

A Peer-to-Peer System as an Exchange Economy

Christina Aperjis
Stanford University
Stanford, CA, 94305
caperjis@stanford.edu

Ramesh Johari
Stanford University
Stanford, CA, 94305
ramesh.johari@stanford.edu

ABSTRACT

We formulate a peer-to-peer system for filesharing as an exchange economy: a price is associated with each file, and users exchange files only when they can afford it. This formulation solves the free-riding problem, since uploading files is a necessary condition for being able to download. However, we do not explicitly introduce a currency; users must upload files in order to earn a budget for downloading. We discuss existence, uniqueness, and dynamic stability of the competitive equilibrium, which is always guaranteed to be Pareto efficient. In addition, a novel aspect of our approach is an allocation mechanism for clearing the market *out of equilibrium*. We analyze this mechanism when users can anticipate how their actions affect the allocation mechanism (price anticipating behavior). For this regime we characterize the Nash equilibria that will occur, and show that as the number of users increases, the Nash equilibrium rates become approximately Pareto efficient.

Categories and Subject Descriptors

C.2.4 [Computer-Communication Networks]: Distributed Systems—Distributed Applications

General Terms

Design, Economics, Performance, Theory

Keywords

Peer-to-peer systems, exchange economy

1. INTRODUCTION

In peer-to-peer systems, users share files or resources with each other. By sharing, a user incurs a cost (because uploading a file consumes network resources), but no direct benefit. Thus, if there is no mechanism that stimulates sharing, a user has a strong incentive to free ride, i.e., use the resources of other peers without contributing his own. Such behavior is observed in existing peer-to-peer systems; for instance, early data showed that nearly 70 percent of users of Gnutella were sharing no files, and nearly 50 percent of all responses were returned by the top 1 percent of sharing

hosts [1]. A more recent study shows that 85 percent of Gnutella users share no files [15]. Even worse, according to [1], there were users in Gnutella who were free riding on the system *despite* sharing files: the files that they were sharing were unpopular, and hence not widely uploaded.

Incentive mechanisms that penalize free riders or reward users that share have been proposed. In [5] users enjoy different levels of service according to how much they share their resources, while in [8] free riders are excluded from the system with some probability. In [7], a distributed rating scheme for tackling the free-rider problem is suggested. More general reputation mechanisms, such as those proposed in [12], can be used to obtain a system-wide reputation for each user. Using this information, each user will give priority to users with high reputation.

An alternate approach is to design a system where resource sharing is *required* to be able to use the resources of other users. This is the case in BitTorrent [6], where users download pieces of the file and at the same time upload the pieces they already have. Analogously, in [2] users directly trade resources between themselves.

Another option is to use monetary incentives to solve the problem of free riding. In this case, users must pay to download files from other peers. The payments may either be in monetary terms (e.g., [11]), or in an internal non-monetary currency. In the latter case, the budget of a user decreases every time he downloads a file, and increases every time he uploads a file. Such models are considered in [13], [16] and [19]. Recent work of Friedman et al. studies system performance as a function of the total amount of internal currency available [9].

In our model, we consider an internal currency and associate a price with each file. Users decide which files they are willing to upload, and the total upload rate they are willing to serve. In return, the system uses the current prices to provide a menu to the users of files available for download. The upload rate of a user generates a “budget” that can be spent to download available files. By maintaining different prices for different files, we avoid situations where users free-ride the system because the files they are sharing are unpopular. In particular, unpopular files will be assigned low prices.

We consider the utility of a user as a function of the rates at which he is downloading and uploading. It is reasonable to assume that the utility is increasing in the download rates. In particular, when the download rates are higher, the user gets the file sooner and is able to download more files in a fixed interval of time. Moreover, if there is some probability that the download will not complete successfully, this probability decreases as the download rate increases.

An important element of our model is that currency is not explicitly tracked; this makes our system lightweight and easily implemented. With this formulation we can also avoid cheap pseudo-

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

GameNets '06 Pisa, Italy

Copyright 2006 ACM 1-59593-507-X ...\$5.00.

nym [10], which are a drawback in most approaches for solving free-riding. Users cannot benefit by leaving the system and joining with a new identity, since user performance is determined only by the files uploaded. This naturally introduces a “transaction cost” into the system that prevents users from taking advantage of multiple identities. Of course, one shortcoming here is that users who join the system with little content of interest to others may be unable to download anything. One solution is to require such users to upload a file that is not desired by anyone. The price of this file can be set to be less than the price of any other file. In this way, new users do not get anything for free, and thus existing users do not have any incentive to rejoin the system with a different name.

In Section 2 we describe the model in more detail. In Section 3, we show the existence of a competitive equilibrium: a vector of prices at which demand of each file is equal to the corresponding supply. It is well known that such a vector is Pareto efficient. We derive conditions that guarantee uniqueness of the competitive equilibrium (up to scaling). In Section 4, we study the tâtonnement price adjustment process [4], and show that under some assumptions the rate of convergence around the equilibrium is linear in the number of users. This means that in a large system, the prices will rapidly converge.

A key aspect of our paper consists of a proposal to clear the market even out of equilibrium. In Section 5, we propose an allocation mechanism to allocate rates when demand is not equal to supply. We study the Nash equilibria when users anticipate how their actions affect the resulting allocation, and show that in large peer-to-peer systems, fully strategic behavior by the users will not ultimately cause large deviations from competitive equilibrium behavior. We conclude in Section 6.

Due to space constraints, several proofs are omitted; for these the reader is referred to [3].

2. MODEL

In this section we introduce our basic mathematical model, and connect it with the standard model of an *exchange economy* in microeconomics. We consider a peer-to-peer system with a set of users U who share a set of files F . User i has a subset of the files $S_i \subset F$, and is interested in downloading files in $T_i \subset F \setminus S_i$. Let x_{ij} be the rate at which user i downloads file $j \in T_i$, and let $\mathbf{x}_i = (x_{ij} : j \in T_i)$ be the vector of download rates of user i . Let y_{ij} be the rate at which user i is uploading file $j \in S_i$. The total upload rate of user i is $y_i = \sum_{j \in S_i} y_{ij}$. We assume that user i is indifferent between any two upload vectors $(y_{ij} : j \in S_i)$ and $(y'_{ij} : j \in S_i)$ with $\sum_{j \in S_i} y_{ij} = \sum_{j \in S_i} y'_{ij}$; in other words, his utility only depends on the vector of download rates \mathbf{x}_i and the total upload rate y_i . We make the following assumption.

Assumption 1 *The preference relation of a user on the set of feasible rate vectors is represented by a continuous utility function $v_i : \mathbb{R}_+^{|T_i|+1} \rightarrow \mathbb{R}$, which is strictly increasing in each download rate x_{ij} , $j \in T_i$; and strictly decreasing in the upload rate y_i .*

(Throughout the paper, \mathbb{R}_+ denotes the interval $[0, \infty)$.)

We introduce strictly positive prices in the system and consider a particular user i . Each user is assumed to have a constraint on the available upload rate; let B_i denote this upper bound for user i . A rate vector is feasible for a user as long as the upload rate is at most equal to the user’s upload capacity. We assume that users do not face any constraint on their download rate; this is consistent with most end user connections today, where upload capacity is far

exceeded by download capacity.¹ Given a vector of prices $\mathbf{p} \gg 0$ (i.e. $p_j > 0$ for $j \in F$), user i can find the upload rate y_i and vector of download rates \mathbf{x}_i that maximize his utility by solving the following optimization problem:

User Optimization:

$$\text{maximize} \quad v_i(\mathbf{x}_i, y_i) \quad (1)$$

$$\text{subject to} \quad \sum_{j \in T_i} x_{ij} \cdot p_j \leq (\max_{j \in S_i} p_j) \cdot y_i; \quad (2)$$

$$y_i \leq B_i; \quad (3)$$

$$y_i \geq 0; \quad x_{ij} \geq 0, \text{ for all } j \in T_i. \quad (4)$$

By assumption, the utility function of a user only depends on his upload rate and not on which files he is uploading. Thus user i will only choose to upload files that have the highest price among all files in S_i . The constraint (2) guarantees that the expenses of a user are at most equal to his revenue from uploading. The constraint (3) guarantees that a user does not upload at a higher rate than his upload capacity B_i . Finally, all rates must be non-negative (constraint (4)). For any price vector $\mathbf{p} \gg 0$, the feasible region of the User Optimization problem is compact and by Assumption 1 the objective function is continuous; thus an optimal solution exists for any price vector $\mathbf{p} \gg 0$. The following lemma captures an important feature of this optimal solution; the proof follows immediately using strict monotonicity of the utility function v_i .

