Addressing the Free-Rider Problem in File-Sharing Systems: A Mechanism-Design Approach

Ryan Porter Yoav Shoham

Computer Science Department
Stanford University
Stanford, CA 94305
{rwporter,shoham}@stanford.edu

Abstract

We consider popular file-sharing systems such as Kazaa from a game theoretic perspective. Such systems are inherently vulnerable to free-riding, that is, to users who download files but do not contribute in return. The result is that these systems do not maximize the number of files exchanged; in the parlance of mechanism design, these mechanisms do not implement the desired social choice function. In the spirit of mechanism design, we search for alternative, more efficient mechanisms. An interesting aspect of the mechanism-design problem is that, in this setting, mechanisms do not have the luxury of ruling out free-riding as an available strategy. More annoyingly, a mechanism cannot distinguish between a free-rider and a good citizen who happens to not have goods to contribute. We nonetheless manage to identify more efficient mechanisms, including one that is provably optimal within the constraints we define.

1 Introduction

It is hard to overstate the impact file sharing systems such as Kazaa [2] and Gnutella [1] have had on electronic commerce and on the pattern of internet usage in general. However, as wildly popular as they have been, they have also been plagued by the problem of free-riding. It is perfectly possible for individuals to download from others but not allow others to access their own repositories, leading to a "tragedy of the commons" [7]; and, there is ample evidence that a significant fraction of the users do just that [3, 10]. Of course, the fact that file-sharing systems are thriving is evidence that enough people do share, but the existence of free-riders means that the exchange of files is not as efficient as it could be. In particular, there are additional files that would be downloaded, had the free-riders made them available. The question we ask is whether, and to what extent, the free-riders can be incented to share their files.

Of course, we are not the first to ask this question, or to offer a solution. Specifically, micro-payment schemes (see, e.g., [6], along with [8], which presents an escrow service that enables payments in a P2P system) and internal currencies (such as the "mojo" of the now extinct MojoNation) have been proposed as means of incenting self-interested users to allow downloads, and [4] calls for a comparison of how well these approaches can overcome the self-interest of agents. Our work differs from most of the previous work in this area because we do not allow either monetary payments (due to the unpopularity of micro-payments) or internal currencies with no monetary value (because we focus on single-shot interactions).

While we are directly motivated by entertainment-based systems such as Kazaa and Gnutella, it should be evident that the problem we are addressing is quite generic, and pertains to any barter-like system for exchanging resources. For example, in distributed storage systems that invite individuals to store redundant copies of others' files, the question is how to incent people to accept these files. Similarly, in inter-domain routing, the question is how to incent one network operator to accommodate the packets of another domain. However, both for concreteness and because each application exhibits some special properties, we will be couching the discussion specifically in terms of Kazaa- and Gnutella-style file-sharing systems.

[1]This work was supported in part by DARPA grant F30602-00-2-0598.

To understand such a system we need a formal model. We will give our model in the next section, but here is the intuition. We posit two types of users: one that enjoys having others download his songs (the altruistic type) and one that does not (the selfish type). Now imagine a one-shot interaction between two agents. In existing file-sharing systems, each agent has two options – to share (by placing the files in a particular folder) or not. Without being too precise about the game theoretic setting (we will be precise later), one can imagine this giving rise to the following normal-form game structure:

	Share	Free-ride
Share		
Free-ride		

Table 1: Structure of a file-sharing game.

To analyze this game we need to be precise about the payoffs as well as the probabilities involved. However, even absent those, intuitively speaking, it is obvious that both the altruistic types and the selfish types have dominant strategies – the altruistic types always share, and the selfish types always free-ride. The question is what mechanisms can be created that will give rise to other games, ones in which free-riders have the incentive to allow others to download their songs.

This mechanism design problem is interesting in several ways. First, no mechanism can eliminate free-riding from the repertoire of player actions. The reason is that free-riding is observationally indistinguishable from sharing by a player that happens to have nothing to share. In concrete terms, if the (e.g., music) sharing folder of a user is empty, it might be because the user doesn't have any songs in his file system, or because he has but has chosen to put them in a different folder. As a result, the action space in the mechanisms we will consider will always be a superset of $\{$Share,Free-ride$\}$.

The same indistinguishability of types by the mechanism gives rise to another constraint – the mechanism cannot directly penalize a free-rider, since it cannot tell free-riders from impoverished sharers. This property, which we call free-ride ignorance, is a critical design constraint in our setting.

With this as background, here is a synopsis of the rest of the paper. We first present a formal model of the situation. Certain properties of this model to keep in mind are the following:

1. We are only modelling a one-shot interaction between two randomly chosen users. Of course one wants ultimately to reason about the repeated interaction among numerous, randomly matched individuals. We make some remarks about this at the end of the paper, but the formal analysis is for the single interaction.

2. We focus on mechanisms that allow users, once told the transaction that will occur (which specifies, in each direction, whether a song will be transferred), to accept or reject the entire transaction. The rationale behind this restriction is that a user can alter his client (or download a modified version) that would disconnect from the server when an unfavorable transaction begins. In the parlance of game theory, by doing so we ensure ex post individual rationality.

