can hope to gain by having money is $\sum_{k=1}^{\infty} (1 - \beta)^a \delta^k = (1 - \beta)^a/(1 - \delta)$. If $(1 - \beta)^a/(1 - \delta) > \alpha$ or, equivalently, if $a > \log_{1-\beta}(\alpha(1 - \delta))$, never volunteering is a dominant strategy. □

Consider the following reasonable values for our parameters: $\beta = .01$ (so that each player can satisfy 1% of the requests), $\alpha = .1$ (a low but non-negligible cost), $\delta = .9999$/day (which corresponds to a yearly discount factor of approximately 0.95), and an average of 1 request per day per player. Then we only need $a > 1145$. While this is a large number, it is small relative to the size of a large P2P network.

Current systems all have a pool of users behaving like our altruists. This means that attempts to add a reputation system on top of an existing P2P system to influence users

[2]For large n, our model converges to one in which players make requests in real time, and the time between a player's requests is exponentially distributed with mean 1. In addition, the time between requests served by a single player is also exponentially distributed.

to cooperate *will have no effect on rational users.* To have a fair distribution of work, these systems must be fundamentally redesigned to eliminate the pool of altruistic users. In some sense, this is not a problem at all. In a system with altruists, the altruists are presumably happy, as are the standard agents, who get almost all their requests satisfied without having to do any work. Indeed, current P2P network work quite well in terms of distributing content to people. However, as we said in the introduction, there is some reason to believe these altruists may not be around forever. Thus, it is worth looking at what can be done to make these systems work in their absence. For the rest of this paper we assume that all agents are standard, and try to maximize expected utility.

We are interested in equilibria based on a scrip system. Each time an agent has a request satisfied he must pay the person who satisfied it some amount. For now, we assume that the payment is fixed; for simplicity, we take the amount to be \$1. We denote by M the total amount of money in the system. We assume that $M > 0$ (otherwise no one will ever be able to get paid).

While a decision to volunteer can in general depend on many factors, such as the name of the agent making the request or the history of requests, we now define a simple class of *threshold strategies*, where a player volunteers if and only if he has less than some fixed threshold \$$k$, and show in Section 4 that, in a precise sense, we can restrict to these strategies without loss of generality. Let K_i^t denote the amount of money agent i has at time t. Clearly $K_i^{t+1} = K_i^t$ unless agent i has a request satisfied, in which case $K_i^{t+1} = K_i^{t+1} - 1$ or agent i fulfills a request, in which case $K_i^{t+1} = K_i^{t+1} + 1$. Formally,

$$K_i^{t+1} = \begin{cases} K_i^t - 1 & \text{if } i = p^t, \sum_{j \neq i} V_j^t B_j^t > 0, \text{ and } K_i^t > 0 \\ K_i^t + 1 & \text{if } i = v^t \text{ and } K_{p^t}^t > 0 \\ K_i^t & \text{otherwise.} \end{cases}$$

The threshold strategy S_k is the strategy such that

$$V_i^t = \begin{cases} 1 & \text{if } K_{p^t}^t > 0 \text{ and } K_i^t < k \\ 0 & \text{otherwise.} \end{cases}$$

Thus, with S_k, a player volunteers iff he has less \$$k$ (and the requester has at least \$1, so that he can actually pay to have the request fulfilled).

The intuition behind using a threshold strategy is easy to explain: A rational agent with too little money will be concerned that he will run out and then want to make a request; on the other hand, a rational agent with plenty of money would not want to work, because by the time he has managed to spend all his money, the util will have less value than the present cost of working. By choosing an appropriate threshold, a rational agent can deal with both concerns. For the rest of the paper, we focus on threshold strategies.

3. THE GAME UNDER NONSTRATEGIC PLAY

Before we consider strategic play, we examine what happens in the system if everyone just plays the same strategy S_k. Our overall goal is to show that there is some distribution over money (i.e., the fraction of people with each amount of money) such that the system "converges" to this distribution in a sense to be made precise shortly.

Suppose that everyone plays S_k. For simplicity, assume that everyone has at most k dollars. We can make this assumption with essentially no loss of generality, since if someone has more than k dollars, he will just spend money until he has at most k dollars. After this point he will never acquire more than k. Thus, eventually the system will be in a state where everyone has at most \$$k$. If $M \geq kn$, no agent will ever be willing to work. Thus, for the purposes of this section we assume that $M < kn$.

Rather than thinking in terms of the total amount of money in the system, it will prove more useful to think in terms of the average amount of money each player has. Of course, the total amount of money in a system with n agents is M iff the average amount that each player has is $m = M/n$.

We want to understand how the distribution of money in the system changes over time if everyone uses a fixed threshold strategy S_k. To do this, it is helpful to think of the system as a Markov chain, where the state of the system is characterized by the amount of money each agent has. Note that if everyone has at most k dollars, then the amount of money that an agent has is an element of $\{0, \ldots, k\}$. If there are n agents, then the state of the system can be described by identifying how much money each agent has, so we can represent it by an element of $\mathcal{S}_{k,n} = \{0, \ldots, k\}^{\{1,\ldots,n\}}$. Since the total amount of money is constant, not all of these states can arise in the game. For example the state where each player has \$0 is impossible to reach in any game with money in the system. Let $m_{\mathcal{S}}(s) = \sum_{i \in \{1\ldots n\}} s(i)$ denote the total mount of money in the game at state s, where $s(i)$ is the number of dollars that agent i has in state s. We want to consider only those states where the total money in the system is $M = mn$, namely

$$\mathcal{S}_{k,n,m} = \{s \in \mathcal{S}_{k,n} \mid m_{\mathcal{S}}(s) = mn\}.$$

Under the assumption that all agents use strategy S_k, the evolution of the system can be treated as a Markov chain $\mathcal{M}_{k,n,m}$ over the state space $\mathcal{S}_{k,n,m}$. It is possible to move from a state s to a state t in one step iff there exist agents i and j such that $t(i) = s(i) + 1 \leq k$, $t(j) = s(j) - 1 \geq 0$, and $s(i') = t(i')$ for all i' other than i and j. Therefore, the probability of a transition from a state s to t is 0 unless there exist two agents i and j such that $s(i') = t(i')$ for all $i' \notin \{i, j\}$, $t(i) = s(i) + 1$, and $t(j) = s(j) - 1$. In this case the probability of transitioning from s to t is the probability of j being chosen to spend a dollar. It is straightforward to calculate this probability, but we do not actually need the transition probabilities for our results, so we do not write them down here.

Let Δ^k denote the set of probability distributions on $\{0, \ldots, k\}$. We can think of an element of Δ^k as describing the fraction of people with each amount of money. This is a useful way of looking at the system, since we typically are not interested in exactly who has each amount of money, but just the fraction of people that have each amount. As before, not all elements of Δ^k are possible, given our constraint that the total amount of money is M. Let Δ_m^k denote all distributions $d \in \Delta^k$ such that $\sum_{j=0}^k d(j)j = m$ (i.e., the expected value of d must be m). Given a state $s \in \mathcal{S}_{k,n,m}$, let $d^s \in \Delta_m^k$ denote the distribution of money in s. Our goal is to show that, if n is large, then there is a distribution $d^* \in \Delta_m^k$ such that, with high probability, the Markov chain $\mathcal{M}_{k,n,m}$ will almost always be in a state s such that d^s is close to d^*.

Thus, agents can base their decisions about what strategy to use on the assumption that they will be in such a state.

We can in fact completely characterize the distribution d^*. Given a distribution $d \in \Delta^k$, let

$$H(d) = -\sum_{\{j:d(j)\neq 0\}} d(j)\log(d(j))$$

denote the *entropy* of d. If Δ is a closed convex set of distributions, then it is well known that there is a unique distribution in Δ at which the entropy function takes its maximum value in Δ. Since Δ^k_m is easily seen to be a closed convex set of distributions, it follows that there is a unique distribution in Δ^k_m, which we denote $d^*_{k,m}$, whose entropy is greater than that of all other distributions in Δ^k_m. We now show that, for n sufficiently large, the Markov chain $\mathcal{M}_{k,n,m}$ is almost surely in a state s such that d^s is close to $d^*_{k,m}$. The statement is correct under a number of senses of "close". For definiteness, we consider the Euclidean distance. This presents us with a problem with our use of m as the exact average amount of money. We have no particular reason to require m to be an integer; a system with 2000 people and \$ 3000 seems perfectly reasonable. However, there are also m which correspond to non-integer total amounts of money. Thus for many values of n, the set of states with n people and mn total dollars can turn out to be $\emptyset$ for the trivial reason that mn is not an integer.

