Daan van de Meeberg

Game Theory in Peer-to-Peer File Sharing Networks

Master thesis, defended on June 18, 2013

Thesis advisor: dr. F.M. Spieksma

Specialisation: Applied Mathematics

Mathematisch Instituut, Universiteit Leiden

Contents

Preface

Since the creation of the internet, people have been using it to share with each other. Information, resources, files, etc. One of the most interesting ways in which this is done are so called Peer-to-Peer networks, where participants act simultaneously as receiver and distributor. These networks have sprung up in the last fifteen years or so and are becoming more popular by the day. This brings with it, the classic conflict, where participants like to get as much as they can from the network, while giving away as little as possible. Of course, if everybody would act this way, nobody would be left to act as distributor and the whole network would collapse.
This principle, where individuals act rational and in their self interest, while knowingly going against the populations long-term best interest, is an economical concept, often referred to as the *tragedy of the commons*.
In this thesis, I will examine this conflict from a game theoretical perspective. In the first chapter I will discuss different measuring technique to determine the quality of such systems and their Nash equilibria. Then after a short introduction to Peer-to-Peer networks in the second chapter, I'll use those methods to examine existing Peer-to-Peer models and investigate the effect that implementing incentive mechanisms have on the quality of the equilibria produced by these systems.
In the third chapter I'll study a model focused on maximizing the available resources in the system. Then in chapter four, the focus shifts from resources to burden sharing. Here we find that sometimes in order to increase the number of contributors to the system, we have to incur a social loss: the price we pay for selfish behavior.
Finally in chapter 5 I'll study the effect caused by incomplete information to the efficiency of a system.

Chapter 1

Qualifying Efficiency in equilibria

In studying noncooperative games, we often find that the outcome, resulting from players only looking out for themselves, i.e. maximizing their own goals, is not nearly as good as one which is reached by some dictatorial designer. More generally speaking, we see that (Nash) Equilibria are sometimes inefficient. We can analyze inefficiency either 'qualitatively' or 'quantitatively'. Qualitatively meaning that cost or payout of all players is abstract and only tells us the players preferences between outcomes. For those models we say that an outcome is *Pareto inefficient* if there is another outcome in which all players enjoy a better result (be it cost or payout).
When payout or cost have a more concrete interpretation, like money, we can consider the quantitative (in)efficiency by considering an objective function, expressing the "social cost" of an outcome. The most common of these functions are the *utilitarian* and the *egalitarian*.
We define the *utilitarian objective function* as the sum of all players' costs.
The *egalitarian objective function* is defined as the maximum cost over all players.
Perhaps the best known game theoretical model is the *Prisoner's Dilemma*, where two prisoners are on trial for a crime and both of them have the option to either stay quiet or confess. If they both stay quiet, the prosecutor will not be able to prove the charges and they'll only serve short sentences, say 2 years. If one of them confesses, while the other stays quiet, the confessor's term will be reduced to 1 year, while the other will get a full sentence of 5 years. If they both confess, they'll both get 4 years (they get a year off for cooperating). Now the unique Nash Equilibrium is reached when they both confess, leading to both of them having to serve 4 years. However, for both functions (utilitarian or egalitarian), a cooperated effort leading to both of them keeping their mouth shut, leads to a better outcome for both players.
By using objective functions, we are able to quantify the inefficiency of equilib-

ria. By doing that, we can distinguish between models where game-theoretical equilibria are guaranteed to be optimal or close to it, and other models. In other words, we differentiate between models where selfish behaviour doesn't have severe consequences and models where it does.
Such guarantees are particularly useful in computer science applications, where reaching an optimum can be near impossible or extremely expensive (in terms of time and/or CPU usage) [12].

1.1 Measuring efficiency

We can measure the inefficiency of the equilibria of a game in many different ways. All these measures are roughly defined as the ratio between the objective value of the optimal solution and that of some equilibrium. To specify, we have to set some basic modeling definitions.
First of all, we have to clearly express a players payoff or cost. For the topics later to be discussed, we will set these either as some concrete payoff (money earned) that players aim to maximize, or a cost that the players try to minimize (network delay for example). We will call an outcome optimal, if it optimizes the chosen objective function.
The second thing to consider is what kind of objective function to use to compare different outcomes. As mentioned before, the most common are the *utilitarian* and the *egalitarian*. Here we'll focus on the utilitarian objective function.
Next we need to define what we mean by *approximately optimal*. We use the ratio between the objective function of the optimal solution and a given outcome as a way to quantify how close the given outcome approximates the optimal solution. Because we will only use non-negative objective functions, this ratio will always be non-negative. For minimization objectives this ratio will be at least one and for maximization objectives the ratio will be at most one. In both cases, a value close to 1 indicates that the given outcome is approximately optimal.
We also have to define what we mean by an equilibrium. There are many different equilibrium concepts, but here we will focus only on Nash equilibria. In this situation we assume that local uncoordinated optimization by players will lead to an equilibrium. This isn't always the case, as some important classes of network games do not guarantee such convergence.
Finally we have to decide, in case of multiple equilibria, which one to consider. This can vary given the choice of objection function and the equilibrium concept.

We'll now introduce the two most commonly used measures.

1.1.1 Price of Anarchy

The *Price Of Anarchy* (POA), the most popular measure of the inefficiency of equilibria, takes a worst-case approach. It is defined as the ratio between the worst objective function value of an equilibrium of the game and that of an optimal outcome [13]. By using this measurement, we try to find games with a price of anarchy close to one, meaning that all possible equilibria are good approximations of the optimal outcome. In other words: selfish behavior is benign and we don't benefit much from a dictatorial control over players, which can be costly or even infeasible. Let $o(x)$ be the objective function, o_{opt} the optimal outcome and $Nash$ the set of all Nash equilibria. This gives the following definition.