After defining the formal model, we consider a sequence of three games. The first is of the form given above, and serves as the baseline. We next add the action "Trade", in which an agent offers to transfer a song (if required by his opponent) in order to get a song in return. This action allows selfish agents to swap songs with each other while still free-riding off of altruistic agents who play "Share". However, this game has the undesirable property that altruistic agents may now have incentive to play "Trade" instead of "Share", in order to extract a song from a selfish agent playing "Trade". For this reason, we then a fourth action, "Conditional-Share", which allows an altruistic to give away songs as if he were playing "Share", but also allows him to extract songs from selfish agents as if he were playing "Trade". In each game, we characterize the set of all equilibria. We prove that the third game maximizes the expected number of transfers. We then discuss the fact that a game with this property must have other equilibria that yield a lower number of expected transfers, and present a possible means of steering the system towards the desired equilibrium.

2 Problem Formulation

The problem formulation consists of two parts: the setting, which defines the agents and the possible outcomes, and the mechanism, which defines how the agents will interact to arrive at an outcome. We now formally define these two components. Good introductions to game theory and mechanism design can be found in [5] and [9], respectively.

2.1 Setting

Formally, the setting is defined by the tuple $\langle N, \Theta, P_s, O, T, \overline{u}(\cdot), p, u(\cdot)\rangle$.

- N: The set of agents is $N = \{a_1, a_2\}$. When referring to an arbitrary agent a_i, we will use a_{-i} to refer to the other agent.

- Θ: Each agent a_i is characterized by a type $\theta_i \in \Theta$, where a type is defined by the tuple: $\langle v_i, c_i\rangle$. An agent gains a value of v_i if it receives a desired item from the other agent, and incurs a cost of c_i if it delivers an item to the other agent. In general, we expect v_i to be significantly greater than the absolute value of c_i. While v_i will always be positive, c_i can either be positive or negative. A positive cost results from the actual cost of delivering the good (e.g., the bandwidth used to upload a song), while a negative cost results from utility gained from contributing to the system. The type defines the utility function of the agent, as described below.

  For expositional purposes, we will only consider two possible types: $\Theta = (\theta_i^a, \theta_i^s)$. For both types, $v_i = v$ and $|c_i| = c$, where v and c are positive constants. The difference is that $c_i = c$ for θ_i^s and $c_i = -c$ for θ_i^a. In the sequel we will refer to agents with type θ_i^s as selfish and those with type θ_i^a as altruistic.

- P_s: For each agent a_i, $\theta_i = \theta_i^s$ with probability P_s, and $\theta_i = \theta_i^a$ with probability $1 - P_s$. Each agent's type is drawn independently according to this probability. We assume that probability for each type is greater than zero ($0 < Pr < 1$).

  In general, if we do not limit consideration to these two types, the type of an agent is drawn from a probability density function Pr over Θ. We will use this more general notation for convenience in later definitions.

- O: The set of possible outcomes is $O = \{(\leftarrow_\perp, \rightarrow_\perp), (\leftarrow_\perp, \rightarrow_1), (\leftarrow_\perp, \rightarrow_{1,2}), (\leftarrow_2, \rightarrow_\perp), (\leftarrow_2, \rightarrow_1),$ $(\leftarrow_2, \rightarrow_{1,2}), (\leftarrow_{1,2}, \rightarrow_\perp), (\leftarrow_{1,2}, \rightarrow_1), (\leftarrow_{1,2}, \rightarrow_{1,2})\}$ to represent the nine different combinations of trading possibilities. The left and right arrows represent an item transferred from agent 2 to 1 and from agent 1 to 2, respectively. The subscript on the arrow denotes the agents that must own an item that the other agent desires in order for the transfer to occur, with "$\perp$" denoting that the transfer never occurs. For example, $(\leftarrow_{1,2}, \rightarrow_1)$ means that both agents must have an item that the other desires in order for agent 2 to transfer an item to agent 1, but that agent 1 will transfer an item desired by agent 2 even if agent 2 does not have an item for agent 1.

  Let $Prereq(o, \leftarrow)$ and $Prereq(o, \rightarrow)$ denote the prerequisites for transferring a file in the direction of the arrow. For example, $Prereq((\leftarrow_2, \rightarrow_{1,2}), \leftarrow) = \{2\}$.

- T: An outcome does not specify which transfers will occur– it only defines the condition for both transfers, based on whether each agent has a song for the other agent. The four possible transactions are $T = \{-, \leftarrow, \rightarrow, \leftrightarrow\}$, to represent, respectively: no transfers, a transfer from agent 2 to 1 only, a transfer from agent 1 to 2 only, and a transfer in both directions.

- $\overline{u}(\cdot)$: An agent's valuation for particular transactions are defined by the function $\overline{u}_i : T \rightarrow \Re$, which is specified in Table 2.

- p: With probability p, an arbitrary agent a_i has an item desired by the other agent, a_{-i}. The probability that a_{-i} has an item that a_i desires is identical and independent. Together with an outcome $o \in O$, this probability induces a lottery over transactions. We will use $p(t|o)$ to denote the probability of transaction t occurring, given outcome o.

  A potential extension of this setting would also include two more terms in the type of an agent: one that models observing the set of items owned by an agent (from which the agent computes a posterior probability that it owns an item that the other agent desires), and another that models observing the set of items that the agent itself desires (from which it computes a posterior probability that it desires an item owned by the other agent). However, while this extension would create a more realistic setting, it would not change the nature of our results.