The problem boils down to the assumption that the of m to have the unwanted implication that the number of people must be such that mn is an integer. This problem appears in other contexts using maximum entropy (see [13]), however, we will adopt a different solution. Given $\epsilon > 0$, let $S_{k,n,m,\epsilon}$ denote the set of states s in $\mathcal{S}_{k,n,\lceil mn\rceil/n}$ such that $\sum_{j=0}^{k} |d^s(j) - d^*_{k,m}(j)|^2 < \epsilon$. Note that, rather than defining $S_{k,n,m,\epsilon}$, to be constrained to have a total of mn dollars, we will define it to be constrained to have a total of $\lceil mn\rceil$. This clearly solves the problem of the set being empty for the trivial reason that we wanted to have a non-integer total amount of money, but creates a new problem. Now the average amount of money for some particular value of n will be not be m but will instead be $\lceil mn\rceil/n$. Thus this will give slightly different constraints and lead to a slightly different maximum entropy distribution. However, the important point is that these differences are slight enough that they are smaller than the noise inherent in the system. Intuitively, the addition of a single dollar to a large system ought not to make any significant difference.

Given a Markov chain $\mathcal{M}$ over a state space $\mathcal{S}$ and $S \subseteq \mathcal{S}$, let $X_{t,s,S}$ be the random variable on executions of the Markov chain that is 1 if $\mathcal{M}$ is in a state of S at time t when started in state s, and 0 otherwise.

THEOREM 3.1. *For all $\epsilon > 0$, all k, and all m, there exists n_ϵ such that for all $n > n_\epsilon$ and all states $s \in \mathcal{S}_{k,n,m}$, there exists a time t^* (which may depend on k, n, m, and ϵ) such that for $t > t^*$, we have* $\Pr(X_{t,s,S_{k,n,m,\epsilon}}) > 1 - \epsilon$.

PROOF. The proof proceeds in three steps. First, using standard techniques, we show that the Markov chain $\mathcal{M}$ has a well-defined *limit distribution* π on $\mathcal{S}_{k,n,m}$; that is, a distribution π such that for all $s, s' \in \mathcal{S}_{k,n,m}$, $\lim_{t\to\infty}\Pr(X_{t,s,s'}) = \pi_{s'}$. We next show that π is uniform. Finally, we use a relatively straighrforward combinatorial argument (which is standard in statistical mechanics) to show that there is a *concentration phenomenon* around the maximum entropy distribution [19]. More precisely, we show that the fraction of states not in $S_{k,n,m,\epsilon}$ is bounded by $p(n)/e^{cn}$, where p is a polynomial. This fraction clearly goes to 0 as n gets large. Thus, for sufficiently large n, $\Pr(X_{t,s,S_{k,n,m,\epsilon}}) > 1 - \epsilon$. We leave the details to the appendix. □

We performed a number of experiments that show that the maximum entropy behavior described in Theorem 3.1 arises quickly for quite practical values of n and t. In the first experiment, we took $n = 1000$, $k = 5$, and $m = 2$. We wanted to examine time to convergence, so we started out in a state that was, intuitively, in an "extreme distribution", one as far from the maximum entropy distribution as possible: each agent had either \$0 or \$5. We averaged 10 runs starting in a state of this type. The results are shown in Figure 1. As we can see, after 2,000 steps, on average, the Euclidean distance from the average distribution of money to the maximum-entropy distribution is .008; after 3,000 steps, the distance is down to .001.

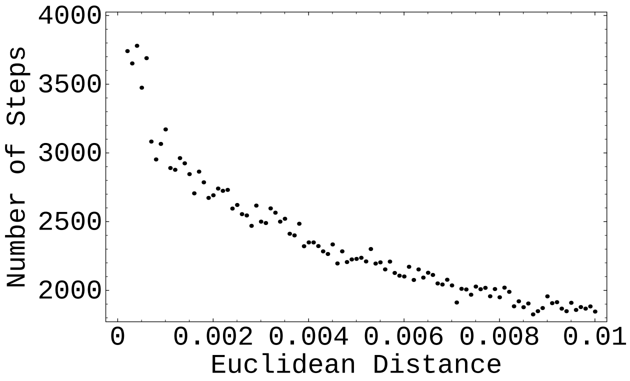

Figure 1: Distance from maximum-entropy distribution with 1000 agents.

With 1000 players, we can think of 1000 steps as taking one real time unit. Thus, with 1000 agents, we can get within .001 of the maximum-entropy distribution within 3 time units. We considered how the time to convergence depended on the number of players. In particular, we considered the time needed to get within .001 of the maximum-entropy distribution. As shown in Figure 2, it seems that we can get to within .001 of the maximum-entropy distribution within three time units (i.e., within $3n$ steps, where n is the number of players) no matter how many players we have. We do not have an analytic proof of this fact (although the graph is very suggestive). We suspect that the Markov chain that arises here is *rapidly mixing*; perhaps some results from the theory of rapidly mixing chains [23] may apply.

Finally, we considered how close the distribution stays to the maximum-entropy distribution. As we would expect, the more players there are, the smaller the variation from the maximum-entropy point. To simplify things, we started the system in a state whose distribution was very close to the maximum-entropy distribution and ran it for 10^6 steps, for various values of n. As Figure 3 shows, the system does not move far from the maximum-entropy distribution once it is close to it. For example, if $n = 5000$, the system is never more than distance .001 from the maximum-entropy

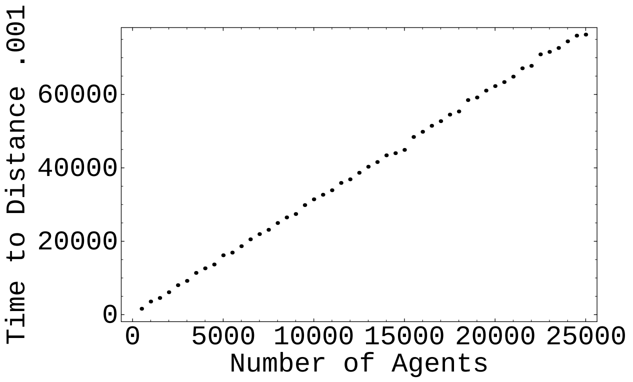

Figure 2: Average number of steps to get within .001 of the maximum-entropy distribution.

distribution; if $n = 25,000$, it is never more than .0002 from the maximum-entropy distribution.

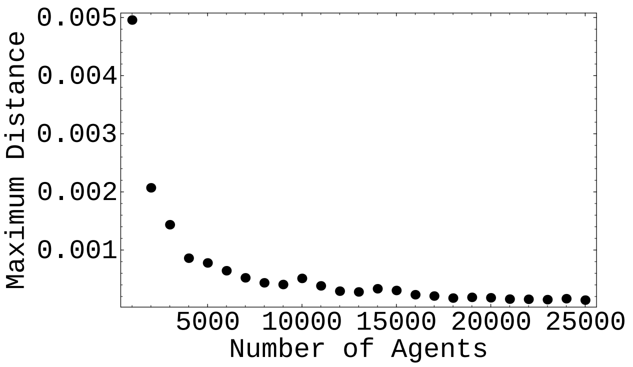

Figure 3: Maximum distance from maximum-entropy distribution over 10^6 timesteps.

4. THE GAME UNDER STRATEGIC PLAY

We have seen that the system is well behaved if the agents all follow a threshold strategy; we now want to show that there is a nontrivial Nash equilibrium where they do so (that is, a Nash equilibrium where all the agents use S_k for some $k > 0$.) This is not quite true, for three reasons. First, if δ is sufficiently small, then the prospect of getting utility in the future does not provide much incentive for the agents to work; all that matters is the present, and no one will volunteer. Thus, with sufficiently small δ, S_0 is the only equilibrium.

Second, we seem to be able to characterize behavior only when the distribution of money in the system is close to the limiting maximum-entropy distribution. Intuitively, this is not a major problem since, as the experiments discussed in Section 3 show, the system converges to the maximum-entropy distribution quite quickly. To deal with what the period of time before the system has converged to the maximum-entropy distribution, we consider a slightly different solution concept: not Nash equilibrium, but (multiplicative) ϵ-Nash equilibrium. For $\epsilon > 0$, we say that a strategy σ for player i is an *(multiplicative)ϵ-best response* to a joint strategy $\vec{\tau}$ for all players other than i if, for all strategies σ' for player i, $u_i(\vec{\tau}, \sigma') < u_i(\vec{\tau}, \sigma)(1+\epsilon)$. A joint strategy $\vec{\sigma} = (\sigma_1, \ldots, \sigma_n)$ is an *ϵ-Nash equilibrium* if σ_i is an ϵ-best response to $\vec{\sigma}_{-i}$ (where $\vec{\sigma}_{-i}$ is the joint strategy for all players other than i where player j plays σ_j). Intuitively, in an ϵ-Nash equilibrium, no no player can do more than ϵ better by choosing a different strategy.