Definition 1.1. *The **Price of Anarchy** POA is defined as the worst ratio between the objective function values of the Nash equilibria and the optimal outcome.*
In formula, we distinguish between two scenarios.
When minimizing cost,

$$POA = \max_{x \in Nash} \frac{o(x)}{o_{opt}};$$

when maximizing payoff,

$$POA = \max_{x \in Nash} \frac{o_{opt}}{o(x)}.$$

1.1.2 Price of Stability

Because a game with multiple equilibria gets a large price of anarchy even if only one of it's equilibria is highly inefficient, the *Price of Stability* (POS) is designed to distinguish between games in which all equilibria are inefficient and those games where only a few of the equilibria are inefficient. The price of stability is defined as the ratio between the best objective function value of one of its equilibria and that of an optimal outcome (first studied by Schultz and Stier Moses in [18], term coined in [2]).
Again, let $o(x)$ be the objective function, o_{opt} the optimal outcome and $Nash$ the set of all Nash equilibria. This gives the following definition.

Definition 1.2. *The **Price of Stability** POS is defined as the best ratio between the objective function value of the Nash equilibria and the optimal outcome.*
In formula, we distinguish between two scenarios.
When minimizing cost,

$$POS = \min_{x \in Nash} \frac{o(x)}{o_{opt}};$$

when maximizing payoff,

$$POS = \min_{x \in Nash} \frac{o_{opt}}{o(x)}.$$

In the POS as well as the POA, the closer the value is to 1, the better.
In situations where there is only one equilibrium, the Price of Stability and the Price of Anarchy are obviously the same. In games with multiple equilibria

however, there can be big differences. The bound of the Price of Stability is always as least as good (i.e. low) as the Price of Anarchy, but can be much closer to 1. This also means that a bound on the Price of Stability provides a much weaker guarantee than a bound on the Price of Anarchy.
Even so, there are still some important reasons for us to examine the Price of Stability. The most practical one is that in some applications a non-trivial bound is only possible for the price of stability (as we'll see in chapters 3 and 4). The other one is that in many network games, the price of stability has a natural interpretation. This is the case if we look at the outcome as initially designed by a central authority, for selfish players to use. Then that authority can propose the best equilibrium as a logical solution.
In many network applications, players are not truly independent, they interact with an underlying protocol. A central entity proposes a solution, which the players can than either accept or defect from. The Price of Stability measures such protocols.
Besides quantifying the inefficiency by the best or the worst equilibria, a third option would be to analyze an "average" equilibrium. However, such analysis is extremely difficult to define and has therefore not yet been successfully used to study the inefficiency of equilibria [12].
Even in the simplest of games, equilibria can be arbitrarily inefficient. Just look at the Prisoner's Dilemma and let the cost for both players in the Nash Equilibrium tend to infinity. Now for any reasonable objective function the objective function value of the Nash Equilibrium is arbitrarily larger than that of the optimal solution.
The inefficiency of equilibria cannot be bounded in general, as I'll show in the next example. Therefore a natural goal is to identify classes of games in which equilibria are guaranteed to be approximately optimal. This is, luckily, the case for a lot of fundamental network models [12].
Next we will look at a couple of examples to illustrate the inefficiency of equilibria.

1.2 Pigou's Example

This example was first discussed by an economist, Pigou, in 1920 [15]. We have a simple network with two disjoint edges connecting a source s and a destination t. Each edge has its own cost function $c(x)$, describing the cost (or travel time) incurred by users of the edge, as a function of the amount of traffic on that edge. As shown in figure 1.1, the upper edge has a constant cost function $c(x) = 1$, representing a route that is relatively long, but immune to congestion. The cost of the lower edge $c(x) = x$ (for $d = 1$), is proportional to the amount of traffic on that edge. The lower edge is cheaper than the upper edge if and only if there is less than one unit of traffic on that edge.

We now suppose there is *one* unit of traffic, representing a very large population of players and each player chooses independently between both edges. If each player aims to minimizes his cost, the lower edge is clearly a dominant

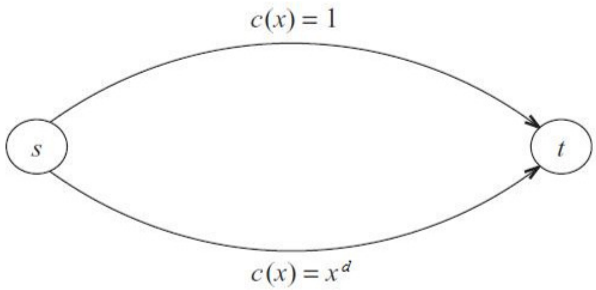

Figure 1.1: Pigou's example

strategy. In the only Nash equilibrium of this model, everyone chooses the lower edge and all of them incur one unit of cost.
For the objective function we take the average cost incurred by players. In this equilibrium the average cost is one. It is easy enough to see that the optimal solution is to distribute the traffic equally over both edges. The objective function is $o(x) = x^2 + (1-x)$ where x is the amount of traffic going over the bottom edge. $o'(x) = 2x - 1$ which is equal to 0 at $x = \frac{1}{2}$. So half of the traffic incurs cost $\frac{1}{2}$ and the other half incurs cost 1. That means the average cost of traffic in this optimal outcome is $\frac{3}{4}$ and thus both the Price of Anarchy and the Price of Stability are equal to $\frac{4}{3}$.
Now we generalize Pigou's example, such that the cost function of the upper edge stays the same, but the cost function of the lower edge changes to $c(x) = x^d$ with $d \geq 1$.
We keep the objective function the same and there is still just one unit of traffic. Clearly choosing the lower edge is still the dominant strategy, so we still have an unique equilibrium with average cost 1. Our objective function has now changed to $o(x) = x^d \cdot x + (1-x) \cdot 1$. Again we find the minimum by setting the derivative equal to zero.

$$\begin{aligned} o'(x) &= (d+1)x^d - 1 \\ 0 &= (d+1)x^d - 1 \\ x^d &= \frac{1}{d+1} \\ x &= \sqrt[d]{\frac{1}{d+1}}. \end{aligned}$$