- u($\cdot$): The utility function for an agent, $u_i : O \times \Theta \rightarrow \Re$, maps each outcome and agent type to an expected value, as specified by Table 3. We assume that each agent is rational, in that it aims to maximize its expected utility. Thus, an agent's utility for a lottery is simply the expected value of $\overline{u}_i(t)$, taken over $p(t|o)$.

Transaction $t \in T$	Valuation $\overline{u}_i(t, \theta_i)$
$-$	0
$\leftarrow$	v_i
$\rightarrow$	$-c_i$
$\leftrightarrow$	$v_i - c_i$

Table 2: Utility to agent a_i for each possible transaction.

Outcome $o \in O$	Lottery Over Transactions				Utility $u_i(o, \theta_i)$
	$p(-\|o)$	$p(\leftarrow \|o)$	$p(\rightarrow \|o)$	$p(\leftrightarrow \|o)$	
$(\leftarrow_\perp, \rightarrow_\perp)$	1	0	0	0	0
$(\leftarrow_\perp, \rightarrow_1)$	$1-p$	0	p	0	$-c_i \cdot p$
$(\leftarrow_\perp, \rightarrow_{1,2})$	$1-p^2$	0	p^2	0	$-c_i \cdot p^2$
$(\leftarrow_2, \rightarrow_\perp)$	$1-p$	p	0	0	$v_i \cdot p$
$(\leftarrow_2, \rightarrow_1)$	$(1-p)^2$	$p \cdot (1-p)$	$p \cdot (1-p)$	p^2	$v_i \cdot p - c_i \cdot p$
$(\leftarrow_2, \rightarrow_{1,2})$	$1-p$	$p \cdot (1-p)$	0	p^2	$v_i \cdot p - c_i \cdot p^2$
$(\leftarrow_{1,2}, \rightarrow_\perp)$	$1-p^2$	p^2	0	0	$v_i \cdot p^2$
$(\leftarrow_{1,2}, \rightarrow_1)$	$1-p$	0	$p \cdot (1-p)$	p^2	$v_i \cdot p^2 - c_i \cdot p$
$(\leftarrow_{1,2}, \rightarrow_{1,2})$	$1-p^2$	0	0	p^2	$v_i \cdot p^2 - c_i \cdot p^2$

Table 3: Utility to agent a_i for each possible outcome.

2.2 Mechanism

The mechanism defines the protocol for interaction between the agents and the center that culminates with the selection of an outcome. It is formally defined by a tuple $\Gamma = (A, g(\cdot))$, where A is the action space of each agent, and $g : A^2 \rightarrow O$ maps the actions of both agents to an outcome $o \in O$. As the mechanism designer, we specify both A and $g(\cdot)$, subject to the constraints below.

The setting and the mechanism are common knowledge among the agents. A mechanism Γ, combined with the setting, induces a Bayesian game between the two agents. In this game, each agent first privately observes its true type θ_i, drawn according to P_s. Then, it selects an action based on this type. Thus, we will speak more generally of an agent's (mixed) strategy $s_i : \Theta \rightarrow \Delta A$, which maps each of its possible types to a distribution over the action that it will take in the game. We will use $s_i(\theta_i, a)$ to denote the probability assigned to action $a \in A$ by the strategy $s_i(\theta_i)$, and we will use $s_i(\theta_i) = a$ to denote the (pure) strategy in which all probability is assigned to action a.

We will represent a Bayesian game in a compact way as a single payoff matrix for the row player (see, for example, Table 4), even though there are two different possible types for both the row and column player. First, note that since the game is symmetric, we only need to list the payoffs for one of the players. Second, we do not need a separate matrix for both possible types of the column player, because only the action of the column player affects the payoff of the row player. Finally, we list the utility to the row player in terms of v_i and c_i so that it is not specific to a single type.

The expected utility, $EU_i(s)$ of agent a_i for a strategy profile $s = (s_1, s_2)$ is computed by taking the expectation over the possible instantiations of agent types and actions.

$$EU_i(s) = \sum_{\theta_i, \theta_{-i} \in \Theta} \sum_{a_i, a_{-i} \in A} Pr(\theta_i) \cdot Pr(\theta_{-i}) \cdot s_i(\theta_i, a_i) \cdot s_{-i}(\theta_{-i}, a_{-i}) \cdot u_i(g(a_i, a_{-i}), \theta_i)$$

Because each agent is self-interested, a strategy profile is not "stable" unless each agent maximizes its utility. The condition for stability we strive for here is that of Bayes-Nash equilibrium.

Definition 1 A strategy profile $s^* = (s_1^*, s_2^*)$ is a Bayes-Nash equilibrium of mechanism Γ if the following condition holds for all i, s_i':

$$EU_i(s_i^*, s_{-i}^*) \geq EU_i(s_i', s_{-i}^*)$$

An equivalent formulation of the condition is that, for each type θ_i, agent i must not be able to increase its expected utility (taken over θ_{-i}) by deviating from playing $s^*(\theta_i)$, holding constant the other agent's strategy as s_{-i}^*.

We will restrict consideration to symmetric equilibria (that is, equilibria in which $s_1 = s_2$), and represent them using a single strategy s_i^*. We make this assumption because our symmetric Bayesian game is intended to model a single interaction between two agents randomly drawn from a large population. Under this model, one interpretation of $s_i^*(\theta_i, a_i)$ is as the fraction of agents in the population of type θ_i who play the pure strategy a_i. Because of the fact that, when two agents are paired up to play the game, there is no notion of which one is a_1 or a_2, their strategies cannot differ depending on which role they take.