Finally, we cannot in general find a Nash (or ϵ-Nash) equilibrium where everyone plays S_k for some natural number k. We need to have strategies that, intuitively, interpolate between S_k and S_{k+1}. More precisely, if $\gamma = k + \gamma'$, where k is a natural number and $0 \leq \gamma' < 1$, let S_γ be the strategy that performs S_k with probability $1 - \gamma'$ and S_{k+1} with probability γ. (Note that we are not considering arbitrary mixed threshold strategies here, but rather just mixing between adjacent strategies.) For the remainder of the paper, when we speak of threshold strategies, we will mean a strategy of the form S_γ for some arbitrary non-negative real value γ. We will show that, if δ is sufficiently close to 1, then there is a nontrivial Nash equilibrium where everyone plays a such generalized threshold S_γ, with $\gamma > 0$. Note for future reference that Theorem 3.1 applies to these generalized threshold strategies S_γ (the same proof goes through without change).

As a first step, we show that, for all $\gamma \geq 0$ there exists an $\epsilon^* > 0$ such that for all $\epsilon < \epsilon^*$ if everyone other than agent i is playing S_γ, then there is a threshold strategy $S_{\gamma'}$ that is an ϵ-best response for agent i. (Of course, it may not be the case that $\gamma = \gamma'$.)

THEOREM 4.1. *Fix m, $epsilon > 0$, threshold strategy S_γ, and an agent i. There exists $\delta^* < 1$ such that for all δ with $\delta^* < \delta < 1$, there exists a threshold strategy $S_{\gamma'}$ such that for all $\epsilon > 0$ there exists $n^* > 0$ such that if $n > n^*$ and $M/n = m$, and every agent other than i is playing S_γ in game $G(n, \delta)$, then $S_{\gamma'}$ is an ϵ-best response for agent i. Moreover, either γ' is unique and is an integer, or there exists an integer k' such that γ' can have any value in $[k', k'+1]$, and these are the only possible values for γ'.*

PROOF. (Sketch:) First assume that the system is characterized by the maximum-entropy distribution and the distribution is not affected by the choice of strategy of one agent. (We can always choose n sufficiently large that the probability of having a distribution of money far from the maximum entropy distribution, either from the influence of the strategy of player i or from the random fluctuations of the system, is arbitrarily small. This means that agents with a large enough δ, using a different strategy has an effect of at most ϵ on the payoffs.) Under these assumptions, the probability of i moving from one state (dollar amount) to another depends only on i's strategy (since we can take the probability that i will be chosen to make a request and the probability that i will be chosen to satisfy a request to be constant). Thus, from i's point of view, the system is a Markov decision process (MDP), and i needs to compute the optimal policy (strategy) for this MDP. It follows from standard results that there is an optimal policy that is a threshold policy. We give details in the appendix, where we also present the uniqueness argument. $\square$

We remark that there may be ϵ-best responses that are not threshold strategies. All that Theorem 4.1 shows is that, among ϵ-best responses, there is at least one that is a threshold strategy. Since we know that there is an ϵ-best response that is a threshold strategy, we can look for an ϵ-Nash equilibrium in the space of threshold strategies.

THEOREM 4.2. *For all m, $\epsilon > 0$, there exists $\delta^* < 1$ such that if $\delta^* < \delta < 1$, then there exists γ such that for all $\epsilon > 0$ there exists $n^* > 0$ such that if $n > n^*$, and $M/n = m$, then the joint strategy where all agents play S_γ is an ϵ-Nash equilibrium of the game $G(n, \delta)$.*

PROOF. Let $br(\delta, \gamma)$ be the minimal best-response threshold strategy if all the other agents are playing S_γ and the system is in its maximum-entropy distribution—that is, $br(\delta, \gamma) = \gamma'$ if S'_γ is the unique best-response threshold strategy (where, by Theorem 4.1, γ' is in fact an integer), and $br(\delta, \gamma) = k'$ if all the strategies $S_{\gamma'}$ for $\gamma' \in [k, k+1]$ are best responses. By Theorem 4.1, for fixed δ, $br(\delta, \cdot)$ is a step function; moreover, if we can join the "steps" by a vertical line, we get a best-response curve. It is easy to see that everywhere that this best-response curve crosses the diagonal $y = x$ defines a Nash equilibrium where all agents are using the same threshold strategy, under the assumption that the system characterized by the maximum-entropy distribution. As we have already observed, one such equilibrium occurs at 0. If there are only \$M in the system, we can restrict to threshold strategies S_γ where $\gamma \leq M + 1$. Since no one can have more than $\$M$, all strategies S_γ for $\gamma > M$ are equivalent to S_M; these are just the strategies where the agent always volunteers in response to a request made by someone who can pay. Clearly $br(\delta, M) \leq M$ for all δ, so the best response function is at or below the diagonal at M, for all δ. If $\lceil\gamma\rceil \leq M/n$, then it is easy to see that, after an initial period, every player will have at least $\lceil\gamma\rceil$ dollars, and so will be unwilling to work. Thus, if everyone but player i plays S_γ, then player i should play S_0; i.e., $br(\delta, \gamma) = 0$ for all $\delta > 0$. It is not hard to show that there exists $\delta^* > 0$ such that for all $\delta \geq \delta^*$, $br(\delta, k^*) \geq k^*$. Thus, either $br(\delta, k^*) \geq k^*$, in which case S_{k^*} gives the desired ϵ-Nash equilibrium, or $br(\delta, k^*) > k$, in which case, by continuity, the curve $br(\delta, \cdot)$ must cross the diagonal (since it is above the diagonal at k^* and below or at it at M). In either case, we get an ϵ-Nash equilibrium. $\square$

This argument also shows us that we cannot in general expect fixed points to be unique. Indeed, if $br(\delta, k^*) > k^*$, then there will be a fixed point in the interval $(0, k^*)$ and another fixed-point in the interval (k^*, M). Figure 4 illustrates the situation when $n = 1000$ and $M = 3000$.

Although Theorem 4.2 shows that there is always an ϵ-Nash equilibrium using threshold strategies, it does not rule it does not rule out the possibility that there may be other equilibria that do not involve threshold stratgies. It is even possible (although it seems unlikely) that some of these equilibria might be better.

the best responses to integer values are plotted (see figure Although the γ in the ϵ-Nash equilibrium strategy S_γ of Theorem 4.2 can in principle be an arbitrary non-negative real number, our simulations suggest that it is typically a natural number. We now try to explain why.

Our simulations suggest that the best-response function is nondecreasing (at least for a large part of its domain). As the threshold increases, more agents will want to work, so work will be harder to find. This means agents will try to

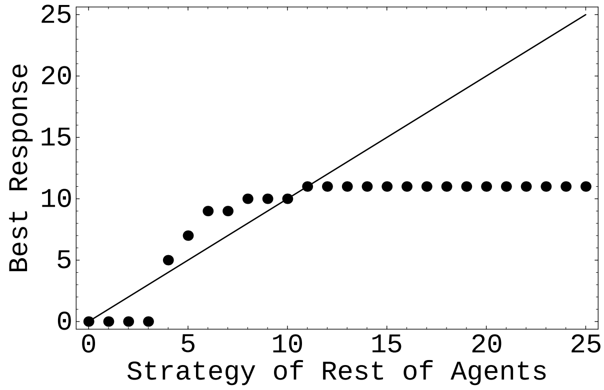

Figure 4: The best response function for $n = 1000$ and $M = 3000$.

save more money for a rainy day, and thus a best response would have a higher threshold. As we now show, if the best-response curve is indeed nondecreasing, then we are guaranteed to have a pure ϵ-Nash equilibrium strategy.

THEOREM 4.3. *Fix M, n, and δ. If the threshold strategy S_γ is an ϵ-Nash equilibrium, the best-response curve is nondecreasing in the interval $[\lfloor\gamma\rfloor, \gamma]$, and $br(\delta, \lfloor\gamma\rfloor) \geq \lfloor\gamma\rfloor$, then $S_{\lfloor\gamma\rfloor}$ is also an ϵ-Nash equilibrium.*

PROOF. If γ is an integer we are done, for then $\gamma = \lfloor\gamma\rfloor\gamma$. Otherwise, suppose by way of contradiction that $S_{\lfloor\gamma\rfloor}$ is not an ϵ-Nash equilibrium. Then $br(\delta, \lfloor\gamma\rfloor) > \lfloor\gamma\rfloor$ so $br(\delta, \lfloor\gamma\rfloor) \geq \lceil\gamma\rceil > \gamma > br(\delta, \gamma)$. The last inequality follows since, for S_γ to be an equilibrium, it must be the case that the best-response curve crosses the diagonal at γ. Since $br(\delta, \gamma)$ takes on only integer values, this can happen only if there exists an integer k' such that (a) all strategies $S_{\gamma'}$ with $\gamma' \in [k', k' + 1]$ are best responses to S_γ, and (b) $\gamma \in [k', k' + 1]$. But this means that $br(\delta, \gamma) = k' < \gamma$. It follows that $br(\delta, \gamma) < br(\delta, \lfloor\gamma\rfloor)$, which contradicts the assumption that br is nondecreasing. $\square$

As we have observed, we can always choose δ such that $br(\delta, \lceil M/n\rceil) > \lceil M/n\rceil$. Thus, if γ is the first place in the interval $[\lceil M/n\rceil, \infty]$ where the best-response curve crosses the diagonal, then we must have $br(\delta, \lfloor\gamma\rfloor) \geq \lfloor\gamma\rfloor$. Thus, if the best-response curve is nondecreasing, it follows the smallest $\gamma \geq \lfloor M/n\rfloor$ such that S_γ is an equilibrium must be an integer.