Lemma 1 *If Assumption 1 is satisfied, the budget constraint will bind in the User Optimization Problem for any price vector $\mathbf{p} \gg 0$.*

To simplify our analysis, we also make the following assumption.

Assumption 2 *For every user $i \in U$, the corresponding User Optimization problem has a unique solution (\mathbf{x}_i, y_i) for any price vector $\mathbf{p} \gg 0$.*

For instance, Assumption 2 is satisfied if each utility function is strictly concave, since the feasible region of the optimization problem of each user is convex. Let $x_{ij}(\mathbf{p})$ and $y_i(\mathbf{p})$ be the optimal values of x_{ij} and y_i respectively when the price vector is $\mathbf{p} \gg 0$.

We now define exchange economy [18] and relate it to our model. In an *exchange economy* there is a finite number of agents and a finite number of commodities. Each agent is endowed with a bundle of commodities, and has a preference relation on the set of commodity vectors. Given a price vector, each agent finds a vector of commodities to exchange that maximizes his utility. In particular, if \mathbf{p} is the vector of prices and agent i has endowment \mathbf{w}_i , he sells it at the market and obtains wealth $\mathbf{p} \cdot \mathbf{w}_i$. Then the agent buys goods for his consumption at the same price (he may buy back some of the goods he sold).

A straightforward reformulation reveals that our model shares much in common with a standard exchange economy. Consider the constraints of the user optimization problem (1)-(4). The constraint $y_i \geq 0$ is implied by the other constraints as long as all prices are non-negative. The remaining constraints can equivalently be written as:

¹While in practice a constraint on download rate exists, we remove it for the purposes of analysis since in practice the binding constraint on user behavior is likely to be the upload rate constraint.

$$\begin{aligned} \sum_{j \in T_i} x_{ij} \cdot p_j + (\max_{j \in S_i} p_j) \cdot (B_i - y_i) &\leq (\max_{j \in S_i} p_j) \cdot B_i; \\ B_i - y_i &\geq 0; \\ x_{ij} &\geq 0, \text{ for all } j \in T_i. \end{aligned}$$

This appears much like the optimization that an agent performs in an exchange economy: it is as if agent i has B_i units of his own “good”, priced at $\max_{j \in S_i} p_j$. He can trade this for other goods on the open market at prices \mathbf{p} . With this interpretation, $B_i - y_i$ is the amount of his own good that he chooses to keep. However, notice that this is not a standard exchange economy, as the upload rate is not a true commodity; rather, the commodities are the rates of specific files that are uploaded. Since B_i imposes a *joint* constraint on the upload rates of these files, our model is a generalization of the standard exchange economy. In the following two sections, we adapt some results about exchange economies to our model.

3. COMPETITIVE EQUILIBRIUM

In this section we define competitive equilibrium. In Section 3.1, we then proceed to show that there always exists at least one for the model described in Section 2. In Section 3.2, we give conditions that guarantee uniqueness.

We start by defining the *aggregate excess demand vector*.

Definition 1 Given a vector of prices $\mathbf{p} \gg 0$, a vector $(z_j, j \in F)$ is an aggregate excess demand vector if there exist $y_{ij}, i \in U, j \in F$, such that:

1. $z_j = \sum_{i \in U: j \in T_i} x_{ij}(\mathbf{p}) - \sum_{i \in U} y_{ij}$, for $j \in F$.
2. $\sum_{j \in F} y_{ij} = y_i(\mathbf{p})$, for $i \in U$.
3. $y_{ij} \geq 0$, for $i \in U$ and $j \in F$.
4. $y_{ij} = 0$, if $j \notin \arg \max_{k \in S_i} p_k$.

We denote the set of all excess demand vectors given \mathbf{p} by $\mathbf{z}(\mathbf{p})$.

If $|S_i| = 1$ for all $i \in U$ (i.e., each user has exactly one file available for upload), then for all $i \in U, j \in F$, the required value y_{ij} is uniquely defined for any price vector \mathbf{p} : in particular, the only way to satisfy Conditions 2, 3 and 4 is to set $y_{ij} = y_i(\mathbf{p})$ if $S_i = \{j\}$ and $y_{ij}(\mathbf{p}) = 0$ otherwise. Thus, when $|S_i| = 1$ for all $i \in U$ the excess demand is a *function* of \mathbf{p} . On the other hand, if there are users uploading multiple files, the excess demand is a *correspondence*. In particular, suppose there is some user i with $|S_i| \geq 2$ and choose $j, k \in S_i$ with $j \neq k$. Then, for a price vector \mathbf{p} with $p_k = p_j = \max_{l \in S_i} p_l$, there are multiple ways to choose $(y_{il}, l \in S_i)$ that satisfy Conditions 2, 3 and 4, and thus there are multiple excess demand vectors. Our definition of aggregate excess demand vector ensures that we capture all possible means of dividing the upload rate of user i among available files.

Definition 2 The rate allocation $(x_i^*, i \in U)$ and $(y_i^*, i \in U)$ and the price vector $\mathbf{p}^* \gg 0$ constitute a *competitive equilibrium* if the following conditions are satisfied:

1. Utility maximization: For each user i , (x_i^*, y_i^*) solves the corresponding User Optimization problem for $\mathbf{p} = \mathbf{p}^*$, i.e. $x_{ij}^* = x_{ij}(\mathbf{p}^*)$ and $y_i^* = y_i(\mathbf{p}^*)$.
2. Market Clearing: $\mathbf{0} \in \mathbf{z}(\mathbf{p}^*)$; i.e., the total upload rate y_i can be split among the highest price files in S_i , so that for each file the aggregate excess demand is zero.

Note that because of Assumption 1, at competitive equilibrium all prices are strictly positive; otherwise users would want to download all free files at unboundedly large rates. For this reason, we can restrict competitive equilibria to strictly positive price vectors without loss of generality.

Our goal is to show that a competitive equilibrium exists. We *emphasize that competitive equilibria are desirable because they are Pareto efficient*; this is the content of the first fundamental theorem of welfare economics [18]. However, we do not expect equilibria to exist without any restrictions on the sets S_i and T_i of files being uploaded and downloaded, respectively, by user i . For example, suppose there is a file that some users want to download, but no user has available for upload. Then in general, such a file will have positive demand, while supply will always be zero. Thus the excess demand for such a file will be positive unless its price is sufficiently high. Setting a sufficiently high price is equivalent to considering a system without that file.

To avoid such pathological situations, we introduce a natural diversity assumption. We define the *user-file graph* as the directed graph $G = (V, E)$ with $V = U \cup F$, and $E = \{(i, j) : i \in U, j \in T_i\} \cup \{(j, i) : i \in U, j \in S_i\}$. In other words, G is a bipartite graph where nodes correspond to *users* and *files*. There is a directed edge from a user to each of the files he wants, and there is a directed edge from a file to every user that has it.

Assumption 3 The user-file graph is strongly connected.

If Assumption 3 is not satisfied, then either the system can be decomposed to subsystems that satisfy the Assumption 3, or an equilibrium may not exist. We will therefore assume that Assumption 3 holds.

3.1 Existence of Competitive Equilibrium

We will adapt standard arguments from microeconomics to establish existence of a competitive equilibrium. We begin with the following basic definitions.