For the Bayesian games we analyze below, it will often be the case that for a particular type θ_i, there exists an action a which always yields at least as high of a utility as any other action a', regardless of the action played by the other agent. If it is also the case that, for each alternative action a', there exists an action by the opponent such that action a yields a strictly higher utility than a', then action a is called a weakly dominant strategy for that type.

Definition 2 An action $a \in A$ is a weakly dominant strategy for type θ_i, if $\forall a',\ (\forall b \in A,\ u_i(a,b) \geq u_i(a',b)) \wedge (\exists c \in A,\ u_i(a,c) > u_i(a',c))$.

If it is the case that a yields a higher expected utility for all actions by the opponent, then it is also a strictly dominant strategy.

Definition 3 An action $a \in A$ is a strictly dominant strategy for type θ_i, if $\forall a',\ (\forall b \in A,\ u_i(a,b) > u_i(a',b))$.

While it is natural to expect a rational agent to play according to a weakly dominant strategies if it exists, the fact that the inequality is not necessarily strict for possible actions by your opponent means that it can be rational to play another action.

2.2.1 Implementation

A social choice function $f : \Theta^2 \rightarrow O$ maps every profile of agent types to an outcome. Our goal as mechanism designer is to implement a social function that maximizes some objective function. A mechanism $\Gamma = (A, g(\cdot))$ implements the social choice function $f(\cdot)$ in Bayes-Nash equilibrium if there exists a Bayes-Nash equilibrium s^* of Γ such that $g(s^*(\theta)) = f(\theta)$ for all $\theta \in \Theta$. That is, if each agent plays according to its equilibrium strategy, then, for every possible profile of types, the outcome of the mechanism will be the outcome of the social choice function. Intuitively, if a social choice function is implemented, then the mechanism has overcome its lack of knowledge of the agents' private information.

Note that implementation only requires existence, and not uniqueness, of an equilibrium with the desired property. It turns out that the social choice function we implement cannot be implemented under the additional requirement of uniqueness, and later we will address the issue of how agents may converge to the desired equilibrium.

The objective function $z(f(\cdot))$ we use to evaluate an implemented social choice function $f(\cdot)$ is the expected number of transfers, which is computed using the following equation.

$$z(f(\cdot)) = \sum_{\theta_1,\theta_2 \in \Theta} Pr(\theta_1) \cdot Pr(\theta_2) \cdot [p(\leftarrow | f(\theta_1,\theta_2)) + p(\rightarrow | f(\theta_1,\theta_2)) + 2 \cdot p(\leftrightarrow | f(\theta_1,\theta_2))]$$

2.2.2 Mechanism Requirements

The setting we model imposes two important constraints on our mechanism: one on the set of social choice functions that we can implement, and another on which sets of actions A are valid.

The first restriction is due to the fact that agents can choose to reject the transaction that results from the mechanism, instead achieving the baseline utility of zero. For example, a selfish agent who is told to transfer a song to the other agent while he is not receiving one in return would choose to stop uploading the file. Only if he were receiving a song in return (which we assume occurs simultaneously, so that an agent who stops uploading would not be allowed to complete the download), would he accept the transaction.

Thus, we require that the social choice function we implement always yield an agent a valuation of at least zero. Formally, this requirement is called ex post individual rationality.

Definition 4 A social choice function $f(\cdot)$ satisfies ex post individual rationality (ex post IR) if for all θ, i:

$$\nexists t \in T, (p(t|f(\theta)) > 0) \wedge (\overline{u}_i(t,\theta_i) < 0)$$

It is common to discuss individual rationality as a property of a mechanism instead of the social choice function. A mechanism Γ satisfies ex post IR if it implements a social choice function that satisfies ex post IR. However, we require that the social choice function we implement satisfy this requirement instead of just the mechanism, because a mechanism can implement multiple social choice functions. Thus, the fact that a mechanism implements a social choice function $f(\cdot)$ and is ex post IR does not imply that $f(\cdot)$ satisfies ex post IR.

The second restriction is imposed by our assumption that an agent can undetectably "sabotage" its chances of matching an item it owns with an item that the other agent desires. This assumption means that an agent can play each action in the Bayesian game without the possibility of transferring an item to the other agent, and without preventing a transfer from the other agent if that agent is willing to transfer the item without receiving one in return. Of course, if the other agent does demand an item in return, then the free-riding agent has no chance of receiving the item.

There are two possible ways to model this restriction. One option is to expand the action space of an agent so that it can privately set the probability of having an item that the other agent desires to be 0 or p (while the agent may in fact be able to set the probability to any value between 0 and p, he will never have incentive to do so, because his utility varies linearly with this probability). A second option, and the one we will use, is to constrain the mechanism to include a "free-ride" variant of each action. Under this option, an agent essentially declares to the mechanism that it is free-riding when he plays such a variant. However, we do not allow the mechanism to exploit this knowledge, because the action is actually a proxy for "free-riding" while playing a different action.

The formal construction of a "free-riding" variant a' of each action a is specified by the following definition.