5. SOCIAL WELFARE AND SCALABILITY

Our theorems show that for each value of M and n, for sufficiently large δ, there is a nontrivial Nash equilibrium where all the agents use some threshold strategy $S_{\gamma(M,n)}$. From the point of view of the system designer, not all equilibria are equally good; we want an equilibrium where as few as possible agents have \$0 when they get a chance to make a request (so that they can pay for the request) and relatively few agents have more than the threshold amount of money (so that there are always plenty of agents to fulfill the request). There is a tension between these objectives. It is not hard to show that as the fraction of agents with \$0 increases in the maximum entropy distribution, the fraction of agents with the maximum amount of money decreases. Thus, our goal is to understand what the optimal amount of money should be in the system, given the number of agents. That is, we want to know the amount of money M that maximizes *efficiency*, i.e., the total expected utility if all the agents use $S_{\gamma(M,n)}$.[3]

We first observe that the most efficient equilibrium depends only on the ratio of M to n, not on the actual values of M and n.

THEOREM 5.1. *There exists n^* such that for all games $G(n_1, \delta)$ and $G(n_2, \delta)$ where $n_1, n_2 > n^*$, if $M_1/n_1 = M_2/n_2$, then $S_{\gamma(M_1,n_1)} = S_{\gamma(M_2,n_2)}$.*

PROOF. Fix $M/n = r$. Theorem 3.1 shows that the maximum-entropy distribution depends only on k and the ratio M/n, not on M and n separately. Thus, given r, for each choice of k, there is a unique maximum entropy distribution $d_{k,r}$. The best response $br(\delta, k)$ depends only on the distribution $d_{k,r}$, not M or n. Thus, the Nash equilibrium depends only on the ratio r. That is, for all choices of M and n such that n is sufficiently large (so that Theorem 3.1 applies) and $M/n = r$, the equilibrium strategies are the same. □

In light of Theorem 5.1, the system designer should ensure that there is enough money M in the system so that the ratio between M and n is optimal. We are currently exploring exactly what the optimal ratio is. As our very preliminary results for $\beta = 1$ show in Figure 5, the ratio appears to be nondecreasing in δ, which matches the intuition that we should provide more patient agents with the opportunity to save more money.

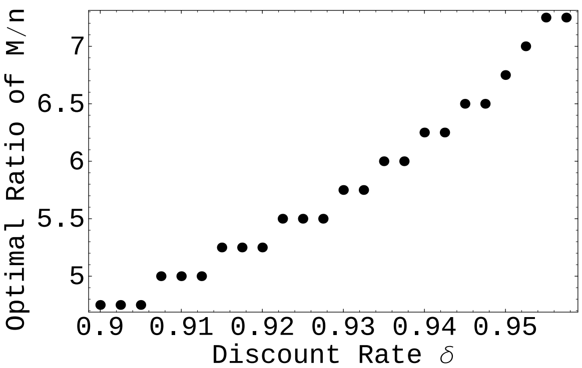

Figure 5: Optimal average amount of money to the nearest .25 for $\beta = 1$ as a function of the discount factor δ.

In practice, it may be easier for the designer to vary the price of fulfilling a request rather than injecting money in the system. This produces the same effect. For example, changing the cost of fulfilling a request from \$1 to \$2 is equivalent to halving the amount of money that each agent has. Similarly, halving the the cost of fulfilling a request is equivalent to doubling the amount of money that everyone has. With a fixed amount of money M, there is an optimal product nc of the number of agents and the cost c of fulfilling a request.

Theorem 5.1 also tells us how to deal with a dynamic pool of agents. Our system can handle newcomers relatively easily: simply allow them to join with no money. This gives existing agents no incentive to leave and rejoin as newcomers. We then change the price of fulfilling a request so that the optimal ratio is maintained. This method has the nice feature that it can be implemented in a distributed fashion; if all nodes in the system have a good estimate of n then they can all adjust prices automatically. (Alternatively, the number of agents in the system can be posted in a public place.) Approaches that rely on adjusting the amount of money may require expensive system-wide computations (see [31] for an example), and must be carefully tuned to avoid creating incentives for agents to manipulate the system by which this is done.

The realization that the cost of fulfilling a request can change can affect an agent's strategy. For example, if an agent expects the cost to increase, then he may want to defer volunteering to fulfill a request. However, if the number of agents in the system is always increasing, then the cost always decreases, so there is never any advantage in waiting to volunteer. On the other hand, there may be an advantage to delaying a request if the agent expects the cost of fulfilling a request to decrease. However, due to the discount factor, having a request satisfied tomorrow results in less utility than having a request satisfying today. If the rate of growth in the number of agents is sufficiently slow relative to δ, then there will not be an advantage to waiting. However, if an agent anticipates a short-term sudden increase in the number of agents, there may well be an advantage to waiting. (Related issues are discussed in [10].)

We ultimately hope to modify the mechanism so that the price of a job can be set endogenously within the system (as in real-world economies), with agents bidding for jobs rather than there being a fixed cost set externally. However, we have not yet explored the changes required to implement this change. Thus, for now, we assume that the cost is set as a function of the number of agents in the system (and that there is no possibility for agents to satisfy a request for less than the "official" cost or for requesters to offer to pay more than it).

6. SYBILS AND COLLUSION

In a naive sense, our system is essentially sybil-proof. To get d dollars, an agent's sybils together still have to perform d units of work. Moreover, since newcomers enter the system with \$0, there is no benefit to creating new agents simply to take advantage of an initial endowment. Nevertheless, there are some less direct ways that an agent could take advantage of sybils. First, by having more identities he will have a greater probability of getting chosen to make a request. It is easy to see that this will lead to the agent having higher total utility. However, this is just an artifact of our model. To make our system simple to analyze, we have assumed that request opportunities came uniformly at random. In practice, requests are made to satisfy a desire. Our model implicitly assumed that all agents are equally likely to have a desire at any particular time. Having sybils should not increase the need to have a request satisfied. Indeed, it would be reasonable to assume that sybils do not

[3]If there are multiple equilibria, we take $S_{\gamma(M,n)}$ to be the Nash equilibrium that has highest efficiency for fixed M and n.

make requests at all.

Second, having sybils makes it more likely that one of the sybils will be chosen to fulfill a request. This can allow a user to increase his utility by setting a lower threshold; that is, to use a strategy $S_{k'}$ where k' is smaller than the k used by the Nash equilibrium strategy. Intuitively, the need for money is not as critical if money is easier to obtain. Unlike the first concern, this seems like a real issue. It seems reasonable to believe that when people make a decision between a number of nodes to satisfy a request they do so at random, at least to some extent. Even if they look for advertised node features to help make this decision, sybils would allow a user to advertise a wide range of features.

Third, an agent can drive down the cost of fulfilling a request by introducing many sybils. Similarly, he could increase the cost (and thus the value of his money) by making a number of sybils leave the system. Conceivably he could alternate between these techniques to magnify the effects of work he does. We have not yet calculated the exact effect of this change (it interacts with the other two effects of having sybils that we have already noted). Given the number of sybils that would be needed to cause a real change in the perceived size of a large P2P network, the practicality of this attack depends heavily on how much sybils cost an attacker and what resources he has available.

The second point raised regarding sybils also applies to collusion if we allow money to be "loaned". If k agents collude, they can agree that, if one runs out of money, another in the group will loan him money. By pooling their money in this way, the k agents can again do better by setting a higher threshold. Note that the "loan" mechanism doesn't need to be built into the system; the agents can simply use a "fake" transaction to transfer the money. These appear to be the main avenues for collusive attacks, but we are still exploring this issue.

7. CONCLUSION

We have given a formal analysis of a scrip system and have shown that the existence of a Nash equilibrium where all agents use a threshold strategy. Moreover, we can compute efficiency of equilibrium strategy and optimize the price (or money supply) to maximize efficiency. Thus, our analysis provides a formal mechanisms for solving some important problems in implementing scrip systems. It tells us that with a fixed population of rational users, such systems are very unlikely to become unstable. Thus if this stability is common belief among the agents we would not expect inflation, bubbles, or crashes because of agent speculation. However, we cannot rule out the possibility that that agents may have other beliefs that will cause them to speculate. Our analysis also tells us how to scale the system to handle an influx of new users without introducing these problems: scale the money supply to keep the average amount of money constant (or equivalently adjust prices to achieve the same goal).