So we get an optimal solution of

$$\begin{aligned} o_{opt} &= \sqrt[d]{\frac{1}{d+1}} \cdot \left(\sqrt[d]{\frac{1}{d+1}}\right)^d + \left(1 - \sqrt[d]{\frac{1}{d+1}}\right) \\ &= \sqrt[d]{\frac{1}{d+1}} \cdot \frac{1}{d+1} + 1 - \sqrt[d]{\frac{1}{d+1}} \end{aligned}$$

$$= 1 - \sqrt[d]{\frac{1}{d+1}} \cdot \Big(1 - \frac{1}{d+1}\Big)$$
$$= 1 - \sqrt[d]{\frac{1}{d+1}} \cdot \Big(\frac{d}{d+1}\Big).$$

This gives us

$$\textit{Price of Anarchy} = \frac{1}{1 - \sqrt[d]{\frac{1}{d+1}} \cdot (\frac{d}{d+1})}. \tag{1.1}$$

We now look what happens to the optimal solution when $d \to \infty$

$$1 - \sqrt[d]{\frac{1}{d+1}} \cdot \Big(\frac{d}{d+1}\Big) = 1 - \Big(\frac{1}{d+1}\Big)^{1/d} \cdot \Big(\frac{d}{d+1}\Big).$$

For $d \to \infty$, $(\frac{1}{d+1})^{1/d} \to 1$ and $(\frac{d}{d+1}) \to 1$, so the optimal solutions approaches zero while the unique Nash equilibrium stays at one. This means that for $d \to \infty$, $PoA \to \infty$, in other words, a very inefficient equilibrium.

1.3 Shapley network design game

A *Shapley network design game* (first studied by Anshelevich et al. (2004) [2]) is defined as follows. The game takes place in an (un)directed graph G, with fixed non-negative edge costs c_e (these costs can represent the cost of installing infrastructure, or leasing a large amount of bandwidth on an existing link).
There are k players and each player i is associated with a source s_i and a destination (sink) t_i. Each player i chooses a path P_i from s_i to t_i to establish connection, so his strategies are the set of available $s_i - t_i$ paths in G. Given the choice of each player, the network is formed as the union of these paths, $\cup_i P_i$. The cost of the network is the sum of the costs of all edges in the network $\sum_{e \in \cup_i P_i} c_e$. The social objective is to minimize these costs.
The key assumption here is that the cost of the network formed is passed onto the players in such a way that the cost of each edge c_e is shared equally among the players who use it.
Formally, each player i incurs cost $\frac{c_e}{f_e}$ for each edge e on P_i where f_e denotes the number of players using edge e. Each individual player chooses a path to minimize the sum of his cost shares.

1.3.1 The H_k example

We consider the network as shown in figure 1.2. Let there be k players, all with the same destination, but with individual sources. $\epsilon > 0$ arbitrarily small. For each $i \in \{1, 2, ..., k\}$, the edge (s_i, t) has cost $\frac{1}{i}$.
The optimal outcome is reached when all players choose the path $s_i \to v \to t$, which gives a total cost of $1 + \epsilon$ for the network. However, this is not a Nash equilibrium, since every player incurs a cost of $\frac{1+\epsilon}{k}$. Hence for player k, choosing the direct path $s_i \to t$ is a dominant strategy (resulting in a cost

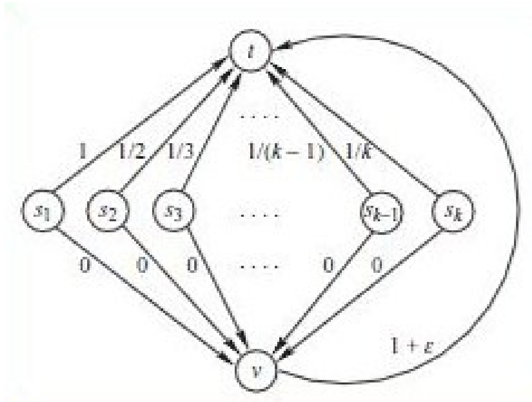

Figure 1.2: H_k example

$\frac{1}{k}$). By induction, going backwards from player $k-1$ to player 1 we find that the unique Nash equilibrium is the outcome where every player i chooses the direct path $s_i \rightarrow t$. The cost of this network is exactly the kth harmonic number $H_k = \sum_{i=1}^{k}(\frac{1}{i})$, which is roughly $\ln k$. We find that, with an optimal outcome of $1+\epsilon$, the Price of Stability can be (arbitrarily close to) H_k in Shapley network design games. Let us now change the network so that all of its edges are undirected, while all costs stay the same. Now in the optimal outcome, each player i chooses the path $s_i \rightarrow v \rightarrow s_k \rightarrow t$ and the cost of the formed network is $\frac{1}{k}$. This is an equilibrium. However, this network possesses multiple Nash equilibria. For $j \in 1, .., k$ fixed, every strategy that has all players i take the route $s_i \rightarrow v \rightarrow s_j \rightarrow t$ is a Nash equilibrium. For those networks the cost is $\frac{1}{j}$ and thus for each individual player the cost is $\frac{1}{j*k}$.
Because the optimal outcome is an equilibrium, the Price of Stability is 1.
The Price of Anarchy however, looks at the worst-case Nash equilibrium. In this case, that is the one where every player i takes the route $s_i \rightarrow v \rightarrow s_1 \rightarrow t$. The cost in this network is 1, so the price of anarchy is $\frac{1}{1/k} = k$.

1.4 Resource Allocation Games

Finally, we'll consider a game that is induced by a natural protocol for allocating resources to players with heterogenous utility functions. We consider a single divisible resource, in this example the bandwidth of a single network link with capacity $C > 0$, that has to be allocated to a finite number $n > 1$ of competing players. For this we assume that each player i has a concave, strictly increasing and continuously differentiable utility function U_i.
The outcome of such a game is a nonnegative allocation vector $(x_1, ..., x_n)$ with $\sum_i x_i = C$. Here x_i denotes the amount of bandwidth allocated to player i. We will study an utilitarian objective, so we're interested in maximizing the sum $\sum_i U_i(x_i)$ of the players' utilities.