Let f be a correspondence defined on a subset $A \subset \mathbb{R}^N$. The correspondence f is *homogeneous of degree zero* if for every $t > 0$, we have $f(tx_1, \dots, tx_N) = f(x_1, \dots, x_N)$. The correspondence f is *convex valued* if $f(x)$ is convex for every $x \in A$. Given the closed set $Y \subset \mathbb{R}^K$, a correspondence $f : A \rightarrow Y$ has a *closed graph* if for any two sequences $x^m \rightarrow x \in A$ and $y^m \rightarrow y$, with $x^m \in A$ and $y^m \in f(x^m)$ for every m , we have $y \in f(x)$. Given the closed set $Y \subset \mathbb{R}^K$, the correspondence $f : A \rightarrow Y$ is *upper hemicontinuous* if it has a closed graph and the images of compact sets are bounded.

The following proposition shows properties of the aggregate excess demand correspondence that are used to prove existence of a competitive equilibrium. The proof is an extension of an argument typically used to prove existence of competitive equilibrium in exchange economies. The key difficulty is in addressing the fact that users may simultaneously upload multiple files; as discussed in Section 2, this feature means our basic model is not a standard exchange economy. Nevertheless, the proof can be appropriately adapted.

Proposition 1 If Assumptions 1, 2 and 3 hold, then the aggregate excess demand correspondence $\mathbf{z}(\cdot)$ defined on $(0, \infty)^F$ satisfies the following properties:

1. For every $\mathbf{p} \gg 0$ and $\mathbf{z} \in \mathbf{z}(\mathbf{p})$, $\mathbf{p} \cdot \mathbf{z} = 0$.
2. $\mathbf{z}(\cdot)$ is convex-valued.
3. $\mathbf{z}(\cdot)$ is homogeneous of degree 0.

4. $z(\cdot)$ is upper hemicontinuous.
5. There is an $s > 0$ such that $z_j > -s$ for any $z \in z(\mathbf{p})$, for every file $j \in F$ and every price vector $\mathbf{p} \gg 0$.
6. If $\mathbf{p}^m \rightarrow \mathbf{p} \neq \mathbf{0}$, $\mathbf{z}^m \in z(\mathbf{p}^m)$ and $p_j = 0$ for some j , then $\max\{z_j^m : j \in F\} \rightarrow \infty$.

Now the existence of a competitive equilibrium follows from standard results in microeconomics; see, e.g., [18], Exercise 17.C.2

Theorem 1 *If Assumptions 1, 2 and 3 hold, then there exists a competitive equilibrium.*

Corollary 1 *If the utility function of each user is strictly concave, and Assumptions 1 and 3 are satisfied, then there exists a competitive equilibrium.*

In Section 5, we assume that each user has a separable utility function, and experiences a cost of uploading that is linear in the upload rate. In this case, the utility function is not strictly concave. The following corollary of Theorem 1 shows existence of a competitive equilibrium for that case.

Corollary 2 *If the utility function of user i is $v_i(\mathbf{x}_i, y_i) = u_i(\mathbf{x}_i) - y_i$, where $u_i(\mathbf{x}_i)$ is continuous, strictly concave, and strictly increasing in each x_{ij} , and Assumption 3 is satisfied, then there exists a competitive equilibrium.*

3.2 Uniqueness of Competitive Equilibrium

We now study uniqueness of the competitive equilibrium. Note that, as is standard, we discuss uniqueness *up to scaling of the price vector*: since z is homogeneous of degree zero, if \mathbf{p}^* is a competitive equilibrium price vector, then so is $t\mathbf{p}^*$. We first define the gross substitutes property.

Definition 3 *The function $z(\cdot)$ has the gross substitutes property if whenever $\mathbf{p}' \gg 0$ and $\mathbf{p} \gg 0$ are such that for some l , $p'_l > p_l$ and $p_k = p'_k$ for $k \neq l$, we have $z_k(\mathbf{p}') > z_k(\mathbf{p})$ for $k \neq l$.*

If the aggregate excess demand is a function that satisfies the gross substitutes property, then there is at most one competitive equilibrium up to scaling of the price vector [18]. In our model, the aggregate excess demand is a function if and only if each user is uploading exactly one file, i.e. $|S_i| = 1$ for all $i \in U$. When some users i have $|S_i| > 1$, the aggregate excess demand is a correspondence, so the preceding result does not apply. In order to adapt that result, we use the following definition.

Definition 4 *The Optimization Problem of user i satisfies the gross substitutes property if whenever $\mathbf{p}' \gg 0$ and $\mathbf{p} \gg 0$ are such that for some l , $p'_l > p_l$ and $p'_k = p_k$ for $k \neq l$, the following conditions hold:*

1. For $l \in T_i$, $x_{ij}(\mathbf{p}') > x_{ij}(\mathbf{p})$ for $j \neq l$, $j \in T_i$ and $y_i(\mathbf{p}') \leq y_i(\mathbf{p})$.
2. If $l \in S_i$ and $p'_l > \max_{k \in S_i} p_k$, then $x_{ij}(\mathbf{p}') > x_{ij}(\mathbf{p})$ for $j \in T_i$.

We interpret this definition as follows. When the price of a file that is relevant to user i increases, user i demands more of all other files he is downloading, and supplies less of the file he is uploading. As one example, it is straightforward to verify that user i 's

optimization problem satisfies gross substitutes if $T_i = \{j\}$ and $v_i(x_{ij}, y) = x_{ij}^\alpha / \alpha - y$, where $0 < \alpha < 1$.

Under a slightly stronger diversity assumption, we can establish the following proposition. The key step in the proof is to show that despite the fact that users may upload multiple files, the monotonicity of excess demand implied by the usual gross substitutes condition continues to hold.

Proposition 2 *If the optimization problem of each user satisfies the gross substitutes property, and $\forall j, k \in F$ there exists $i \in U$ such that $j, k \in T_i$, then there is at most one competitive equilibrium up to scaling of the price vector.*

PROOF. It suffices to show that whenever $\mathbf{p} \gg 0$ and $\mathbf{p}' \gg 0$ are two price vectors that are not collinear, any corresponding aggregate excess demand vectors can not be equal, i.e. $z(\mathbf{p}) \cap z(\mathbf{p}') = \emptyset$. Since z is homogeneous of degree zero, we can assume that $p'_k \geq p_k$ for all k , and $p_l = p'_l$ for some l . Let $S = \{j : p_j = p'_j\}$.

Consider altering the price vector \mathbf{p} to obtain the price vector \mathbf{p}' , by increasing (or keeping unaltered) the price of every file $k \notin S$, one file at a time. As we increase p_k for some $k \notin S$, for every file $j \in S$ (including l), there is a user i who wants both j and k ; i.e., $j, k \in T_i$. The optimization problem of that user satisfies the gross substitutes property, so the total demand (i.e., $\sum_{i:k \in T_i} x_{ik}$) for each file $k \in S$ does not decrease in any step, and if there is a file $k \notin S$ with $p_k < p'_k$, the total demand for file k will strictly increase in at least one step. Thus, the total demand $\sum_{k \in S} \sum_{i:k \in T_i} x_{ik}$ for files in S increases (or remains the same, if $p_k = p'_k$ for all $k \notin S$), while by a similar argument the total supply for files in S decreases (or remains the same).

We now consider $j \in S$ such that $p'_j > p_j$; if no such file exists, then there must be some file $k \notin S$ with $p_k < p'_k$, so for every $\mathbf{w} \in z(\mathbf{p})$ and every $\mathbf{w}' \in z(\mathbf{p}')$, $\sum_{j \in S} w_j < \sum_{j \in S} w'_j$. Thus suppose that $p'_j > p_j$ for some $j \in S$; we increase the price of every such file j from p_j to p'_j , one at a time. By gross substitutes, the total demand for each file in $S - \{j\}$ will strictly increase. On the other hand, each user that was previously uploading either j or some other file in S , will now only upload j , while each user that was only uploading files in $S - \{j\}$ will upload those files at most as much as he was uploading before (again by gross substitutes). Thus the total excess demand for files in $S - \{j\}$, i.e., $\sum_{k \in S - \{j\}} z_k(\cdot)$, will strictly increase. We repeat this procedure for every file $j \in S$ with $p'_j > p_j$. Let $S' = \{j \in S : p'_j = p_j\}$; S' is nonempty since $l \in S'$. Then, for every $\mathbf{w} \in z(\mathbf{p})$ and every $\mathbf{w}' \in z(\mathbf{p}')$, $\sum_{j \in S'} w_j < \sum_{j \in S'} w'_j$, so $z(\mathbf{p}) \cap z(\mathbf{p}') = \emptyset$. \square

4. TÂTONNEMENT PROCESS

In this section we restrict our attention to the case where every user is uploading a single file (i.e., $|S_i| = 1$ for all i), and consider convergence of prices to a competitive equilibrium price vector. We describe a price adjustment mechanism, and show that under some assumptions the rate of convergence of this process will be linear in the number of users. This means that in a large system, the prices will rapidly converge to equilibrium.