There are a number of theoretical issues that are still open, including a characterization of the number of equilibria—are there usually 2? In addition, we expect that one should be able to compute analytic estimates for the best response function and optimal pricing which would allow us to understand the relationship between pricing and various parameters in the model.

It would also be of great interest to extend our analysis to handle more realistic settings. We mention a few possible extensions here:

- We have assumed that the world is homogeneous in a number of ways, including request frequency, utility, and ability to satisfy requests. It would be interesting to examine how relaxing any of these assumptions would alter our results.

- We have assumed that there is no cost to an agent to be a member of the system. Suppose instead that we imposed a small cost simply for being present in the system to reflect the costs of routing messages and overlay maintainance. This modification could have a significant impact on sybil attacks.

- We have described a scrip system that works when there are no altruists and have shown that no system can work once there there are sufficiently many altruists. What happens between these extremes?

- One type of "irrational" behavior encountered with scrip systems is hoarding. There are some similarities between hoarding and altruistic behavior. While an altruist provide service for everyone, a hoarder will volunteer for all jobs (in order to get more money) and rarely request service (so as not to spend money). It would be interesting to investigate the extent to which our system is robust against hoarders. Clearly with too many hoarders, there may not be enough money remaining among the non-hoarders to guarantee that, typically, a non-hoarder would have enough money to satisfy a request.

- Finally, in P2P filesharing systems, there are overlapping communities of various sizes that are significantly more likely to be able to satisfy each other's requests. It would be interesting to investigate the effect of such communities on the equilibrium of our system.

There are also a number of implementation issues that would have to be resolved in a real system. For example, we need to worry about the possibility of agents counterfeiting money or lying about whether service was actually provided. Karma [31] provdes techniques for dealing with both of these issues and a number of others, but some of Karma's implementation decisions point to problems for our model. For example, it is prohibitively expensive to ensure that bank account balances can never go negative, a fact that our model does not capture. Another example is that Karma has nodes serve as bookkeepers, keeping track of other nodes' account balances. Like maintaining a presence in the network, bookkeeping imposes a cost on the node, but unlike maintaining a presence, it is a cost that can be avoided with no obvious harm. Karma suggests several ways to incentivize nodes to perform these duties. We have not investigated whether these mechanisms can be incorporated without disturbing our equilibrium.

8. ACKNOWLEDGEMENTS

We would like to thank Emin Gün Sirer, Shane Henderson, Jon Kleinberg, and three anonymous referees for helpful suggestions. EF, IK and JH are supported in part by NSF under grant ITR-0325453. JH is also supported in part by NSF under grants CTC-0208535 and IIS-0534064, by ONR under grant N00014-01-10-511, by the DoD Multidisciplinary University Research Initiative (MURI) program administered by the ONR under grants N00014-01-1-0795 and N00014-04-1-0725, and by AFOSR under grant F49620-02-1-0101.

APPENDIX

A. PROOFS OF THEOREMS

In this section we provide the details of the proofs of Theorems 3.1, 4.1, and 4.2. We restate the theorems for the readers' convenience.

THEOREM 3.1. *For all $\epsilon > 0$, all k, and all m, there exists n_ϵ such that for all $n > n_\epsilon$ and all states $s \in \mathcal{S}_{k,n,m}$, there exists a time t^* (which may depend on k, n, m, and ϵ) such that for $t > t^*$, we have $\Pr(X_{t,s,S_{k,n,m,\epsilon}}) > 1 - \epsilon$.*

PROOF. Given a Markov chain $\mathcal{M}$ over a state space $\mathcal{S}$ and state $s \in \mathcal{S}$, let $X_{t,s,s'}$ be the random variable on executions of the Markov chain that is 1 if $\mathcal{M}$ is in state s' at time t (given that $\mathcal{M}$ is in state s at time 0) and is 0 otherwise. Then $\lim_{t\to\infty} \Pr(X_{t,s,s'} = 1)$ is the limit probability of being in state s' given that the Markov chain starts in state s. This limit may not in general exist. However, there are well known conditions under which the limit exists, and is independent of the initial state s. A Markov chain is said to be *irreducible* if every state is reachable from every other state; it is *aperiodic* if, for every state s, there exist two cycles from s to itself such that the gcd of their lengths is 1.

THEOREM A.1. *[28, Corollary to Proposition 2.13.5] If $\mathcal{M}$ is a finite, irreducible, and aperiodic Markov chain over state space $\mathcal{S}$, then there exists a d such that, for all $s \in \mathcal{S}$, we have $\lim_{t\to\infty} \Pr(X_{t,s,s'} = 1) = d$.*

Thus, if we can show that $\mathcal{M}_{k,n,m}$ is finite, irreducible, and aperiodic, then the limit distribution exists and is independent of the start state. We now show this.

LEMMA A.1. *If there are at least three agents, then $\mathcal{M}_k$ is finite, irreducible, and aperiodic and therefore has a limit distribution π.*

PROOF. $\mathcal{M}_{k,n,m}$ is clearly finite since $\mathcal{S}_{k,n,m}$ is finite. We prove that it is irreducible by showing that state s is reachable from state s' by induction on the distance $\sum_{i=1}^{n} |s(i) - s'(i)|$ (i.e., the sum of the absolute differences the amount of money each person has). If the distance is 0 then $s = s'$ so we are done. Suppose that $d > 0$ and all pairs of states that are less that d apart are reachable from each other. Consider a pair of states s and s' such that the distance from s to s' is d. Since $d > 0$ and the total amount of money is the same in all states, there must exist i_1 and i_2 such that $s(i_1) > s'(i_1)$ and $s(i_2) < s'(i_2)$. We must have $s'(i_1) < k+1$ and $s'(i_2) > 0$, so it is possible for i_2 to do work for i_1 and get paid a dollar. The resulting state s'' is of distance $d-2$ from s. By the induction hypothesis, s'' is reachable from s. Since s' is clearly reachable from s'', s' is reachable from s. This completes the induction proof. Since the distance between states is finite (and, indeed, is bounded by $(k+1)^n$), the desired result follows.

Finally, we must show that $\mathcal{M}_{k,n,m}$ is aperiodic. Suppose s is a state such that there exist three agents p_1, p_2, and p_3 where p_1 has more than 0 dollars and p_2 and p_3 have less than k dollars. Clearly there is a cycle of length 2 from s to itself: p_2 does work for p_1 and then p_2 does work for p_1. There is also a cycle of length 3: p_2 does work for p_1, p_3 does work for p_2, then p_1 does work for p_3. If there do not exist such agents p_1, p_2, and p_3 in s, then it must be the case that, in state s, all but one of the agents in s has \$$k$. (The remaining agent cannot have \$$k$, given our assumption that the total money supply is $< kn$). But, in that case, suppose that p_1 and p_2 have \$$k$ and p_3 has less than \$$k$. Then we can have p_3 doing work for p_2, then p_2 doing work for p_3, giving us a cycle of length 2; and p_3 doing work for p_2, p_2 doing work for p_1, and p_1 doing work for p_3, giving us a cycle of length 3. □

The next step is to show that the limit distribution π of $\mathcal{M}_{k,n,m}$ is in fact uniform. The following well-known result gives a sufficient condition for the limit distribution

THEOREM A.2. *[?, Theorem xxx] If $\Pr(X_{1,s,s'} = 1) = \Pr(X_{1,s',s} = 1)$ for all states s and s', then the limit distribution is uniform.*

As the following lemma shows, our chain satisfies this condition.

LEMMA A.2. *For all states s and $s' \in \mathcal{S}_{k,n,m}$, $\Pr(X_{1,s,s'} = 1) = \Pr(X_{1,s',s} = 1)$.*

PROOF. This says that the probability of transitioning from s to s' is the same as transitioning from s' to s. If the states are not adjacent, then both probabilities are zero and we are done. Otherwise there must be some i and j such that i being selected to make a request and j being selected to satisfy it results in a transition from s to s'. This means that j being chosen to make a request and i being chosen to satisfy it results in a transition from s' to s. Thus we simply need to check that these probabilities are the same.