We will use a *proportional sharing* protocol here. This means that each user expresses his interest in receiving bandwidth by submitting a nonnegative bid b_i. The protocol then allocates all of the bandwidth in proportion to the bids, so that each user i receives

$$x_i = \frac{b_i}{\sum_{j=1}^{n} b_j} \cdot C$$

units of bandwidth. Player i is then charged his bid b_i. We assume that the payoffs for the players are quasilinear, i.e. the payoff Q_i to a player is defined as his utility for the bandwidth he receives, minus the price he has to pay:

$$Q_i(b_1, ..., b_n) = U_i(x_i) - b_i = U_i\Big(\frac{b_i}{\sum_{j=1}^{n} b_j} \cdot C\Big) - b_i.$$

If all players bid zero, we assume that they will all receive zero payoff.
Because of the posed restrictions on the utility function U_i it's ensured that the payoff function Q_i is continuously differentiable and strictly concave in the bid b_i for every fixed vector b_{-i} with at least one positive component. b_{-i} denotes the vector of bids of players other than i, in formula $b_{-i} = \sum_{j \neq i} b_j$.

Theorem 1.1. *Q_i is continuously differentiable and strictly concave in b_i for a fixed $b_{-i} > 0$.*

Proof. Let b_i be fixed and nonnegative with at least one strictly positive component, so $b_{-i} > 0$.

$$\begin{aligned} Q_i(b_i, b_{-i}) &= U_i(x_i) - b_i = U_i\Big(\frac{b_i}{\sum_{j=1}^{n} b_j} \cdot C\Big) - b_i \\ &= U_i\Big(\frac{b_i}{\sum_{j \neq i} b_j + b_i} \cdot C\Big) - b_i. \end{aligned} \tag{1.2}$$

Then

$$Q_i(b_i, b_{-i}) = U_i\Big(\frac{b_i}{b_{-i} + b_i} \cdot C\Big) - b_i.$$

Because U_i is continuously differentiable, so is $Q_i(b_i, b_{-i})$. Next we need to prove that Q_i is also strictly concave. For that we have to show that

$$\frac{\partial^2 Q_i(b_i, b_{-i})}{\partial b_i} < 0.$$

First

$$\begin{aligned} \frac{\partial Q_i(b_i, b_{-i})}{\partial b_i} &= \frac{\partial U_i}{\partial b_i}\Big(\frac{b_i}{b_{-i} + b_i} \cdot C\Big) - 1 \\ \frac{\partial^2 Q_i(b_i, b_{-i})}{\partial b_i^2} &= \frac{\partial^2 U_i}{\partial b_i^2}\Big(\frac{b_i}{b_{-i} + b_i} \cdot C\Big) \end{aligned}$$

$$\begin{aligned}
\frac{\partial^2 U_i}{\partial b_i^2}\Big(\frac{b_i}{b_{-i}+b_i}\cdot C\Big) &= \frac{\partial}{\partial b_i}\left(\frac{\partial U_i}{\partial b_i}\Big(\frac{b_i}{b_{-i}+b_i}\cdot C\Big)\right) \\
&= \frac{\partial}{\partial b_i}\Big(\frac{C\cdot b_{-i}}{(b_{-i}+b_i)^2}\cdot\frac{\partial U_i(x_i)}{\partial x_i}\Big) \\
&= \frac{-2C\cdot b_{-i}}{(b_{-i}+b_i)^3}\cdot\frac{\partial U_i(x_i)}{\partial x_i}+\frac{\partial^2 U_i(x_i)}{\partial x_i^2}\cdot\Big(\frac{C\cdot b_{-i}}{(b_{-i}+b_i)^2}\Big)^2.
\end{aligned}$$

$U_i(x_i)$ is a concave, strictly increasing continuously differentiable function, so $\frac{\partial U_i(x_i)}{\partial x_i} > 0$ and $\frac{\partial^2 U_i(x_i)}{\partial x_i^2} \leq 0$. $b_{-i} > 0$, so the first part is strictly negative and the second part is non-positive. Thus:

$$\frac{\partial^2 Q_i(b_i, b_{-i})}{\partial b_i} < 0.$$

So Q_i is strictly concave in b_i. □

In this resource allocation game, we define a **Nash equilibrium** as a bid vector in which every user bids optimally, given the bids of the other users.

Definition 1.3. *A bid vector* $(b_i^*, ..., b_n^*)$ *is an* equilibrium *of the resource allocation game* $(U_1, ..., U_n, C)$ *if for every user* $i \in \{1, 2, ..., n\}$,

$$Q_i(b_i^*, b_{-i}^*) = \sup_{b_i \geq 0} Q_i(b_i, b_{-i}^*).$$

Now that we have this definition, we'll prove that such an equilibrium can only be reached if there are at least 2 positive bids.

Theorem 1.2. *Every equilibrium of a resource allocation game has* at least *2 strictly positive components.*

Proof. Firstly, let $(b_1, ..., b_n)$ be an equilibrium without strictly positive components. In other words, there are *no* bids.
This means that all users receive payoff zero. Now let user i change his bid b_i to $0 < \tilde{b}_i < U_i(C)$. Because the restrictions of the utility function U_i this $\tilde{b}_i$ exists. Now the payoff for player i changes to

$$Q_i(\tilde{b}_i, b_{-i}) = U_i(C) - \tilde{b}_i > 0.$$

So the bidvector $(b_i, ..., b_n)$ with all zero bids is not an equilibrium.
Now let $(b_1, ..., b_n)$ be an equilibrium with only one positive bid. W.l.o.g. we assume b_1 to be strictly positive and $b_2, ..., b_n = 0$. Player 1 will get all the bandwidth for price b_i. Now as said before, player 1 can lower his bid and still get all the bandwidth allocated. This increases his and therefore the total payoff. So for $0 < \tilde{b}_i < b_i$ we have that

$$Q_i(\tilde{b}_i, b_{-i}) > Q_i(b_i, b_{-i}).$$

So $(b_1, ..., b_n)$ is *not* an equilibrium.
We conclude that there are no equilibria with just one or zero strictly positive components, so there have to be at least 2. □

With the next example we'll demonstrate that equilibria in resource allocation games can be inefficient.