When every user is uploading a single file, the aggregate excess demand is a function. A reasonable way to adjust the prices in order to reach a competitive equilibrium is to increase the prices of the files whose excess demand is positive, and decrease the prices of the files whose excess demand is negative. This motivates the *tâtonnement process* [18], where the price adjustment rate is equal

(or in general proportional) to excess demand:

$$\frac{dp_j}{dt} = z_j(\mathbf{p}). \quad (5)$$

The following theorem is a restatement of Proposition 17.H.1 [18] for our model.

Theorem 2 *If $|S_i| = 1$ for all $i \in U$, the gross substitutes property holds for the aggregate excess demand function, and Assumptions 1, 2 and 3 are satisfied, then the relative prices of any solution trajectory of (5) converge to the unique equilibrium (up to scaling of the price vector).*

We next show a result about the rate of convergence of the tâtonnement process. Suppose that a unique competitive equilibrium exists, up to scaling of the price vector. Without loss of generality, we fix a file $f_0 \in F$, and fix $p_{f_0}(t) = 1$ for all times t . This determines the relative values of all other prices at the unique competitive equilibrium; furthermore, the standard tâtonnement dynamics described above will converge to the unique competitive equilibrium price vector where $p_{f_0} = 1$. The following theorem shows that under some assumptions about the structure of the system, the rate of convergence near this equilibrium is linear in N . The proof is straightforward, as the dynamics scale linearly with N .

Theorem 3 *Suppose that $|S_i| = 1$ for all $i \in U$, the gross substitutes property holds for the aggregate excess demand function, and Assumptions 1, 2 and 3 are satisfied. Suppose also that $U = U_1 \cup \dots \cup U_K$ with $U_k \cap U_l = \emptyset$ whenever $k \neq l$, and:*

1. $S_i = S_k, T_i = T_k, v_i(\cdot) = v_k(\cdot)$ and $B_i = B_k, \forall i \in U_k$.
2. $|U_k| = r_k N$ with $r_k \geq 0$, for $k = 1, \dots, K$.

Consider the tâtonnement dynamics with $p_{f_0}(t) = 1$ for all t . If the tâtonnement process converges to some price vector \mathbf{p}^ and $x_{ij}(\mathbf{p}), y_i(\mathbf{p})$, are differentiable at \mathbf{p}^* , the rate of convergence near the equilibrium price vector is linear in N .*

The preceding result assumes that users can be partitioned into identical sets; in this case, the tâtonnement dynamics scale linearly with the number of users. We believe the preceding result can be extended to more general assumptions about the user population; the key requirement is that as the system becomes large, the excess demand level should increase proportionally for any given price level $\mathbf{p} \gg 0$. From a system design point of view, this type of a result suggests that a large peer-to-peer system operating as an exchange economy will have fast convergence in a neighborhood of the equilibrium point.

5. PROPORTIONAL ALLOCATION

Although the tâtonnement process provides a price adjustment mechanism that (under reasonable assumptions) ensures that a competitive equilibrium is reached, it has a serious shortcoming from a system design standpoint: the tâtonnement process does not specify how agents should engage in trade *before* equilibrium is reached. Thus, in addition to adjusting the price vector according to the tâtonnement process (5), our system design should specify a mechanism for allocating rates out of equilibrium.

Mechanisms for exchange out of equilibrium have been proposed for an exchange economy in the economics literature [14], but do not directly apply to our model. These mechanisms work by performing part of the exchange, then updating the endowment of each user. However, in a peer-to-peer system, the endowment of a

user at any given time is determined by the file with the maximum price, and the amount he owns is his upload capacity. Therefore, the amount remains the same even after a user uploads the file once.

We consider an alternate system design: we will ask users to report the *total upload rate* they are willing to allow, and the *proportions* of their budget they wish to spend on downloading various files. For analytical simplicity, we consider the case where each user uploads a single file, i.e. $|S_i| = 1$ for all i . Let $f(i)$ denote the file user i is uploading. Suppose that each user i optimizes with respect to the current prices $\mathbf{p} \gg 0$, and reports his optimal upload rate to the system, $y_i(\mathbf{p})$. If user i is interested in downloading multiple files, i.e. $|T_i| > 1$, then he also reports what *proportion* $\pi_{ij}(\mathbf{p})$ of his budget he wants to spend on each file in $j \in T_i$. In terms of $x_i(\mathbf{p})$ and $y_i(\mathbf{p})$, if $y_i(\mathbf{p}) > 0$, we have:

$$\pi_{ij}(\mathbf{p}) = \frac{p_j x_{ij}(\mathbf{p})}{p_{f(i)} y_i(\mathbf{p})}. \quad (6)$$

However, unless the current price corresponds exactly to a competitive equilibrium, it will not be possible to give to every user the download rates he desires. Informally, we will use the proportions π_{ij} to allocate rates based on the proportion of the budget each agent i intended to spend on downloading the files in T_i ; this is called the *proportional allocation mechanism*.

In order to formally motivate the proportional allocation mechanism, we first consider the outcome at a competitive equilibrium. From Definition 2 we know that the rates $(x_{ij}^*, i \in U, j \in T_i), (y_i^*, i \in U)$ and the price vector \mathbf{p}^* constitute a competitive equilibrium if the following conditions are satisfied.

1. Each user optimizes, i.e. $x_{ij}^* = x_{ij}(\mathbf{p}^*), y_i^* = y_i(\mathbf{p}^*)$ for all $i \in U$, for all $j \in T_i$.
2. The market clears, i.e. $\sum_{i:j \in T_i} x_{ij}^* = \sum_{i:f(i)=j} y_i^*$ for all $j \in F$.

Since we know that it is not possible to satisfy both conditions out of equilibrium, we will relax one of these conditions. If Condition 2 is not satisfied for some file j , then either the total upload rate of the file is strictly less than its total download rate, which is infeasible, or the total upload rate is higher than the total download rate, which means that resources are being wasted. Thus, it is preferable to satisfy Condition 2, and relax Condition 1.

In particular, given prices \mathbf{p} , suppose user i reports his desired upload rate $y_i(\mathbf{p})$ and what proportion of his budget he wants to spend on each $j \in T_i$, π_{ij} ; we do not assume anything about π , other than $\pi_{ij} \geq 0$ and $\sum_{j \in T_i} \pi_{ij} = 1$. Because \mathbf{p} may not be a competitive equilibrium price vector, in general it is not possible to choose download rates at the current prices that ensure each user i spends exactly the desired proportion π_{ij} on file j . Instead, we will use π and $\mathbf{y}(\mathbf{p})$ to compute a *download* rate allocation \hat{x}_i to each user i , together with a *new price vector* $\hat{\mathbf{p}}$ such that each user i earns a budget of $\hat{p}_{f(i)} y_i(\mathbf{p})$, and spends exactly a proportion π_{ij} on file j ; i.e., (6) is satisfied for all i and $j \in T_i$ with $y_i(\mathbf{p}) > 0$. This is a relaxation of Condition 1 above: of course the resulting allocation may not be optimal for each user given the prices $\hat{\mathbf{p}}$; however, the following *budget constraint* will hold:

$$\sum_{j \in T_i} \hat{p}_j \hat{x}_{ij} = \hat{p}_{f(i)} y_i(\mathbf{p}). \quad (7)$$

This ensures every agent has maximally spent their available budget under the new prices $\hat{\mathbf{p}}$; this is a requirement of optimality, cf. Lemma 1.

The existence of such prices $\hat{\mathbf{p}}$ and download rates \hat{x} is summarized in the following proposition.