The person to make a request is chosen at random, so the probability of the required player being chosen is always the same: $1/n$. Therefore we just need to check the the probability of j being chosen to satisfy a request in state s is the same as the probability of i being chosen to satisfy a request in state s'. Let V_s be the set of agents in s who are willing to volunteer to fulfill i's request; that is, $V_s = \{i' \neq i : s(i') < k\}$; similarly, let $V_{s'}$ be the set of agents in s' who are willing to volunteer to fulfill j's request. By assumption, $j \in V_s$. Moreover, since i has one less dollar in s' than s, $i \in V_{s'}$. Since all agents other than i and j have the same amount of money in s and s', it is easy to see that $V_{s'} = V_s \cup \{i\} - \{j\}$, and thus $|V_s| = V_{s'}$. It easily follows that the probability of j being chosen to satisfy a request in s is the same as the probability of i being chosen to satisfy a request in s'. □

The next last step is to show that there is a concentration phenomenon, so that the system is in a state in $S_{k,n,m,\epsilon}$ with high probability. In order to do this, we bound the fraction of states not in $S_{k,n,m,\epsilon}$ as n grows large.

LEMMA A.3. *The fraction of states not in $S_{k,n,m,\epsilon}$ is bounded by $p(n)/e^{cn}$, where p is a polynomial and c is a constant dependant on ϵ*

PROOF. The proof is essentially that of Theorem 3.13 in [13] (though applied to a different type of constraint). Because Δ_m^k is a closed, convex set, it has a unique point that maximizes entropy [13]. Let d^* be this point, and let ρ be the entropy at the point d^*. Consider the set D of points in Δ_m^k whose distance from d^* is greater than or equal to ϵ. Since this is a closed set, the entropy function takes on a maximum in D; suppose this maximum is ρ_l. We want to show that there is a state $s \in S_{k,n,m,\epsilon}$ such that $H(d^s) > \rho_l$. Take ρ_h to be some value in (ρ_l, ρ) (say $\rho_l + \rho/2$). By the continuity of the entropy function, there exists some $\epsilon' > 0$ such that if $|d^s - p| < \epsilon'$, then $H(d^s) \geq \rho_h$. Thus it suffices to show that for any ϵ', there exists an n^* such that, for all $n > n^*$, $S_{k,n,m,\epsilon'}$ is nonempty.

Consider the distribution d_n with $d(i) = \lfloor nd^*(i) \rfloor / n$. Due to taking the floor to guarantee integers, the total amount of money in a state that has this distribution may be less than $\lceil mn \rceil$ dollars. However, it can only differ by at most $O(k^2)$ dollars because the rounding of the ith coordinate loses at most i dollars and one more dollar may be lost if $mn < \lceil mn \rceil$. To find a distribution corresponding to a state with exactly $\lceil mn \rceil$ dollars, we will create a d'_n by subtracting $1/n$ from some $d(i)$ and adding $1/n$ to $d(i+1)$ for each dollar that needs to be added (this is equivalent to picking a person to whom to give an extra dollar). Since each coordinate of d'_n is a multiple of $1/n$, and the total amount of money represented by d'_n is $\lceil mn \rceil$, there is a state s_n in $S_{k,n,m}$ where the fraction of people with i dollars is $d'_n(i)$, for $i = 1, \ldots, \lfloor m/n \rfloor$.

We now want to show that for n sufficiently large, $s_n \in S_{k,n,m,\epsilon'}$, that is, that $|d'_n - d^*| < \epsilon'$. Notice that we defined d'_n from d^* by doing some rounding and then applying $O(k^2)$ changes of $1/n$. Thus $|d'_n(i) - d^*(i)|$ is $O(k^2/n)$. Since k is a constant independent of n, this means that the difference in each dimesion is $O(1/n)$ and the total distance is $O(1/\sqrt{n})$. Thus for sufficently large n, the distance will always be less than ϵ'.

To complete the proof, we use the fact shown in the proof of Lemma 3.11 of [13] that for any $d \in \Delta_m^k$

$$\frac{1}{f(n)} e^{nH(d)} \leq \binom{n}{nd(0), \ldots, nd(k)} \leq g(n) e^{nH(d)}$$

where f and g are polynomial in n. We know that there is some point at which the entropy is ρ_h. There are at least as many total states as there are states at that point, so the total number of states is at least

$$\frac{1}{f(n)} e^{nH(d)} = \frac{1}{f(n)} e^{n\rho_H}$$

The number of states at each point of distance at least ϵ is at most

$$g(n) e^{nH(d)} = g(n) e^{n\rho_l}.$$

There are at most $(n+1)^{k+1}$ points, so the total number of "bad" states is at most

$$h(n) e^{n\rho_l}$$

where $h(n) = g(n)(n+1)^{k+1}$ which is still polynomial in n. Thus the fraction of "bad" states is at most

$$\frac{h(n) e^{n\rho_l}}{\frac{1}{f(n)} e^{n\rho_h}} = \frac{f(n)h(n)}{e^{n(\rho_h - \rho_l)}}$$

$\square$

This is a ratio of a polynomial to an exponential, so it goes to zero in the limit. Thus far we have shown that the Markov chain for our system has a limit distribution in which the probability distribution on states is uniform. We have also shown that, in the limit as $n \to \infty$, the the fraction of states s where d^s is not "close" to d^* is 0. Thus, if the probability distribution on states is exactly the limit distribtuion, then the limit probability of being in a state at distance greater than ϵ from d^* is 0 for any $\epsilon < 0$. After any fixed number of rounds, the probability of being in a given state is not exactly the probability given by the limit distribution, but as the number of rounds increases, the distributions get arbitrarily close. This will be enough for us to complete the proof of Theorem 3.1

By Lemma A.2, the limit distribution of $\mathcal{M}_{k,n,m}$ is uniform independent of the start state. Thus, after a sufficient amount of time, the distribution over states will be arbitrarily close to uniform. In particular, for any c, there exists a t_c such that after t rounds the probability of being in any particular state is within a factor of c of uniform. By Lemma A.3, the fraction of states not is $S_{k,n,m,\epsilon}$ is bounded by $p(n)/e^{c'n}$. The probability of being in one of these states is at most a factor of c beyond the uniform probabiility. Thus the probability of being in one of these states is bounded by $cp(n)/e^{c'n}$, which is still the ratio of a polynomial and an exponential. Thus we simply need to choose a c and n_ϵ such that $cp(n_\epsilon)/e^{c'n_\epsilon} < \epsilon$ and take $t^* > t_c$.

THEOREM 4.1. *Fix a strategy S_γ and an agent i. There exists $\delta^* < 1$ and n^* such that if $\delta > \delta^*$, $n > n^*$, and every agent other than i is playing S_γ in game $G(n, \delta)$, then there is an integer k' such that the best response for agent i is $S_{k'}$. Either k' is unique (that is, there is a unique best response that is also a threshold strategy), or there exists an integer k'' such that $S_{\gamma'}$ is a best response for agent i for all γ' in the interval $[k'', k''+1]$ (and these are the only best responses among threshold strategies).*

We know from Theorem 3.1 that, if all agents play a particular threshold strategy, after a sufficent amount of time the system will be close to the maximum entropy distribution for that strategy. Suppose to start with that the distribution in the system is the maximimum entropy distribution, regardless of the strategy of a single agent. Then the number of other agents volunteering will be constant in every round. Thus, from the perspective of a single agent, p_u (the probability he will satisfy a request given that he volunteers) and p_d (the probability that he will have a request satisfied) are time-independent constants. Therefore the strategy for an agent is a policy for a relatively simple Markov descision process. There are states 0 to M representing the amounts of money he could have. In each round he chooses either to volunteer (action 1) or not volunteer (action 0). He then transitions to a new state and gets a payoff based on the results of that action in our game. Conditions for MDPs having optimal threshold policies have been extensively studied (see [26]).

We define some notation that will be helpful in proving the result.

- $S = \{0, \ldots, M\}$ is the set of possible states for the MDP (i.e., the possible amounts of money);

- $A = \{0, 1\}$ is the set of possible actions, where 0 denotes not volunteering and 1 denotes volunteering;
- $r(s, a)$ The (immediate) expected reward for performing action a in state s;
- $p(j \mid s, a)$ The probability of being in state j after performing action a in state s;
- $q(k \mid s, a) = \sum_{j=k}^{\infty} p(j \mid s, a)$ the probability of being in a state at least k after performing action a in state s;
- $u^*(s)$ The expected utility of being in state s (assuming the optimal strategy is being followed);
- $u(s, a) = r(s, a) + \sum_{j=0}^{\infty} p(j \mid s, a)u^*(j)$. The expected utility for performing action a in state s (given that the optimal strategy will be followed after this action);

In proving that the decision problem has an optimal threshold policy, we will want to make use if the fact that $u^*(s)$ is monotone increasing. A well known sufficent condition for this is

Theorem A.3. *[26, Proposition 4.7.3] $u^*(s)$ is monotone increasing if*

(a) $r(s, a)$ is nondecreasing in s for all a;

(b) $q(k \mid s, a)$ is nondecreasing in s for all $k \in S$ and $a \in A$;

Thus, to show that u^* is monotone, it suffices to show that the conditions of Theorem A.3 hold. This is straightforward. Since the cost of volunteering is constant and the payoff for getting a request satisfied is constant, $r(s, a)$ is independent of s (except that it is less in state 0, which is consistent) so (a) holds. Clearly the amount of money you have after deciding wheter to volunteer is nondecreasing in your current amount of money. Since volunteering earns you money, and volunteering is 1 while not volunteering is a 0, the probability of having at least k dollars is also not decreasing in a. Thus (b) also holds and we know that u^* is monotone.