Definition 5 A mechanism Γ satisfies free-ride ignorance if $\forall a \in A, \exists a' \in A, \forall b \in A$, the following conditions hold:

- $(Prereq(g(a', b), \rightarrow) = \{\perp\})$
- $(Prereq(g(a, b), \leftarrow) = \{2\}) \implies (Prereq(g(a', b), \leftarrow) = \{2\})$
- $(Prereq(g(a, b), \leftarrow) = \{1, 2\}) \vee (Prereq(g(a, b), \leftarrow) = \{\perp\}) \implies (Prereq(g(a', b), \leftarrow) = \{\perp\})$

3 Initial Game

We begin by examining mechanism Γ_{SF}, which has two actions, share (S) and free-ride (F), and induces the game specified in Table 4. In our setting, this is the most basic game one could examine: one action allows agents to share files, while the other is a consequence of the requirement that the mechanism satisfy free-ride ignorance.

	Share (S)	Free-ride (F)
Share (S)	$v_i \cdot p - c_i \cdot p$ $(\leftarrow_2, \rightarrow_1)$	$-c_i \cdot p$ $(\leftarrow_\perp, \rightarrow_1)$
Free-ride (F)	$v_i \cdot p$ $(\leftarrow_2, \rightarrow_\perp)$	0 $(\leftarrow_\perp, \rightarrow_\perp)$

Table 4: Outcomes, along with the utilities of the row player a_i, in the game induced by mechanism Γ_{SF}.

As was anticipated in the Introduction, in this game it is indeed a strictly dominant strategy for altruistic agents to play S and for selfish agents to play F. Thus, the unique Bayes-Nash equilibrium is for both agents to play according to the strategy $s_i^*(\theta_i^a) = S$ and $s_i^*(\theta_i^s) = F$.

4 Step I: Promoting Trades

In order to induce selfish agents not to free-ride, we add the action Trade (T). This action is similar to F, and it is equivalent to this action (for both directions of a possible transfer) when played against an agent who plays either S or F. However, when both agents play T, and when both agents have an item that the other desires, a trade occurs. Intuitively, an agent who plays T is only willing to transfer an item if doing so is necessary in order to receive one in return, an offer which is consistent with our expectation that $v_i > |c_i|$. The game induced by this mechanism, which we label Γ_{STF}, is specified by Table 5. Note that this mechanism satisfies free-ride ignorance, because F is the free-ride variant of T.

	Share (S)	Trade (T)	Free-ride (F)
Share (S)	$v_i \cdot p - c_i \cdot p$ $(\leftarrow_2, \rightarrow_1)$	$-c_i \cdot p$ $(\leftarrow_\perp, \rightarrow_1)$	$-c_i \cdot p$ $(\leftarrow_\perp, \rightarrow_1)$
Trade (T)	$v_i \cdot p$ $(\leftarrow_2, \rightarrow_\perp)$	$(v_i - c_i) \cdot p^2$ $(\leftarrow_{1,2}, \rightarrow_{1,2})$	0 $(\leftarrow_\perp, \rightarrow_\perp)$
Free-ride (F)	$v_i \cdot p$ $(\leftarrow_2, \rightarrow_\perp)$	0 $(\leftarrow_\perp, \rightarrow_\perp)$	0 $(\leftarrow_\perp, \rightarrow_\perp)$

Table 5: Outcomes, along with the utilities of the row player a_i, in the game induced by mechanism Γ_{STF}.

Now, it is no longer a strictly dominant strategy for selfish agents to play F, and in fact it is a weakly dominant strategy for them to play T. However, depending on both p and P_s, it may also no longer be a dominant strategy for altruistic agents to play S. Specifically, an altruistic agent may have incentive to play T in order to have a chance of extracting an item from an opponent who plays T, even though playing this action may prevent him from giving away an item.

The characterize all Bayes-Nash equilibria, we first analyze the case in which all selfish agents play their weakly dominant strategy T. Since F is a strictly dominated strategy for an altruistic agent, we can restrict consideration to S and T for this type of agent.

In this case, the action T yields a strictly higher expected utility than S for an altruistic agent, regardless of the strategy adopted by the opposing agent for its altruistic type, if the following condition holds:

$$\begin{aligned} v \cdot p \cdot (1 - P_s) + (v + c) \cdot p^2 \cdot P_s &> (v + c) \cdot p \cdot (1 - P_s) + c \cdot p \cdot P_s \\ (v + c) \cdot p^2 \cdot P_s &> c \cdot p \\ p &> \frac{c}{(v + c) \cdot P_s} \end{aligned}$$

By examining the payoffs for (S, T) and (T, T), we find that the condition for the action S to be a strictly dominant strategy for an altruistic agent is that $p < \frac{c}{v+c}$.