Proposition 3 Suppose $|S_i| = 1$ for all $i \in U$. Suppose each user i reports an upload rate y_i , and a vector π_i describing the proportion π_{ij} of his eventual budget to be spent on file j . Then there exists a pair \hat{p} and $\hat{x} = (\hat{x}_i, i \in U)$ such that:

1. For each user i and $j \in T_i$, if $\pi_{ij}\hat{p}_{f(i)}y_i = 0$, then $\hat{x}_{ij} = 0$.
2. For each user i and $j \in T_i$, if $\pi_{ij}\hat{p}_{f(i)}y_i > 0$, then $\hat{p}_j\hat{x}_{ij} = \pi_{ij}\hat{p}_{f(i)}y_i$ for all $j \in T_i$.
3. The market clears where possible; i.e., we have:

$$\sum_{i:j \in T_i} \hat{x}_{ij} = \sum_{i:f(i)=j} y_i$$

for all $j \in F$ with $\hat{p}_j > 0$.

Further, the vectors \hat{x}_i are uniquely determined.

PROOF. If we multiply through the third condition by \hat{p}_j , and substitute from the second condition, we obtain:

$$\sum_{i:j \in T_i} \pi_{ij}\hat{p}_{f(i)}y_i = \hat{p}_j \sum_{i:f(i)=j} y_i. \quad (8)$$

Consider a continuous time Markov chain on the state space F , where the transition rate from state (file) j to state (file) k is $Q_{jk} = \sum_{i:f(i)=j} \pi_{ik}y_i$. (Note that $\pi_{ik} = 0$ if $k \notin T_i$.) Let $Q_{jj} = -\sum_{k \neq j} Q_{jk}$. Then (8) can be rewritten as:

$$\sum_{k \in F} Q_{kj}\hat{p}_k = \hat{p}_j \sum_{k \in F} Q_{jk}.$$

Note that these are the balance equations for the continuous time chain, and so at least one nonnegative solution \hat{p} exists. Further, if the communicating classes of Q are $C_1 \cup \dots \cup C_K = F$, then \hat{p} is unique up to scaling by a positive constant on each communicating class C_l .

If $\pi_{ij}\hat{p}_{f(i)}y_i = 0$, then we define $\hat{x}_{ij} = 0$ (Condition 1). Note that if $\hat{p}_{f(i)} = 0$, then $f(i)$ is transient; thus $f(i)$ will have zero mass in any stationary distribution, and thus \hat{x}_{ij} is uniquely determined in this case. On the other hand, suppose $\pi_{ij}\hat{p}_{f(i)}y_i > 0$ for some i and j . Let $k = f(i)$; then $Q_{kj} > 0$, and $\hat{p}_k > 0$. This, together with the balance equations, implies that $\hat{p}_j > 0$. We conclude that there exists a positive value of \hat{x}_{ij} such that Condition 2 in the proposition is satisfied. Further, since \hat{p} is uniquely defined up to scaling on each communicating class, \hat{x}_{ij} is uniquely determined. This completes the proof. \square

The first condition in the preceding proposition ensures that when either a user i is not interested in downloading a file j ($\pi_{ij} = 0$); his upload rate is zero ($y_i = 0$); or the eventual price of the file he is uploading is zero ($\hat{p}_{f(i)} = 0$), then the download rate \hat{x}_{ij} is zero. In all other cases, the download rates \hat{x}_{ij} are uniquely determined by this procedure. Further, this allocation ensures that all users split their budget in accordance with their desired proportions.

In practice, such a mechanism suggests a natural means to adapt prices as well as allocations. In particular, the following corollary ensures that if the upload rates and requested proportions correspond to those in a competitive equilibrium, then the allocation mechanism given in Proposition 3 will yield the competitive equilibrium allocation.

Corollary 3 Suppose $p^* \gg 0$ is a competitive equilibrium, and $y^* = y(p^*)$ and $\pi^* = \pi(p^*)$. Let \hat{p} and \hat{x} the corresponding prices and download rates, respectively, of Proposition 3. Then $\hat{x} = x(p^*)$.

PROOF. It suffices to note that p^* and $x(p^*)$ satisfy Conditions 1-3 of Proposition 3. Since \hat{x} is uniquely determined, it must be the case that $x(p^*) = \hat{x}$. \square

The preceding result demonstrates that the proportional allocation mechanism is a generalization of the competitive equilibrium allocation, to ensure the market clears even out of equilibrium. However, if users anticipate that the market will be cleared using the proportional allocation mechanism, they may not report their true optimal upload rates $y_i(p)$ or desired proportions $\pi_{ij}(p)$; they may have an incentive to try to “game” the system. In this case they will anticipate that prices and rates are chosen using the proportional allocation mechanism, and choose their declarations strategically. In the remainder of this section, we consider a special case of this game, and prove a competitive limit theorem: in the large system limit, it is as if each user optimizes as a price taker.

5.1 Two Files, Two User Types

We consider a system consisting of two files, and two types of users. Users of type 1 have file 1 and want file 2, while users of type 2 have file 2 and want file 1. We assume that there are at least two users of each type. We will use the subscript ki to denote user i of type k . The upload rate constraint for user i of type k is B_{ki} . By x_{ki} and y_{ki} we denote the upload and download rates, respectively, of user i of type k . Throughout the remainder of the section, we make the following assumption about the utility functions.

Assumption 4 The utility of user i of type k when he is downloading at rate $x_{ki} \geq 0$ and uploading at rate $y_{ki} \geq 0$ is $u_{ki}(x_{ki}) - y_{ki}$, where $u_{ki}(x_{ki})$ is continuously differentiable, strictly concave, and strictly increasing.²

In the next section, we characterize competitive equilibria for this system. In Section 5.1.2, we study Nash equilibria of a game where users have utilities that satisfy Assumption 4, and anticipate that prices and allocations are chosen according to the proportional allocation mechanism. We establish a competitive limit theorem: in the large system limit, it is as if each user optimizes as a price taker. In Section 5.2, we specialize further to a case where all users of the same type share the same utility function. This allows us to establish uniqueness of the Nash equilibrium as well, and gives a more precise characterization of the Nash equilibrium rates. Finally, in Section 5.3, we study the efficiency of the rate allocation obtained at a Nash equilibrium.

5.1.1 Competitive Equilibrium

We denote by p_1 the price of file 1, i.e., the file that type 1 users have, and by p_2 the price of file 2, i.e., the file that type 1 users want. Since only relative prices matter, without loss of generality we normalize $p_2 = 1$. User i of type 1 solves the following problem:

$$\begin{aligned} & \text{maximize} && u_{1i}(x_{1i}) - y_{1i} \\ & \text{subject to} && x_{1i} \leq p_1 y_{1i}; \\ & && x_{1i} \geq 0; y_{1i} \leq B_{1i}. \end{aligned}$$

Since the budget constraint will be binding, this problem is equivalent to:

$$\max_{0 \leq y_{1i} \leq B_{1i}} u_{1i}(p_1 y_{1i}) - y_{1i}.$$

²This model can be extended so that different users have different linear costs for uploading: when the utility of user i of type k is $\hat{u}_{ki}(x_{ki}) - c_{ki} \cdot y_{ki}$ where $\hat{u}_{ki}(x_{ki})$ is continuously differentiable, strictly concave, and strictly increasing, the results of this Section hold for $u_{ki}(x_{ki}) = \hat{u}_{ki}(x_{ki})/c_{ki}$.