1.4.1 Example

We consider a resource allocation game in which we have n players and the capacity $C = 1$, the first player has the utility function $U_1(x_1) = 2x_1$ and the other $n-1$ users have utility function $U_i(x_i) = x_i$. It is quite clear that in the optimal allocation, player 1 will receive all of the bandwidth, resulting in an objective function value of 2. This allocation, unfortunately, is not an equilibrium. Because of the *proportional sharing* protocol, the only bid vectors that would result in such an optimal allocation are the ones where only player 1 submits a positive bid. As we've shown in Theorem 2, such a bid vector can never be an equilibrium because the first player can always bid a smaller amount and continue to receive the same bandwidth.
To find the equilibrium, we start from the defining equation of the function Q_i. Let $B = \sum_{i=1}^{n} b_i$, the sum off all bids. Using that $U_1(x_1) = 2x_1$ and $C = 1$, we obtain the condition $\frac{2(B-b_1)}{B^2} = 1$ for player 1.
For players $i > 1$ the same calculations give us the condition $\frac{B-b_i}{B^2} = 1$. Subtracting the second condition from the first gives $2b_1 - b_i = B$ for every $i = 2, 3, ..., n$. Adding these $n-1$ equations together gives $2(n-1)b_1 - (B - b_1) = (n-1)B$. Solving this equation we find that the first player's bid is only a $\frac{n}{2n-1}$ fraction of the sum of all bids:

$$b_1 = nB/(2n-1).$$

This results in player 1 only getting a $\frac{n}{2n-1}$ fraction of the bandwidth. The other $(n-1)$ players equally divide the other $\frac{n-1}{2n-1}$ part of the bandwidth.
The sum of the players utility in this equilibrium is:

$$\begin{aligned}\sum_i U_i(x_i) &= 2 \cdot \frac{n}{2n-1} + 1 \cdot \frac{n-1}{2n-1} \\ &= \frac{3n-1}{2n-1} \\ &= 1 + \frac{n}{2n-1}.\end{aligned}$$

As n grows large, this means that roughly half of the bandwidth will be allocated to the first player, while the rest is split evenly among the other $n-1$ players. For $n \to \infty$, $\sum_i U_i(x_i) \to 3/2$. The optimal solution gave a value of 2, so $POA = \frac{2}{3/2} = 4/3$.
This is the same POA as for Pigou. This is no coincidence. For affine cost functions, the price of anarchy in network games is always at most 4/3 [17].

Chapter 2

An introduction to Peer-to-Peer Networks

After the emergence of the internet, people soon started to use it as a way to share resources, such as files and processing power. The first networks created for this purpose were centralized server-based models. The workings of these networks are pretty straightforward. One big server functions as a sort of warehouse, where all files are stored and people (or *clients*) can access the server and then download the file they want to their computer.
Although these networks work quite well there is one practical problem. These servers have to be very big to service all clients and this costs money.
To split these costs a new model was created: Peer-to-Peer (P2P) networks.
In peer-to-peer networks, users act simultaneously as clients, as well as servers. Each user sets aside a portion of his resources, be it files, network bandwidth, processing power or disk storage, without central coordination.
The peers now act as both consumers and suppliers in this system.

2.1 P2P File-Sharing Game

In a P2P File-Sharing system, a peer has two roles: on the one hand he is a client who wants to download a file and who benefits from a successful download, on the other hand he is a server who is asked to upload (part of) a file, which will, if he agrees, cost him something in the form of bandwidth and CPU usage.
It is clear that if we look at this as a one-shot game, from a game theoretical standpoint, the strategy where we download while being a client and refuse to upload while being a server, would be the dominant strategy. However, with the knowledge that the player may want to download a file from the other in the future, he might be inclined to agree to upload the file in fear of repercussion.
In large p2p networks there is no guarantee that two peers will have multiple transactions in a lifetime. And even if they do so, there is no guarantee that the roles will be such that reciprocity or retaliation is possible. In such a large