For all other values of p, there exist two pure strategy Bayes-Nash equilibria: one in which altruistic agents play S, and another in which they play T. Also, there exists a mixed strategy equilibrium in which altruistic agents randomize between these two actions. The condition for these two actions to yield an equal expected utility for an altruistic agents is as follows:

$$\begin{aligned} v \cdot p \cdot (1 - P_s) \cdot s_i^*(\theta_i^a, S) + (v + c) \cdot p^2 \cdot [P_s + (1 - P_s) \cdot (1 - s_i^*(\theta_i^a, S))] = \\ (v + c) \cdot p \cdot (1 - P_s) \cdot s_i^*(\theta_i^a, S) + c \cdot p \cdot [P_s + (1 - P_s) \cdot (1 - s_i^*(\theta_i^a, S))] \end{aligned}$$

$$\begin{aligned} (v + c) \cdot p \cdot [1 - (1 - P_s) \cdot s_i^*(\theta_i^a, S)] &= c \cdot (1 - P_s) \cdot s_i^*(\theta_i^a, S) + c \cdot [1 - (1 - P_s) \cdot s_i^*(\theta_i^a, S)] \\ (v + c) \cdot p - (v + c) \cdot p \cdot (1 - P_s) \cdot s_i^*(\theta_i^a, S) &= c \\ s_i^*(\theta_i^a, S) &= \frac{(v + c) \cdot p - c}{(v + c) \cdot p \cdot (1 - P_s)} \end{aligned}$$

To complete the characterization of the equilibria, we note that selfish agents will never play S, since it is strictly dominated for them, and that there cannot exist a mixed strategy equilibrium in which they mix between T and F, because, if there is any possibility that the opponent of a selfish agent plays T, then playing T himself yields a strictly higher expected utility than playing F. Thus, the only other equilibrium is the one from the previous game. Table 6 summarizes these results, characterizing all Bayes-Nash equilibria for each region of a partition of the parameter space for our setting.

5 Step II: Maximizing Transfers

In this section, we address the problem we created in the previous section– namely, that altruistic agents may no longer have incentive to share their items. To do this, we add a fourth action, "Conditional-Share" (C), which allows an agent

Condition	All Bayes-Nash Equilibria
$p < \frac{c}{v+c}$	(1) $s_i^*(\theta_i^a) = S$ and $s_i^*(\theta_i^s) = T$ (2) $s_i^*(\theta_i^a) = S$ and $s_i^*(\theta_i^s) = F$
$\frac{c}{v+c} \leq p \leq \frac{c}{(v+c)\cdot P_s}$	(1) $s_i^*(\theta_i^a) = S$ and $s_i^*(\theta_i^s) = T$ (2) $s_i^*(\theta_i^a) = T$ and $s_i^*(\theta_i^s) = T$ (3) $s_i^*(\theta_i^a) = [S : \frac{(v+c)\cdot p - c}{(v+c)\cdot p\cdot(1-P_s)}; T : 1 - \frac{(v+c)\cdot p - c}{(v+c)\cdot p\cdot(1-P_s)}]$ and $s_i^*(\theta_i^s) = T$ (4) $s_i^*(\theta_i^a) = S$ and $s_i^*(\theta_i^s) = F$
$p > \frac{c}{(v+c)\cdot P_s}$	(1) $s_i^*(\theta_i^a) = T$ and $s_i^*(\theta_i^s) = T$ (2) $s_i^*(\theta_i^a) = S$ and $s_i^*(\theta_i^s) = F$

Table 6: A characterization of the Bayes-Nash equilibria for the game induced by mechanism Γ_{STF}. A mixed strategy $s_i(\theta_i)$ is specified in the form of a lottery, $[a : s_i(\theta_i, a); a' : s_i(\theta_i, a'); \ldots]$, where each action that is played with positive probability is followed by its associated probability.

to give away an item whenever possible, but demands an item in return if he transfers an item to the other agent and the other agent has an item he desires. Of course, if the other agent free-rides, then an agent who plays C will give away an item if possible and never get one in return, because it appears to him that the other agent simply did not have an item for him. However, an opponent who plays T must, if possible, give him an item in return for receiving one. Thus, an agent who plays C extracts an item from his opponent under the exact same conditions as an agent who plays T does, but has a greater chance of giving away an item for each action of his opponent. The game induced by this mechanism, which we call Γ_{SCTF}, is specified by Table 7. Note that this mechanism also satisfies free-ride ignorance, because F is the free-ride variant of C.

	Share (S)	Conditional-Share (C)	Trade (T)	Free-ride (F)
Share (S)	$v_i \cdot p - c_i \cdot p$ $(\leftarrow_2, \rightarrow_1)$	$v_i \cdot p - c_i \cdot p$ $(\leftarrow_2, \rightarrow_1)$	$-c_i \cdot p$ $(\leftarrow_\perp, \rightarrow_1)$	$-c_i \cdot p$ $(\leftarrow_\perp, \rightarrow_1)$
Conditional-Share (C)	$v_i \cdot p - c_i \cdot p$ $(\leftarrow_2, \rightarrow_1)$	$v_i \cdot p - c_i \cdot p$ $(\leftarrow_2, \rightarrow_1)$	$v_i \cdot p^2 - c_i \cdot p$ $(\leftarrow_{1,2}, \rightarrow_1)$	$-c_i \cdot p$ $(\leftarrow_\perp, \rightarrow_1)$
Trade (T)	$v_i \cdot p$ $(\leftarrow_2, \rightarrow_\perp)$	$v_i \cdot p - c_i \cdot p^2$ $(\leftarrow_2, \rightarrow_{1,2})$	$(v_i - c_i) \cdot p^2$ $(\leftarrow_{1,2}, \rightarrow_{1,2})$	0 $(\leftarrow_\perp, \rightarrow_\perp)$
Free-ride (F)	$v_i \cdot p$ $(\leftarrow_2, \rightarrow_\perp)$	$v_i \cdot p$ $(\leftarrow_2, \rightarrow_\perp)$	0 $(\leftarrow_\perp, \rightarrow_\perp)$	0 $(\leftarrow_\perp, \rightarrow_\perp)$

Table 7: Outcomes, along with the utilities of the row player a_i, in the game induced by mechanism Γ_{SCTF}.