The optimization problem for a type 2 user is symmetrically defined with p_1 replaced by $1/p_1$. Given price p_1 , the optimality conditions for a user i of type 1 are:

$$p_1 u'_{1i}(0) \leq 1, \text{ if } y_{1i} = 0; \quad (9)$$

$$p_1 u'_{1i}(p_1 y_{1i}) = 1, \text{ if } 0 < y_{1i} < B_{1i}; \quad (10)$$

$$p_1 u'_{1i}(p_1 B_{1i}) \geq 1, \text{ if } y_{1i} = B_{1i}. \quad (11)$$

The optimality conditions for a user of type 2 are symmetrically defined, with p_1 replaced by $1/p_1$. The conditions above give the optimal upload rates $y_{1i}(p_1)$ and $y_{2i}(p_1)$. The optimal download rates are $x_{1i}(p_1) = p_1 y_{1i}(p_1)$ and $x_{2i}(p_1) = (1/p_1) y_{2i}(p_1)$. At a competitive equilibrium, the market must clear: the total upload rate of type 1 users must equal to the total download rate of type 2 users (and vice versa). So, the price vector $(p_1, 1)$ is a competitive equilibrium if:

$$\sum_i x_{1i}(p_1) = \sum_i y_{2i}(p_1), \text{ and} \\ \sum_i y_{1i}(p_1) = \sum_i x_{2i}(p_1).$$

Note that given the relationship between x_{ki} and y_{ki} , each of these conditions implies the other.

We know from Corollary 2 that a competitive equilibrium always exists. The following proposition characterizes the competitive equilibria.

Proposition 4 *If $\sup_i u'_{1i}(0) \cdot \sup_i u'_{2i}(0) \leq 1$, then at any competitive equilibrium $y_{1i} = y_{2i} = 0$ for all i . On the other hand, if $\sup_i u'_{1i}(0) \cdot \sup_i u'_{2i}(0) > 1$, then at any competitive equilibrium there exist i, j such that $y_{1i} > 0$ and $y_{2j} > 0$.*

PROOF. We first show that if $\sup_i u'_{1i}(0) \cdot \sup_i u'_{2i}(0) \leq 1$, there does not exist a competitive equilibrium where some upload rate is strictly positive. If such an equilibrium exists, then at least one user from each type must be uploading at a strictly positive rate. Suppose such an equilibrium exists and let $(p_1, 1)$ be the corresponding price vector. Then, there exist users i and j such that:

$$u'_{1i}(0) > \frac{1}{p_1}, \text{ and } u'_{2j}(0) > p_1.$$

By multiplying the two inequalities, we see that the assumption $\sup_i u'_{1i}(0) \cdot \sup_i u'_{2i}(0) \leq 1$ is contradicted.

Now we assume that $\sup_i u'_{1i}(0) \cdot \sup_i u'_{2i}(0) > 1$. Suppose that there exists a competitive equilibrium where $y_{ki} = 0$ for all k and i . Let $(p_1, 1)$ be the corresponding price vector. Then, for all i ,

$$u'_{1i}(0) \leq \frac{1}{p_1}, \text{ and } u'_{2i}(0) \leq p_1.$$

If we take the supremum in both inequalities and multiply the result, we see that the assumption $\sup_i u'_{1i}(0) \cdot \sup_i u'_{2i}(0) > 1$ is contradicted. \square

5.1.2 Nash Equilibrium

We use the proportional allocation mechanism to clear the market out of equilibrium. The results in Proposition 3 are simplified in this case, because there are only two files ($|F| = 2$) and each user is downloading a single file. Thus users only report upload rates; it is clear that they will spend their entire budget on the single file they wish to download. Let y_{ki} be the upload rate that user i of type k reports and $Y_k = \sum_i y_{ki}$. If $Y_1 > 0$ and $Y_2 > 0$, it is straightforward to check that the proportional allocation mechanism will use

the following price to clear the market:

$$\hat{p}_1 = \frac{Y_2}{Y_1}. \quad (12)$$

If either $Y_1 = 0$ or $Y_2 = 0$, then all agents receive zero download rate. When users anticipate that the price to clear the market will be set in this way, they play a game, where the strategy is the declared upload rate. The strategy space of user i of type 1 is $[0, B_{1i}]$. The payoff of user i of type 1 is:

$$\Pi_{1i}(y_{1i}) = \begin{cases} u_{1i}(y_{1i} Y_2 / Y_1) - y_{1i}, & \text{if } Y_1 > 0; \\ u_{1i}(0), & \text{if } Y_1 = 0. \end{cases} \quad (13)$$

If $Y_1 - y_{1i} = \sum_{j \neq i} y_{1j} > 0$, the preceding payoff is continuous and differentiable on $[0, B_{1i}]$. A symmetric expression holds for users of type 2.

We first observe that $y_{1i} = y_{2i} = 0$ for all i is a Nash equilibrium. In particular, if $Y_2 = 0$, the optimal upload rate for any type 1 user is zero, and symmetrically, if $Y_1 = 0$, the optimal upload rate of any type 2 user is zero. However, such a Nash equilibrium is degenerate; it exploits the fact that the system exhibits a strong complementarity between users. Such a situation will be trivially avoided if a small amount of upload rate of each type of file is always available.

Now suppose that $Y_2 > 0$ and $Y_1 - y_{1i} = 0$. Then for any $y_{1i} > 0$, the utility of user i is $u_{1i}(Y_2) - y_{1i}$, while if $y_{1i} = 0$ his utility is $u_{1i}(0)$; in this case his utility is discontinuous, and no best response exists for user i . Thus there does not exist an equilibrium where $Y_1 - y_{1i} = 0$ and $Y_2 > 0$. A symmetric argument shows that there does not exist an equilibrium where, for some user i of type 2, $Y_2 - y_{2i} = 0$ and $Y_1 > 0$. Thus in searching for nonzero Nash equilibria, we can assume that $Y_1 - y_{1i} > 0$ and $Y_2 - y_{2i} > 0$ for all users i of types 1 and 2, respectively.

When $Y_1 - y_{1i} > 0$ and $Y_2 > 0$, the optimality conditions for user i of type 1 become:

$$u'_{1i}(0) \leq \frac{Y_1}{Y_2}, \text{ if } y_{1i} = 0; \quad (14)$$

$$u'_{1i} \left(y_{1i} \frac{Y_2}{Y_1} \right) \left(1 - \frac{y_{1i}}{Y_1} \right) = \frac{Y_1}{Y_2}, \text{ if } 0 < y_{1i} < B_{1i}; \quad (15)$$

$$u'_{1i} \left(B_{1i} \frac{Y_2}{Y_1} \right) \left(1 - \frac{B_{1i}}{Y_1} \right) \geq \frac{Y_1}{Y_2}, \text{ if } y_{1i} = B_{1i}. \quad (16)$$

Symmetric optimality conditions hold for a user i of type 2, when $Y_2 - y_{2i} > 0$ and $Y_1 > 0$.

Let N_1 and N_2 be the number of type 1 and type 2 users respectively. The following theorem shows that under reasonable conditions, a non-zero Nash equilibrium exists. This result is not straightforward, as the payoff function is typically discontinuous at $\mathbf{y} = 0$, so a direct fixed-point argument does not suffice. We instead use a perturbation approach: we introduce two "virtual" users who upload at a rate ε for each file. In this regime a Nash equilibrium always exists; and by considering the limit as ε approaches zero we are able to establish existence of a Nash equilibrium for the original game.

Theorem 4 *If Assumption 4 is satisfied, $u'_{1i}(0) > N_1 / (N_1 - 1)$ for all i , and $u'_{2i}(0) > N_2 / (N_2 - 1)$ for all i , then there exists a Nash equilibrium $(\mathbf{y}_1, \mathbf{y}_2)$ at which not all rates are equal to zero.*

We now develop a *competitive limit*, where the number of users of each type becomes large. Suppose that $N_1, N_2 \rightarrow \infty$, and consider a sequence of Nash equilibria \mathbf{y}^N indexed by $N = N_1 + N_2$; by taking subsequences if necessary, we can assume the Nash

equilibria converge, say to \mathbf{y} . Let $Y_k^N = \sum_i y_{ki}^N$. Suppose that $y_{ki}^N/Y_k^N \rightarrow 0$ for all users i of type k , but that $Y_2^N/Y_1^N \rightarrow p_1 \in (0, \infty)$; we normalize $p_2 = 1$. Under these assumptions, since the optimality conditions (14)-(16) are continuous, they become identical to the optimality conditions (9)-(11) for a competitive equilibrium. Thus informally, we expect that the Nash equilibrium rates should approach competitive equilibrium rates.