Lemma A.4. *This MDP has an optimal threshold policy*

Proof. Suppose the optimal action for state $i + 1$ is 0. Then for state i, if we always choose action 1 (volunteering) until we get to state $i - 1$, we may cycle between states i and $i + 1$ until eventually we end up in state $i - 1$ and then follow the optimal strategy. We can write the expected utility of this strategy $E[U_c + \delta^{t_c} + \delta^{t_c}u^*(i - 1)]$ where U_c is the utility received while cycling, and t_c is the time until the cycle ends by spending a dollar in state i. Similarly we can write the expected utility of doing the same with action 0 as $E[\delta^{t_d} + \delta^{t_d}u^*i - 1]$ where t_d is the time to spend a dollar. The difference in expected values is therefore $E[U_c + \delta^{t_c} - \delta^{t_d} + (\delta^{t_c} - \delta^{t_d})u^*i - 1]$. Note that U_c, t_c, and t_d are constants independant of s and that $\delta^{t_c} - \delta^{t_d}$ is negative. Thus volunteering will be superior whenever $u^*(i - 1)$ is smaller than some threshold u. The only possible action in state M is 0. Therefore the action 0 will be optimal in every state down from the M until the first time $u^*(i - 1) < u$, at which point action 1 will be optimal for that state. So we know that in state i, the optimal action will be 1. Now we want to determine the optimal action for state $i - 1$. In state $i - 1$, if we choose action 0, the expected utility is independent of the expected utility of state i. If we choose action 1, it does depend on it. By the same argument we used for state i, if we decided to always use action 0 in state i (instead of the optimal action 1), we would still prefer to volunteer in state $i - 1$ because $u^*(i - 2) \leq u^*(i - 1) < u$ by A.3. Thus performing the optimal action of 1 in state i can only improve the utility of choosing action 1 in state $i - 1$ while leaving the utility of action 0 unchanged. Thus action 1 is optimal in state $i - 1$. By induction, we know that 1 will be optimal for every state from i down to 1. State 0 is a special case. However note that for 1 to be optimal in state i, $U_c > 0$. Thus the value of action 1 in state 0 will be > 0 while the value of action 0 will be 0, so 1 is optimal in state 0 as well. □

We have shown that, if we assume that the distribution of money is fixed at exactly the maximum entropy distribution, there is a best response that is a threshold strategy. However, the system will in general not be at exactly the maximum entropy distribution. Still, as Theorem 3.1 shows, after a sufficent amount of time the system is very unlikely to be very far away. Intuitively, if an agent is patient enough then all that will matter is the future state of the system, when it will be very close to the maximum entropy distribution. Similarly, we might worry that the strategy of the agent will alter the maximum entropy distribution so that the MDP we have examined is no longer valid. Intuitively, if there are enough people in the system, any one will be unable to alter the distribution. We formalize these intuitions in the following lemma. [[Put Lemma Here]]

We have observed that, when the system is stable, threshold strategies are a best response. We have also noted that, after some amount of time of agents playing threshold strategies, the system will be stable with very high probability. In section 3, we conjectured that the Markov chain is rapidly mixing, so that we expect these periods to be relatively brief. As a result, it seems reasonable to not worry about the incentives of agents during these transient time periods, especially as it is not clear how to determine what would be a better strategy for them. We can also make this intuitive appeal more precise using a slightly non-standard notion of an ϵ-Nash. The basic idea behind this is that, we can get a finite bound on to maximum utility possible for the time until a stable state is reached, and if δ is large enough, that becomes an insignificant portion of the utility over the course of the infinite game. Thus our notion of an ϵ best response asserts that an agent cannot do better by a multiplicative factor of ϵ rather than an additive one as is more commonly meant.

Lemma A.5. *Let S_k be a best response as in Theorem 4.2. For every ϵ, there exists a δ^* such that if $\delta > \delta^*$ then S_k is an ϵ best response to S_γ.*

Proof. We know from Theorem 3.1, that there exists some t such that the probability of being in a state far from the maximum entropy distribution is at more is arbitrarily small. Since the value of a strategy is finite, we can choose a t such that the expected loss from playing a suboptimal strategy in this event is at more $\epsilon/2$. The maximum payoff possible is a payoff of 1 each round. Thus until time t, an upper bound on the maximum payoff is $\sum_{i=0}^{t} \delta^i$, or if we desire to make it independent of δ we can also simply use t and ignore the discouting. Then we simply need to choose a δ^* such that t is at most an $\epsilon/2$ fraction of the expected future value of playing S_k in a state near the maximum entropy distribution. Note that this value can be arbitrarily large because in the limit as *delta* goes to 1, the value of any non-zero threshold strategy is infinite. $\square$

We now know that there is a best response that is a threshold strategy. However, there may be multiple threshold strategies that are best responses. Intuitively, an agent with 5 dollars might be ambivalent about paying α to have 6 dollars, in which case S_5 and S_6 would both be optimal. Similarly, any mixture of them would also be optimal. The following lemma shows that this is the only way for the best response to be non-unique.

LEMMA A.6. *The best response is either unique or consists of all γ in the interval $[k', k'+1]$.*

PROOF. Consider the strategies S_k and S_{k+1}. Starting with the same amount of money, an agent using S_k will behave in the same way as an agent using S_{k+1}, except in a state where the agents have \$$k$ and have the opportunity to work. In this case the agent using S_{k+1} will volunteer, but the agent using S_k will not. Afterwards, the two agents will behave identically, but the agent using S_k will have a dollar less than the agent using S_{k+1}. This continues to be the case until the agent using S_k runs out of money and is offered an opportunity to spend. At this point, the agent using S_k will not be able to take the opportunity, but the agent using S_{k+1} will be able to take it. Afterwards, both agents have \$0, so are in the same state again. Thus, S_{k+1} is the better strategy iff α (the cost of working now) $\leq E[\delta^{X_k}]$ (the expected future payoff), where X_k is a random variable that describes the number of steps that it takes an agent playing S_k to run out of money, given that it has \$$k+1$. In order to run out money starting with $k+2$ dollars, some number of rounds must pass after which the agent has $k+1$ dollars. Therefore $E[\delta^{X_k}]$ is strict monotone decreasing (because $\delta < 1$ and the value of the exponent is increasing). Thus it can't be the case that S_k and $S_{k'}$ are equally good but for some $k < k'' < k'$ $S_{k''}$ is worse. $\square$

THEOREM 4.2. *For all M, there exists $\delta^* < 1$ and n^* such that if $\delta > \delta^*$ and $n > n^*$, there exists a Nash equilibrium in the game $G(n, \delta)$ where all agents play S_γ for some integer $\gamma > 0$.*

PROOF. It follows easily from the proof Theorem 4.1 that if $br(\delta, \gamma)$ is the minimal best response threshold strategy if all the other agents are playing S_γ and the discount factor is δ then, for fixed δ, $br(\delta, \cdot)$ is a step function. It also follows from the theorem that if there are two best responses, then a mixture of them is also a best response. Therefore, if we can join the "steps" by a vertical line, we get a best-response curve. It is easy to see that everywhere that this best-response curve crosses the diagonal $y = x$ defines a Nash equilibrium where all agents are using the same threshold strategy. As we have already observed, one such equilibrium occurs at 0. To show that there is another, nontrivial equilibrium, we need to show that there is another crossing. Since we have shown that the best response curve is continuous, it is sufficient to show that there is a point one the curve above the line and a point on the curve below the line. If there are only \$M in the system, we can restrict to threshold strategies S_k where $k \leq M+1$. Since no one can have more than \$M, all strategies S_k for $k > M$ are equivalent to S_M; these are just the strategies where the agent always volunteers in response to request made by someone who can pay. Clearly $br(\delta, S_M) \leq M$ for all δ, so the best response function has a point at or below the line. If $k \leq M/n$, every player will have at least k dollars and so will be unwilling to work and the best response is just 0. Consider k^*, the smallest k such that $k > M/n$. We need to show that if an agent is sufficiently patient, the curve will be above the line at this point. That is, for k^* there exists a δ^* such that for all $\delta \geq \delta^*$, $br(\delta, k^*) \geq k^*$. In the proof of lemma A.6, we showed that S_{k+1} is a better strategy than S_k if $\alpha \leq E[\delta^{X_k}]$. It follows that the optimal strategy is the minimal one such that $\alpha \geq E[\delta_k^X]$. Each X_k is independent of δ. Thus by taking a large enough δ, the expectation can be made larger than α for any chosen k. Thus take δ^* to be any value such that $\alpha \leq E[(\delta^*)^{X_{k^*}}]$. This gives us a point above the line.