Now, it is a weakly dominant strategy for an altruistic agent to play C. The disadvantage of adding this action is that it may undo the work of adding T, because it is no longer a weakly dominant strategy for selfish agents to play T. Instead, a selfish agent must tradeoff the possibility of needing an item to trade with another selfish agent against the possibility of the being forced to give an item to an altruistic agent who would give him an item even if he chooses to free-ride. However, under conditions that we consider likely to be satisfied, an optimal equilibrium exists in this game.

To characterize the set of Bayes-Nash equilibria, we first note that only S and C are not strictly dominated for altruistic agents, and that only T and F are non strictly dominated for selfish agents.

We first find all equilibria in which selfish agents assign some probability to T. When this is true, C always yields a higher expected utility to altruistic agents than S, and thus we can limit consideration to strategies that in which all altruistic agents play C. Here, it is never rational for a selfish agent to play T over F if the following condition holds:

$$\begin{aligned}(v \cdot p - c \cdot p^2) \cdot (1 - P_s) + (v - c) \cdot p^2 \cdot P_s &< v \cdot p \cdot (1 - P_s) \\ v \cdot P_s - c &< 0 \\ P_s &< \frac{c}{v}\end{aligned}$$

On the other hand, when $P_s \geq \frac{c}{v}$ holds, there exist two equilibria in which selfish agents assign some probability to T: one in which T is the pure strategy of the selfish agent, and another in which selfish agents mix between T and F. Specifically, the expected utility for both of these actions are equal when the following condition is satisfied:

$$
\begin{aligned}
(v \cdot p - c \cdot p^2) \cdot (1 - P_s) + (v - c) \cdot p^2 \cdot P_s \cdot s_i^*(\theta_i^s, T) &= v \cdot p \cdot (1 - P_s) \\
s_i^*(\theta_i^s, T) &= \frac{c}{v \cdot P_s}
\end{aligned}
$$

To complete the characterization, we need to find all equilibria in which selfish agents always play F. Since C is equivalent to S when there is no chance that the other agent will play T, any $s_i^*(\theta_i^a)$ that assigns positive probability only to C and S forms an equilibrium with $s_i^*(\theta_i^s) = F$. These results are summarized in Table 8.

Condition	All Bayes-Nash Equilibria
$P_s < \frac{c}{v}$	(1) $\forall x \in [0,1], s_i^*(\theta_i^a) = [S : x; C : (1-x)]$ and $s_i^*(\theta_i^s) = F$
$P_s \geq \frac{c}{v}$	(1) $s_i^*(\theta_i^a) = C$ and $s_i^*(\theta_i^s) = T$ (2) $\forall x \in [0,1], s_i^*(\theta_i^a) = [S : x; C : (1-x)]$ and $s_i^*(\theta_i^s) = F$ (3) $s_i^*(\theta_i^a) = C$ and $s_i^*(\theta_i^s) = [T : \frac{c}{v \cdot P_s}; F : (1 - \frac{c}{v \cdot P_s})]$

Table 8: A characterization of the Bayes-Nash equilibria for the game induced by mechanism Γ_{SCTF}.

We have shown that when $P_s \geq \frac{c}{v}$ is satisfied, the strategy s_i^* such that $s_i^*(\theta_i^a) = C$ and $s_i^*(\theta_i^s) = T$ is an equilibrium and, therefore, Γ_{SCTF} implements the social choice function $f_{SCTF}(\cdot)$ specified in Table 9.

θ_1	θ_2	$f_{SCTF}(\theta_1, \theta_2)$
θ_i^a	θ_i^a	$(\leftarrow_2, \rightarrow_1)$
θ_i^a	θ_i^s	$(\leftarrow_{1,2}, \rightarrow_1)$
θ_i^s	θ_i^a	$(\leftarrow_2, \rightarrow_{1,2})$
θ_i^s	θ_i^s	$(\leftarrow_{1,2}, \rightarrow_{1,2})$

Table 9: Social Choice Function $f_{SCTF}(\cdot)$.

Since in the real-world settings that we are modelling we expect that v is significantly greater than c and that selfish agents make a large fraction of the population, we feel that the condition $P_s \geq \frac{c}{v}$ is likely to be satisfied. We now prove that implementing this social choice function is as well as we can do, even among mechanisms that are not constrained to satisfy free-ride ignorance.

Theorem 1 There does not exist a mechanism Γ' that implements a social choice function $f'(\cdot)$ that satisfies ex post IR and such that $z(f'(\cdot)) > z(f_{SCTF}(\cdot))$.