Formally, recall that we define $x_{ki}(p_1)$ and $y_{ki}(p_1)$ as the optimal solutions for a price taking user (i.e., a user solving (1)-(4)), given a price p_1 . We then have the following theorem.

Theorem 5 *Let $N = N_1 + N_2$ be the total number of users. Suppose that as $N \rightarrow \infty$, both $N_1 \rightarrow \infty$ and $N_2 \rightarrow \infty$. Suppose that Assumption 4 holds for the utility function of each user, $\sup_i B_{ki} < \infty$, $\sup_i u'_{ki}(0) < \infty$, and $\inf_i u'_{ki}(0) > 1$ for $k = 1, 2$. Let \mathbf{y}^N denote a nonzero Nash equilibrium when $N = N_1 + N_2$ users are in the system, and let $p_1^N = Y_2^N/Y_1^N$, where $Y_k^N = \sum_i y_{ki}^N$. Then:*

1. $0 < \inf_N p_1^N$ and $\sup_N p_1^N < \infty$.
2. For all i and k , $y_{ki}^N/Y_k^N \rightarrow 0$ as $N \rightarrow \infty$, while $Y_k^N \rightarrow \infty$ as $N \rightarrow \infty$.
3. Any limit point (p_1, \mathbf{y}) of the sequence (p_1^N, \mathbf{y}^N) satisfies the competitive equilibrium optimality conditions (9)-(11).

PROOF. It suffices to show that $Y_k^N \rightarrow \infty$ as $N \rightarrow \infty$, for $k = 1, 2$. In this case, the second property of the theorem holds simply because each y_{ik} is bounded above by the upload rate constraint. Further, the optimality conditions (14)-(16) will imply both the first and third properties of the theorem: the first property follows because p_1^N cannot go to zero nor become unbounded if some users have positive rate; and the third property follows because the Nash optimality conditions are continuous as long as $Y_1 > 0$ and $Y_2 > 0$.

Suppose that Y_1^N remains bounded as $N \rightarrow \infty$; in this case, taking subsequences if necessary, we can assume that $Y_1^N \rightarrow c_1 < \infty$ as $N \rightarrow \infty$. Suppose also that $\sup_N Y_2^N = \infty$; then for at least one type 2 user i the analogous optimality conditions (15) or (16) hold. Taking subsequences if necessary, we have as $N \rightarrow \infty$, $Y_2^N/Y_1^N \rightarrow \infty$ and $Y_1^N/Y_2^N \rightarrow 0$, which contradicts either (15) or (16) for type 2 user i and the assumption $\sup_i u'_{ki}(0) < \infty$. Thus Y_2^N remains bounded as $N \rightarrow \infty$, so taking subsequences again if necessary, we assume that $Y_2^N \rightarrow c_2 < \infty$ as $N \rightarrow \infty$.

Without loss of generality, we can assume that $c_1/c_2 \leq 1$; otherwise we apply the subsequent argument to type 2 users. Again taking subsequences if necessary, we assume y_{ki}^N converges to y_{ki} for all i and k ; this is straightforward, as the strategy space of each user is compact. Now since Y_1^N remains bounded for large N , there exists at least one user i of type 1 who has $y_{1i} = 0$. For such a user, taking limits in (14)-(15), we conclude we must have:

$$u'_{1i}(0) \leq \frac{c_1}{c_2} \leq 1.$$

This contradicts our assumption that $\inf_i u'_{ki}(0) > 1$. Thus we conclude that in fact $Y_k^N \rightarrow \infty$ as $N \rightarrow \infty$, for $k = 1, 2$, as required; this establishes the theorem. \square

The preceding theorem shows that in the large system limit, it is as if each user optimizes as a price taker. Observe that from the proof, in the limit we have infinite upload rates for both types of files; thus we cannot directly interpret the limit point as a competitive equilibrium. However, we can make the following precise statement: asymptotically, users choose upload rates that are nearly

equal to their optimal upload rate if they were acting as price takers. One way to interpret such a theorem is that in large peer-to-peer systems, fully strategic behavior by the users will not ultimately cause large deviations from competitive equilibrium behavior.

5.2 Homogeneous Utilities

In this section we consider a system where all users have the same utility functions (i.e., $u_{1i}(\cdot) = u_{2i}(\cdot) = u(\cdot)$) and the same rates (i.e., $B_{1i} = B_{2i} = B$). Moreover, we assume that there is the same number of type 1 and type 2 users, denoted N . This is a special case of the model analyzed in the previous subsection.

Throughout this section, to avoid boundary conditions, we will make the following additional simplifying assumption about the utility functions; the analysis can be extended to study the case where the assumption does not hold, but the main insights remain the same.

Assumption 5 *The function $u(\cdot)$ satisfies $u'(x) \rightarrow \infty$ as $x \rightarrow 0$, and $u'(x) \rightarrow 0$ as $x \rightarrow \infty$.*

Note that under this assumption, if u is strictly concave, then $u'^{-1}(x)$ is well defined for $x \in (0, \infty)$. In the next two sections, we study competitive equilibria and Nash equilibria of this model, respectively; our key result is that under the homogeneity assumption, the system has a unique Nash equilibrium.

5.2.1 Competitive Equilibrium

We show that under Assumptions 4 and 5, the price $p_1 = 1$ is always a competitive equilibrium of this economy. First suppose that $u'(B) < 1$. Since $u'(\cdot)$ is continuous, there is a $y \in (0, B)$ such that $u'(y) = 1$. Then, when $p_1 = 1$, a user of type 1 will choose to upload and download $u'^{-1}(1)$, and the same for all users of type 2. Since the total upload and download rates of a file are equal, this is a competitive equilibrium. On the other hand, if $u'(B) \geq 1$, then when $p_1 = 1$, all users will choose to upload B . In this case the upload rate constraint binds, and we again have a competitive equilibrium.

In the special case we are studying here, uniqueness of the competitive equilibrium can be guaranteed via a simple condition on the utility function u .

Proposition 5 *Suppose Assumptions 4 and 5 are satisfied, and $u(\cdot)$ is twice differentiable. Then, the following are equivalent:*

1. For all $B > 0$, the Optimization Problem of each user satisfies the gross substitutes property.
2. $pu'^{-1}(p)$ is nonincreasing on $(0, \infty)$.
3. $xu'(x)$ is nondecreasing.

Under these conditions, the competitive equilibrium is unique.

5.2.2 Nash Equilibrium

The analysis of Nash equilibria is simplified when the system is homogeneous, due to the following lemma.

Lemma 2 *If $u(\cdot)$ is a strictly concave function, then users of the same type will have the same upload rate at any Nash equilibrium.*

For the remainder of this section, we suppose that Assumptions 4 and 5 hold. If $u'(x) \rightarrow \infty$ as $x \rightarrow 0$, the optimality condition (14) will never apply. Let y_1, y_2 be the rates at which users of type 1 and type 2 upload, respectively, at a Nash equilibrium; and recall that N denotes the number of users of each type. If $y_1 > 0$ and

$y_2 > 0$, the optimality conditions (15) and (16) can be equivalently written:

$$u'(y_2) = \frac{N}{N-1} \frac{y_1}{y_2}, \text{ if } 0 < y_1 < B; \quad (17)$$

$$u'(y_2) \geq \frac{N}{N-1} \frac{B}{y_2}, \text{ if } y_1 = B; \quad (18)$$

a symmetric condition holds for users of type 2.

If $u'(0) \leq N/(N-1)$, then $y_1 = y_2 = 0$ is the unique Nash equilibrium. To show this, we first observe that if $y_1 = 0$, then $y_2 = 0$ (and vice versa), i.e., there can not be a Nash equilibrium at which only users of one type are uploading at strictly positive rates. Now suppose there exists a Nash equilibrium at which $y_1 > 0$ and $y_2 > 0$. Then, assuming that $u(\cdot)$ is strictly concave,

$$\left(\frac{N}{N-1}\right)^2 \geq (u'(0))^2 > u'(y_1)u'(y_2) \geq \left(\frac{N}{N-1}\right)^2,$$

a contradiction.