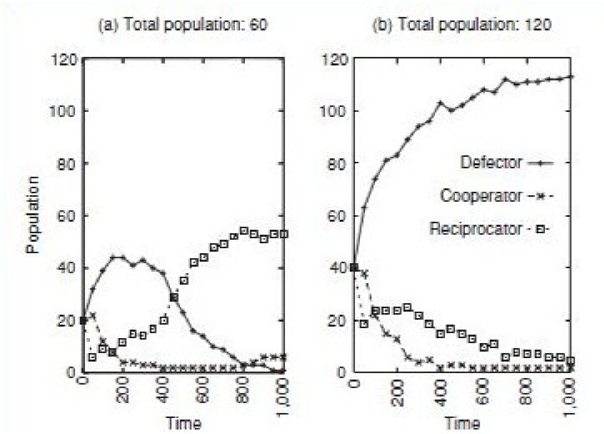

dynamic population with random matching of players, the probability of repeated interaction between players might be too small as to force cooperation and keeps the free-riding strategy dominant.
In the figure above, taken from a study of a p2p file-sharing game (Feldman et al., 2004 [6]), two simulations have been run, with three different kind of strategies, starting with equal share of players. The first strategy is to always defect, the second to always cooperate and the third is a reciprocative strategy. The game consists of 1,000 rounds and the size of the population that plays each strategy is proportional to its success in the previous round.
The simulation on the left shows a game with a population of only 60 where the one on the right has a population double that number, 120. We see that for a relatively small population the reciprocative strategy is dominant after 1,000 rounds, but for the larger population of only twice the size, we see that the interaction between pairs of players is not frequent enough to make this an effective strategy against defectors. These simulations suggest that strategies based on direct reciprocity might not work effectively in these kind of P2P systems with random matching and large populations. To overcome this we have basically two options. We can force repeated interactions with a small number of peers. The other option is to change to a strategy based on indirect reciprocity. For that we need some kind of reputation system, where peers can base their actions on opponents past behaviour.
Lastly, because of the size of these large, dynamic p2p systems with a high turnover rate, it is important to think about how one deals with strangers.
Always cooperating with strangers might encourage strangers to join the system, but it also makes it easily exploitable by so called *whitewashers*, who will simply leave the system when their reputation gets bad and start off fresh with a new identity.
Always defecting against newcomers is robust against such whitewashers, but has the disadvantage of discouraging newcomers [12].

2.2 History of P2P-file sharing

Peer-to-peer networks have evolved enormously in the last decade. To get a better understanding of these models, I'll give a brief summary of the most popular network models and their workings.

2.2.1 Napster

One can not write about peer-to-peer file sharing, without mentioning *Napster*. Although technically not the first, it is generally seen as the pioneer of P2P file sharing systems.
Napster was developed in 1999 by an American teenager, Shawn Fanning, and became a huge success that changed the internet, the music industry (and incidentally also the way we now look at intellectual property) [23]. Napster focused solely on the sharing of MP3 files, i.e. songs.
The network worked as follows: Everybody connected to Napster had a directory on their computer that was shared, so that someone else could access it. The files were catalogued by a central server, so that it had a complete list of songs available at every hard disk connected to Napster at that time. If someone wanted to download a song, he would log in to the system and typed what he was looking for. The central index server (the main program) then searched its index to find all the users who had the song in question.
If a match was found, the user was informed where the requested file could be found. Next the user would connect to one of the peers who had the original file directly and download the file from him.
In this case the central server only functioned as an index, or catalogue, that connected you to a peer who had the file you were looking for. The actual file-transfer went directly peer-to-peer and not through the server.
The fact that a central server was involved, albeit just for indexing and searching, led to Napster eventually being held liable for copyright infringement and Napster shut down in 2001 by court order.
Before that happened though, Napster peaked at over 60 million users per month.

2.2.2 Gnutella

The second generation of peer-to-peer networks found a solution to the problem that eventually led to Napsters downfall. While users still place files on their hard disks that they want to share and make available for downloading like with Napster, the big difference with these networks using the so called *Gnutella* architecture, is that they work without a central database. Instead, all of the machines on the network tell each other about available files using a distributed query approach.
So how does one find the location of all available files without the use of a centralized server?
You send a request for a certain song (or other file) to one or more known

Gnutella users (virtual neighbours), these machines search their hard disks for the requested file. If found, they send back the file name and where to find it to the requester. At the same time, these machines send out the same request to the machines they are connected to, i.e. their neighbours. This process repeats itself. This way, in just a limited number of rounds the request can reach a lot of other machines. Most systems have a *Time to live* mechanism placed on it, that limits the number of levels a request goes. If each machine on the network knows of just four others, in just seven levels over 8.000 machines can be reached (the first level reach 4 new peers, all other levels 3 new peers per current one: $4 \cdot 3^6 = 8749$). In practice this number will be lower, because some neighbours will overlap.
These virtual neighbours are assigned to you when first installing the software and after that update every time you open the software.
All you need is to know of at least one other machine running the Gnutella software and you are able to query practically the whole network. From a legal standpoint it is also very advantageous because there is no single authority that controls everything.
The disadvantages are that the number of machines reached is by no means guaranteed. Because we go several levels deep, it might take some time for your query to find a file. Lastly because you are also a part of the network, your machine also has to answer requests, pass them along and rout back responses. This costs bandwidth.
These disadvantages haven't stopped people using these clients though, as hundreds of millions of copies of Gnutella clients have been downloaded over the years. Some of these clients still exist, although they dropped significantly in popularity in the last few years, due to alternatives, lawsuits and a lack of cooperation between peers [1]. Created in 2000, they reached the top of their popularity in 2006 and were the #1 P2P application as late as 2007.

2.2.3 KaZaA

Another second generation P2P architecture is *KaZaA*. KaZaA, like the Gnutella network, uses a decentralized system. Users contact each other directly online to share content. The biggest difference between the KaZaA architecture and the one of Gnutella is the way in which you contact peers. KaZaA uses a so called *FastTrack* protocol. The system divides KaZaA users in two groups: *supernodes* and *ordinary nodes.*
Supernodes are powerful computers with fast network connections, high bandwidth and quick processing capabilities. These supernodes act like traffic hubs, processing data requests from the slower ordinary nodes. Each supernode may serve between 60 and 150 ordinary nodes.
The KaZaA software comes with a coded list of supernodes. Every time a user launches the KaZaA application his computer chooses from a list of currently active supernodes. Now when the user sends out a request for a file, the request is funneled through that supernode. That supernode communicates (or passes along the request) with other supernodes, which in turn connect to regular nodes

that in turn connect to even more regular nodes, to fulfill the request until the *Time to Live* limit of 7 is reached. Once the correct file is located the user can download the file directly from the owner; it doesn't have to go through a supernode.
KaZaA, like Gnutella served millions of users in their hay days, but eventually succumbed to legal trouble in 2004/2005.