We have shown that there is a point above the line and then a point below the line for sufficiently large δ. It follows by continuity that if $\delta \geq \delta^*$, there must be some γ such that $br(\delta, \gamma) = \gamma$. This is the desired Nash equilibrium. $\square$

Material on Continuity

We want to show that the best response function is continuous. Let $U_{k'}(\gamma)$ be the expected utility of playing $S_{k'}$ given that everyone else is playing S_γ and the agent currently has zero dollars (this last is just for definiteness in the definition, if a strategy is better with zero dollars, it is better regardless of the amount of money). We will show that these utility functions are continuous in γ and as a result the best response function is continuous.

LEMMA A.7. *$U_{k'}(\gamma)$ is continous in γ.*

PROOF. Let p_u be the probability of the player getting picked to do work and $p_d = 1/n$ be the probability of the playet getting picked to spend money (p_u is just a function of γ, so it has the nice property that it is independent both of the strategy chosen by the player and state that the player is currently in). Then we can give the following system of equations for the expected payoffs (P_i is the expected payoff in state i, with $U_{k'}(\gamma) = P_0$):

$$P_0 = p_u(-\alpha + \delta P_1) + (1 - p_u)\delta P_0$$

$$P_i = p_u(-\alpha + \delta P_{i+1}) + p_d(1 + \delta P_{i-1}) + (1 - p_u - p_d)\delta P_i$$

$$P_{k+1} = p_d(1 + \delta P_k) + (1 - p_d)\delta P_{k+1}$$

Since α, δ, and p_d are constants, the coefficients are a continuous function of p_u. Thus the value of P_0 is continuous in γ if p_u is. $p_u = 1/(n - n_0)$, where n_0 is the number of players with zero dollars. The total number of players with zero dollars is the sum of the players playing $S_{\lfloor \gamma \rfloor}$ with zero dollars and the number of players playing $S_{\lceil \gamma \rceil}$ with zero dollars. From Theorem 3.1, we know that these numbers will be given by the maximum entropy distribution. Using lagrange multipliers, we can solve for the exact value:

$$n_0 = n\left(\frac{(1-\gamma-\lfloor\gamma\rfloor)}{\sum_{i=0}^{\lfloor\gamma\rfloor} e^{-\lambda_1 i}} + \frac{\gamma-\lfloor\gamma\rfloor}{\sum_{i=0}^{\lfloor\gamma\rfloor+1} e^{-\lambda_2 i\gamma\rfloor}}\right)$$

Here λ_1 and λ_2 are constants chosen to satisfy the constraint on the amount of money. We are still missing a piece to show that the λs are continuous in γ, but once that is done this proof is complete. □

LEMMA A.8. *The best response curve is continous*

PROOF. The best response to γ is $argmax_{k'}U_{k'}(\gamma)$, so Lemma A.7 tells us we are taking an argmax of continous functions. We also know from Theorem 4.1 that this argmax is either unique or consists of two consecutive values of k' (and all the γs in between them). Suppose the best response discontinuously switches from k' to k''. Then $U_{k'}(\gamma)$ goes from strictly greater than $U_{k''}(\gamma)$ to strictly less than it without crossing, which contradicts their continuity. □

B. REFERENCES

[1] E. Adar and B. A. Huberman. Free riding on Gnutella. *First Monday*, 5(10), 2000.

[2] K. G. Anagnostakis and M. Greenwald. Exchange-based incentive mechanisms for peer-to-peer file sharing. In *International Conference on Distributed Computing Systems (ICDCS)*, pages 524–533, 2004.

[3] BitTorrent Inc. BitTorrent web site. http://www.bittorent.com.

[4] A. Cheng and E. Friedman. Sybilproof reputation mechanisms. In *Workshop on Economics of Peer-to-Peer Systems (P2PECON)*, pages 128–132, 2005.

[5] Cornell Information Technologies. Cornell's ccommodity internet usage statistics. http://www.cit.cornell.edu/computer/students/ bandwidth/charts.html.

[6] J. R. Douceur. The sybil attack. In *International Workshop on Peer-to-Peer Systems (IPTPS)*, pages 251–260, 2002.

[7] G. Ellison. Cooperation in the prisoner's dilemma with anonymous random matching. *Review of Economic Studies*, 61:567–588, 1994.

[8] eMule Project. eMule web site. http://www.emule-project.net/.

[9] M. Feldman, K. Lai, I. Stoica, and J. Chuang. Robust incentive techniques for peer-to-peer networks. In *ACM Conference on Electronic Commerce (EC)*, pages 102–111, 2004.

[10] E. J. Friedman and D. C. Parkes. Pricing wifi at starbucks: issues in online mechanism design. In *EC '03: Proceedings of the 4th ACM Conference on Electronic Commerce*, pages 240–241. ACM Press, 2003.

[11] E. J. Friedman and P. Resnick. The social cost of cheap pseudonyms. *Journal of Economics and Management Strategy*, 10(2):173–199, 2001.

[12] P. Golle, K. Leyton-Brown, and I. Mironov. Incentives for sharing in peer-to-peer networks. In *ACM Conference on Electronic Commerce*, pages 264–267, 2001.

[13] A. J. Grove, J. Y. Halpern, and D. Koller. Random worlds and maximum entropy. *J. Artif. Intell. Res. (JAIR)*, 2:33–88, 1994.

[14] R. Guha, R. Kumar, P. Raghavan, and A. Tomkins. Propagation of trust and distrust. In *Conference on the World Wide Web(WWW)*, pages 403–412, 2004.

[15] M. Gupta, P. Judge, and M. H. Ammar. A reputation system for peer-to-peer networks. In *Network and Operating System Support for Digital Audio and Video(NOSSDAV)*, pages 144–152, 2003.

[16] Z. Gyongi, P. Berkhin, H. Garcia-Molina, and J. Pedersen. Link spam detection based on mass estimation. Technical report, Stanford University, 2005.

[17] D. Hughes, G. Coulson, and J. Walkerdine. Free riding on gnutella revisited: The bell tolls? *IEEE Distributed Systems Online*, 6(6), 2005.

[18] J. Ioannidis, S. Ioannidis, A. D. Keromytis, and V. Prevelakis. Fileteller: Paying and getting paid for file storage. In *Financial Cryptography*, pages 282–299, 2002.

[19] E. T. Jaynes. Where do we stand on maximum entropy? In R. D. Levine and M. Tribus, editors, *The Maximum Entropy Formalism*, pages 15–118. MIT Press, Cambridge, Mass., 1978.

[20] S. D. Kamvar, M. T. Schlosser, and H. Garcia-Molina. The Eigentrust algorithm for reputation management in P2P networks. In *Conference on the World Wide Web (WWW)*, pages 640–651, 2003.

[21] M. Kandori. Social norms and community enforcement. *Review of Economic Studies*, 59:63–80, 1992.

[22] LogiSense Corporation. LogiSense web site. http://www.logisense.com/tm_p2p.html.

[23] L. Lovasz and P. Winkler. Mixing of random walks and other diffusions on a graph. In *Surveys in Combinatorics, 1993, Walker (Ed.), London Mathematical Society Lecture Note Series 187, Cambridge University Press*. 1995.

[24] Open Source Technology Group. Slashdot FAQ - comments and moderation. http://slashdot.org/faq/com-mod.shtml#cm700.

[25] OSMB LLC. Gnutella web site. http://www.gnutella.com/.

[26] M. L. Puterman. *Markov Decision Processes*. Wiley, 1994.

[27] K. Ranganathan, M. Ripeanu, A. Sarin, and I. Foster. 'to share or not to share' an analysis of incentives to contribute in collaborative file sharing environments. In *Workshop on Economics of Peer-to-Peer Systems (P2PECON)*, 2003.

[28] S. I. Resnick. *Adventures in Stochastic Processes*. Birkhauser, 1992.

[29] SETI@home. SETI@home web page. http://setiathome.ssl.berkeley.edu/.

[30] Sharman Networks Ltd. Kazaa web site. http://www.kazaa.com/.

[31] V. Vishnumurthy, S. Chandrakumar, and E. Sirer. Karma: A secure economic framework for peer-to-peer resource sharing. In *Workshop on Economics of Peer-to-Peer Systems (P2PECON)*, 2003.

[32] L. Xiong and L. Liu. Building trust in decentralized peer-to-peer electronic communities. In *Internation Conference on Electronic Commerce Research (ICECR)*, 2002.

[33] H. Zhang, A. Goel, R. Govindan, K. Mason, and B. V. Roy. Making eigenvector-based reputation systems robust to collusion. In *Workshop on Algorithms and Models for the Web-Graph(WAW)*, pages 92–104, 2004.