If $u'(0) > N/(N-1)$, then there exists a Nash equilibrium with $y_1 > 0$ and $y_2 > 0$: for example, $y_1 = y_2 = \min(u'^{-1}(N/(N-1)), B)$. When the upload rates are positive we define the Nash price as $p^{NE} = y_2/y_1$. Thus $p^{NE} = 1$ is a possible Nash price and we know that $p^* = 1$ is a competitive equilibrium price. In particular, if $u'(0) > N/(N-1)$, there exists a Nash equilibrium with the same price as the unique competitive equilibrium. Theorem 5 does not apply here, because of Assumption 5. However, by Lemma 2, $y_{ki}/Y_k = 1/N \rightarrow 0$ as $N \rightarrow \infty$, and thus any limit point (p_1, \mathbf{y}) of the sequence (p_1^N, \mathbf{y}^N) satisfies the competitive equilibrium optimality conditions. Moreover, it can be shown that the rates of any sequence of Nash equilibria converge to the rates of a competitive equilibrium.

The following proposition is our key result for the model with homogeneous users: we show that if gross substitutes holds, there exists a unique Nash equilibrium where the upload rates are strictly positive. The proof uses the characterization of gross substitutes shown in Proposition 5.

Proposition 6 *Suppose that Assumptions 4 and 5 hold and $u(\cdot)$ is twice differentiable. Then, if the Optimization Problem of each user satisfies the gross substitutes property, there is a unique Nash equilibrium with strictly positive rates. At the equilibrium $y_1 = y_2 = u'^{-1}(N/(N-1))$.*

The Nash equilibrium is not always unique, as the following example shows. Let $u(x) = -1/x$, so that $pu'^{-1}(p) = \sqrt{p}$ is strictly increasing. The optimality conditions give $y_1 \cdot y_2 = (N-1)/N$, i.e. there are infinitely many Nash equilibria. In particular, the set of Nash equilibria is $\{(y_1, y_2) : 0 \leq y_1 \leq B, 0 \leq y_2 \leq B, y_1 y_2 = (N-1)/N\}$. For this utility function there are also infinitely many competitive equilibria: since $(1/p)u'^{-1}(1/p) = u'^{-1}(p)$ for every p , any price is a competitive equilibrium.

5.3 Efficiency

We consider a Nash equilibrium of the game that results from the proportional allocation mechanism at which not all rates are zero; i.e., $Y_1 > 0$ and $Y_2 > 0$. Notice that when type 1 users choose their optimal upload rate, they take Y_2 as given. Thus, we can interpret the rates y_{1i} reported by type 1 users as a Nash equilibrium to the following auction game. Suppose that the available upload rate of file 2 is fixed and equal to Y_2 . Type 1 users submit bids to acquire a share of the available file transfer rate for file 2; each user has to pay his bid, and is allocated a download rate proportional to his bid.

In [17], it is shown for this game that if Assumption 4 is satisfied, and for all i $u_{1i}(0) \geq 0$, then:

$$\sum_{i \in U} u_{1i} \left(y_{1i} \frac{Y_2}{Y_1} \right) \geq \frac{3}{4} \max_{\sum_{i \in U} \bar{x}_{1i} = Y_2} \sum_{i \in U} u_{1i}(x_{1i}).$$

A symmetric result holds for type 2 users. This result shows that given the available upload rate of file 2, it is nearly efficiently shared among type 1 users; and similarly for type 2 users.

On the other hand, Nash equilibria need not be Pareto efficient. Suppose that Assumption 4 holds, users are homogeneous, and $1 < u'(0) \leq N/(N-1)$, where N is the number of users in each type. Then at any competitive equilibrium, all upload and download rates will be strictly positive, while at any Nash equilibrium all rates will be zero. Since each user has the option of uploading and downloading zero in the competitive equilibrium, this shows that each user is strictly worse off at the Nash equilibrium.

6. CONCLUSION

This paper presents a model of peer-to-peer filesharing as an exchange economy. Our formulation is novel, and the approach not only controls free-riding, but also ensures that users that provide the most benefit to the system are appropriately rewarded. In this section we briefly comment on two issues that require additional attention: first, the use of a centralized server; and second, the lack of “bankable” currency.

In our model, we use a central server to update prices. However, the information stored at this server scales only with the number of files in circulation—not with the number of users. In this sense the system is highly scalable. Nevertheless, our model has ignored issues of distributed query processing and the impact of network structure on price dynamics; such issues remain fruitful avenues for future research.

A potential implementation problem is that users are not allowed to store currency. This can be problematic, as users cannot leverage valuable uploads today to finance downloads tomorrow. This problem might be resolved by extending our model to allow users to store currency; from a game theoretic standpoint, however, this would require analyzing users’ net present value in finding equilibria. For this reason such a model is a substantial departure from the framework in the current paper, and also remains an open direction.

Acknowledgments

This work was partially supported by the National Science Foundation under grant 0428868. We also gratefully acknowledge helpful conversations with Ashish Goel and Richard Zeckhauser.

7. REFERENCES

- [1] E. Adar and B. Huberman. Free riding on Gnutella. *First Monday*, 5(10), 2000.
- [2] K. 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] C. Aperjis and R. Johari. A peer-to-peer system as an exchange economy. 2006. Stanford University Management Science and Engineering Working Paper.
- [4] K. Arrow, H. D. Block, and L. Hurwicz. On the stability of the competitive equilibrium II. *Econometrica*, 27(1):82–109, 1959.
- [5] C. Buragohain, D. Agrawal, and S. Suri. A game theoretic framework for incentives in P2P systems. In *Third*

International Conference on Peer-to-Peer Computing, pages 48–56, 2003.

- [6] B. Cohen. Incentives build robustness in BitTorrent. In *Workshop on Economics of Peer-to-Peer Systems*, 2003.
- [7] D. Dutta, A. Goel, R. Govindan, and H. Zhang. The design of a distributed rating scheme for peer-to-peer systems. In *Workshop on Economics of Peer-to-Peer Systems*, 2003.
- [8] M. Feldman, C. Papadimitriou, J. Chuang, and I. Stoica. Free-riding and whitewashing in peer-to-peer systems. In *PINS '04: Proceedings of the ACM SIGCOMM workshop on Practice and theory of incentives in networked systems*, pages 228–236, New York, NY, USA, 2004. ACM Press.
- [9] E. J. Friedman, J. Y. Halpern, and I. Kash. Efficiency and Nash equilibria in a scrip system for P2P networks. In *ACM Conference on Electronic Commerce*, 2006.
- [10] E. J. Friedman and P. Resnick. The social cost of cheap pseudonyms. *Journal of Economics and Management Strategy*, 10(2):173–199, 2001.
- [11] 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.
- [12] R. Guha, R. Kumar, P. Raghavan, and A. Tomkins. Propagation of trust and distrust. In *International World Wide Web Conference (WWW2004)*, pages 403–412, 2004.
- [13] M. Gupta, P. Judge, and M. Ammar. A reputation system for peer-to-peer networks. In *NOSSDAV '03: Proceedings of the 13th international workshop on Network and operating systems support for digital audio and video*, pages 144–152, New York, NY, USA, 2003. ACM Press.
- [14] F. Hahn and T. Negishi. A theorem on non-tatonnement stability. *Econometrica*, 30(3):463–469, 1962.
- [15] D. Hughes, G. Coulson, and J. Walkerdine. Free riding on Gnutella revisited: The bell tolls? *IEEE Distributed Systems Online*, 6(6), 2005.
- [16] J. Ioannidis, S. Ioannidis, A. Keromytis, and V. Prevelakis. Fileteller: Paying and getting paid for file storage. In *Financial Cryptography*, pages 282–299, 2002.
- [17] R. Johari and J. N. Tsitsiklis. Efficiency loss in a network resource allocation game. *Mathematics of Operations Research*, 29(3):407–435, 2004.
- [18] A. Mas-Colell, M. D. Whinston, and J. R. Green. *Microeconomic Theory*. Oxford University Press, Oxford, United Kingdom, 1995.
- [19] V. Vishnumurthy, S. Chandrakumar, and E. G. Sirer. Karma: A secure economic framework for peer-to-peer resource sharing. In *Workshop on Economics of Peer-to-Peer Systems*, 2003.