2.2.4 BitTorrent

The *BitTorrent* P2P-method is often called the third generation of file sharing [22].
Unlike the peer-to-peer downloading methods described above, BitTorrent does use some form of a central entity (called a *tracker*) to help with the file searching. It also works on the basic principal of direct reciprocity, i.e. tit-for-tat, thus encouraging cooperation.
The other big difference is that BitTorrent partitions files into smaller pieces, so that one can download different pieces simultaneously from multiple computers. BitTorrent software communicates with a tracker to find other computers running BitTorrent that have the complete file (seeds) and those with portions of the file (peers that are also in the process of downloading the file). The tracker identifies these connected computers and helps the client to trade pieces of the file with other users.
Downloading multiple pieces at the same time solves a common problem with other peer-to-peer downloading methods: peers upload at a much slower rate than they download. By downloading multiple pieces from different sources at the same time, the overall speed can greatly improve. This method is therefore especially useful for large, popular files (like movies, TV series, software and games).
One of the most famous tracker sites is The Pirate Bay, who is in a long legal battle with the authorities. In The Netherlands this has resulted in the decision that several internet providers had to block their clients access to this site.

2.3 Public opinion and usage

Over the years peer-to-peer online file sharing has become more and more popular, recent studies showed that in 2008, over 20 percent of Europeans had used file sharing networks to obtain music at least once. In comparison, only 10 percent ever used paid-for digital services such as iTunes [20].
Public perception towards this is also quite interesting. While an overwhelming majority thinks that stealing a DVD from a store constitutes a serious offense, only half of those people think that way about illegally downloading that same movie from the internet. Among people under 30 this percentage is even lower [21].
Selfishness and convenience seem to get the upper hand over moral objections and 'fairness'.

2.4 Selfish Behaviour and Incentive Mechanisms

It is that sort of selfish behaviour that we also find within the peer-to-peer systems. In all the networks mentioned above, within months of introduction, the majority of all users were *free-riders*, downloading files, but refusing to upload any in return [1]. Because of the great size, high turnover and relative anonymity of the p2p networks, most of its transactions are interactions between strangers who will never meet again in the future. Not surprisingly this makes it very hard to force any kind of cooperation. Making matters worse is the fact that the supply of online identities is limitless, so you can make as many new identities as you want at no cost.
Generally, we can divide the problems caused by selfish behaviour into three categories.
The first problem is, as described above, peers who only take (download) from the system, without giving something back (refusing to upload). Because other peers will still want to download the same files, by refusing to upload, the burden on the system falls on a smaller number of peers.
The second problem is, *resource availability.* Besides uploading and downloading existing files in the system, this is the most important thing in a P2P file sharing system. Where the first problem focuses on the sharing, this one focuses on the *files.* Without a healthy stream of new files regularly entering the system, the system can't survive.
Finally, the third problem concerns the *hidden actions* in the system. Most importantly while sending a search request. As described above, when searching for a file, we send a request to our virtual neighbours, who we then expect to forward that request to their neighbours and so on. The problem here, however, is that we have no way of knowing wether or not they actually forwarded this request. The only information we receive, is wether or not they, as a group, were able to locate the file. Because forwarding requests costs bandwidth and CPU, selfish players would want to do this as little as possible. If too many peers do this, it becomes increasingly difficult to successfully locate files, which can seriously affect the system performance.
In the next three chapter I will discuss all these problems more extensively and examine how incentive schemes can help to resolve them.
To combat these problems the incentive based mechanisms that P2P network developers try to implement in the networks can be divided into three different categories: *currency*, *reputation* and *bartering.*

2.4.1 Currency

A currency scheme is pretty straight forward. Peers can earn currency by contributing resources to the system and in turn spend that currency by obtaining resources from the system. In the past there have been such p2p-systems, most notably MojoNation, but over the years these have lost in popularity in favor of other systems.
It is shown (Friedman et al., 2006 [8]) that for each fixed amount of money

supply in the system, there is a nontrivial Nash equilibrium where all peers play a threshold strategy. This means that a peer will satisfy a request, earning some money, if his balance is under a certain threshold value and refuses the request if his currency is above that value. By comparing the efficiency of equilibria it is possible to determine the money supply level that maximizes efficiency for a system of a given size [8].
It turns out this can be controlled by either injecting or removing currency, or by changing the price of servicing a request. That second option basically says that inflation can be used as a tool to maintain efficiency as the system grows in size.
Also for these kind of systems, robustness against Sybil and whitewashing attacks are still an important requirement [8]. Whitewashing can occur if either new users start with a positive balance or if accounts can have negative balances.

2.4.2 Reputation

Simply put, reputation schemes reward good behavior and punish bad. These kind of schemes have had great success in various areas, including online marketplaces such as Ebay or Marktplaats [12]. In P2P schemes the reward and punishment is the quality of service peers receive from other peers. These service differentiation schemes reward a high reputation with good service from others (for file-sharing this means high priority downloads) while the opposite holds for bad reputations.
There are however a few ways to cheat these kind of systems. By colluding with other peers, it is possible to artificially boost each others ratings (or unjustly smear others). Some other problems are caused by the easy access to pseudonyms. First there are so called *Sybil attacks*, where a single user creates multiple identities that collude with each other (so basically the same as above, but now you don't need coconspirators). Also, just like in currency based schemes, creating a new identity whenever your reputation has fallen too low, let you defect on every transaction without facing the consequences for it (whitewashing).

2.4.3 Bartering: BitTorrent

BitTorrent (BT) is a file-sharing mechanism that has gained greatly in popularity in the last years. Although incentives play a integral part in its design, it differs from other incentive mechanisms because it's based on direct reciprocity rather than indirect reciprocity [22].
In BT, a large file is divided into several small fixed size pieces. A 'seeder' provides different pieces to different peers, who then in turn can exchange those pieces with each other. When a peer has downloaded all the different pieces he can reconstruct the file.
Because the problems that arise with random matchings in large populations, BT forces repeated transactions amongst peers. When a peer starts the download he is matched with a small set of peers who are downloading the same file.