Proof: Assume by contradiction that there does exist a mechanism Γ' that implements a $f'(\cdot)$ such that $z(f'(\cdot)) > z(f_{SCTF}(\cdot)$. Then, it must be the case there exists some pair (θ_1, θ_2) such that the following inequality holds:

$$
\begin{aligned}
&p(\leftarrow | f'(\theta_1, \theta_2)) + p(\rightarrow | f'(\theta_1, \theta_2)) + 2 \cdot p(\leftrightarrow | f'(\theta_1, \theta_2)) > \\
&\qquad p(\leftarrow | f_{SCTF}(\theta_1, \theta_2)) + p(\rightarrow | f_{SCTF}(\theta_1, \theta_2)) + 2 \cdot p(\leftrightarrow | f_{SCTF}(\theta_1, \theta_2))
\end{aligned}
$$

However, for each pair (θ_1, θ_2), $f_{SCTF}(\theta_1, \theta_2)$ has the property that any possible transfer from a selfish agent has $\{1, 2\}$ as the prerequisites, and any possible transfer from a altruistic agent a_i has $\{i\}$ as the prerequisite. Thus, the probability of a transfer from an altruistic agent is already maximal, while the probability of a transfer from a selfish agent a_i can only possibly be increased by changing the prerequisites to $\{i\}$. Since this change would violate ex post IR, we have reached a contradiction. □

An unsatisfying aspect of our implementation of $f_{SCTF}(\cdot)$ under the condition $P_s \geq \frac{c}{v}$ is that two other social choice functions are also implemented that yield a smaller number of expected transfers. Unfortunately, the existence of these other two equilibria is unavoidable, given our requirement of free-ride ignorance. Informally, this is the case because any mechanism that implements a social choice function $f'(\cdot)$ that achieves the maximal value $z(f_{SCTF}(\cdot))$ must be based on an equilibrium in which selfish agents only randomize over actions that will yield the outcome $(\leftarrow_2, \rightarrow_{1,2})$ when paired with an altruistic agent and $(\leftarrow_{1,2}, \rightarrow_{1,2})$ when paired with another selfish agent. The "free-ride" variant of each of these actions will then yield the outcome $(\leftarrow_2, \rightarrow_\perp)$ when paired with an altruistic agent and $(\leftarrow_\perp, \rightarrow_\perp)$ when paired with another selfish agent, which means that the above analysis would apply to this game as well.

6 Discussion

If this single interaction occurs in a broader context that satisfies certain properties, then we can suggest a way to help the agents to converge on the desired equilibrium, when the condition for its existence is satisfied. Consider the expanded setting of a repeated Bayesian game with randomly paired opponents from a large population, in which each agent is myopic, in the sense that if it is selected to participate in the game for a particular round, it aims to maximize its expected utility only for the current Bayesian game, ignoring repeated game effects. Furthermore, assume that each agent forms its expectations over what its opponent will play based on previous rounds of the game. Then, starting from the initial game induced by Γ_{SF}, if we switch to Γ_{STCF}, it is unlikely that we will converge to the desired equilibrium, because agents will expect an opponent of the selfish type to play F, the action they played in the previous game. However, if we first switch to Γ_{STF}, then it is likely that the agents will converge on an equilibrium for the corresponding game in which all selfish agents play T, because it weakly dominates F (while both actions would yield equal utility based on expectations formed from Γ_{SF}, all it takes is a single agent to begin the transition to all agents playing T, and this agent could make this choice either by randomizing over best responses to its expectations or as a best response to an expectation formed from previous play combined with a prior). If we then switch to Γ_{SCTF}, we are more likely to converge to the desired equilibrium, because T is a best response for a selfish agent to the expectation that all other selfish agents will play T and to any expectation over what fraction of altruistic agents play S as opposed to C.

To summarize, we have formalized and studied a mechanism design problem concerning free-riding in P2P systems. Two important properties of our setting are that the mechanism cannot distinguish a free-rider from an agent who simply has nothing to share, and that an agent can choose to reject the transaction that results from the interaction. Starting from the basic file-sharing game, we first added the action “Trade” to promote the trading of songs between selfish users, and then added the action “Conditional-Share” so that there exists an equilibrium in which transfers occur whenever possible, given the constraints imposed by our setting. We characterized the set of all Bayes-Nash equilibria in each game. After proving the optimality of the desired equilibrium in the final game, we discussed how to possibly avoid the undesirable equilibria that must exist in conjunction with it.

References

[1] Gnutella home page, http://www.gnutella.com/.

[2] Kazaa home page, http://www.kazaa.com/.

[3] E. Adar and B. Huberman, Free riding on gnutella, First Monday 5 (2000).

[4] J. Feigenbaum and S. Shenker, Distributed algorithmic mechanism design: Recent results and future directions, Proceedings of the 6th International Workshop on Discrete Algorithms and Methods for Mobile Computing and Communications, 2002, pp. 1–13.

[5] D. Fudenberg and J. Tirole, Game theory, MIT Press, 1991.

[6] P. Golle, K. Leyton-Brown, and I. Mironov, Incentives for sharing in peer-to-peer networks, Proceedings of the Third AMC Conference on Electronic Commerce, 2001.

[7] G. Hardin, The tragedy of the commons, Science 162 (1968), 1243–1248.

[8] B. Horne, B. Pinkas, and T. Sander, Escrow services and incentives in peer-to-peer networks, Proceedings of the Third AMC Conference on Electronic Commerce, 2001.

[9] A. Mas-Colell, M.D. Whinston, and J.R. Green, Microeconomic theory, Oxford University Press, 1995.

[10] S. Saroiu, P. Gummadi, and S. Gribble, A measurement study of peer-to-peer file sharing systems, Proceedings of Multimedia Computing and Networking 2002 (MMCN ’02), 2002.