He then chooses a few peers to interact with, periodically updating as to choose the peers with the best rates. To induce the peers to upload their parts, the download rates of the peers are influenced by their upload rates through direct reciprocity.
For larger files such as movies or software packages, the repeated interaction between two peers can be quite large, which forces them to contribute, for fear of retribution.
A downside of this barter scheme is that is does not address cooperation beyond the file download period. Therefore there is no incentive for peers to upload original files, called *seeding*.
To address this issue a number of communities have implemented some form of reputation scheme on top of the barter scheme already in place, excluding peers with low contribution levels.
In practice it turns out that these schemes result in some highly effective cooperative communities. That being said, theoretically the BT protocol is still vulnerable to manipulations by a selfish peer [12]. Sybil attacks (here a user 'splits' himself, thereby increasing his chance of being selected to upload to), whitewashing and even uploading fake data to boost ones upload rate are still threats to this system.

Chapter 3

Resource Availability

As discussed in the previous section, there are a number of potential problems in Peer-To-Peer Systems, allowing peers to 'not pay their fair share' and threatening the system as a whole. Those problems lower the overall performance and even threaten to destroy it. In this chapter, I will examine a model, first proposed by Buragohain et al. [5], that uses a differential service-based incentive scheme to improve the availability of resources in a P2P model.
In this model we'll focus on the availability of (new) resources. No matter how eager peers are to upload files to each other, if nobody is adding new files to the system, eventually the system will succumb because of this lack of new material.
In this model, we'll assume rational, strategic peers, who want to maximize their utility by participating in the P2P system. The utility depends on the benefit (the resources of the system they can use) and their cost (their contribution to the system). The benefit will be a monotonically increasing function of a peer's contribution.
Because everybody wants to maximize his utility (being the rational players that they are) we have a non-cooperative game amongst peers.
First we'll look at the model with homogeneous peers, where the assumption is that all peers equally benefit from everybody else.
Next, the more general case with heterogeneous peers will be considered. This allows for different benefit functions for each pair of peers, because not everyone values the same resources equally.
The behavior that a player adopts while interacting with other players is known as that player's *strategy*. Here the player's strategy is his level of contribution. The player derives a benefit from his interaction with other players which is termed as a utility. To characterize the dynamics of the system, we will focus on *Nash equilibria*.
In order to render a verdict over the effectiveness of the system and its Nash equilibria, I will introduce a maximum capacity. This will be the maximum a player can contribute to the system. Without this maximum, we'll find that the optimal outcome of the system is usually unbounded when contribution levels grow large. Therefore, adding this maximum will allow us to find meaningful,

non trivial answers while qualifying equilibria.
Fortunately this maximum is not only practical, but also realistic as no one actually has unlimited resources.

3.1 The Model

There are N players/peers, $P_1, P_2, ..., P_N$, in the system. U_i will denote the utility function of the ith peer.
To denote the contribution of P_i, we'll use D_i. This can be disk space one contributes to the system, the number of downloads served or the amount of files shared, over a fixed period of time (hour / day / week).
Each contributed unit of resource incurs a cost c_i, so the total cost for participating in the system is $c_i D_i$. For convenience a normalized contribution of

$$d_i \equiv D_i/D_0$$

will be used, where D_0 is an absolute measure of contribution that the system architect is free to set. The incentive scheme will strive to ensure all peers contribute at least D_0.
Because nobody's resources are limitless, I'll also introduce a maximum capacity, M_i, which is the maximum that player P_i can contribute to the system. So $D_i \leq M_i$. We normalize this such that

$$m_i \equiv M_i/D_0.$$

This results in $d_i \leq m_i$. To measure the benefit of a peer's contribution to the system, a $N \times N$ matrix B is used, where B_{ij} denotes how much the contribution of P_j is worth to player P_i. So if player i has no interest in the contribution of player j, $B_{ij} = 0$. In general, $B_{ij} \geq 0$ and $B_{ii} = 0$, for all i.
Again a set of normalized parameters is defined, corresponding to B_{ij}.

$$b_{ij} = B_{ij}/c_i, b_i = \sum_j b_{ij}, b_{av} = \frac{1}{N} \sum_i b_i,$$

where b_i can be interpreted as the total benefit that P_i can derive from the system if all other users make unit contributions each. This parameter will turn out to be important in determining whether it is worthwhile for a player to join the system. I'll show there exists a critical value b_c of benefit, such that in the homogeneous system for $b_i < b_c$ it's better for the players to not join the system.
Differential service can be handled in different ways, in this model we use reputation. The scheme is implemented as followed: a peer P_j accepts a request for a file from peer P_i with probability $p(d_i)$ and rejects it with probability $1-p(d_i)$, the players contribution d_i is thus directly linked to his reputation in the system and is sent along with his request.
The choice of the exact probability function seems to have no effect on the

qualitative nature of the results, so any reasonable probability function that is monotonically increasing will do [5]. Here we'll use

$$p(d) = \frac{d^\alpha}{1 + d^\alpha}, \alpha > 0.$$

Now $p(0) = 0$ and $p(d) \to 1$ as $d \to \infty$. The choice of α determines the steepness of the probability function.
This gives us the following utility function, defined as U_i, derived for P_i if he decided to join the system:

$$U_i = -c_i D_i + p(d_i) \sum_{j \neq i} B_{ij} D_j.$$

With the use of the normalized parameter

$$u_i = \frac{U_i}{c_i D_o},$$

the utility can be rewritten as

$$u_i = -d_i + \sum_{j \neq i} b_{ij} d_j p(d_i).$$

The $-d_i$ term is the scaled cost to join the system for P_i and increases linearly if more is contributed to the system.

3.2 Homogeneous System of Peers

First we take a look at the homogeneous system, where $b_{ij} = b$ for all $i \neq j$. So all peers derive equal benefit from each other.
Here the utility function can be reduced to

$$u = -d + (N - 1)bdp(d)$$

and $b_i = b_{av} = b(N - 1)$ for all P_i.
Because of symmetry, the problem can be reduced to a 2-person game [5].

3.2.1 Two person game

We can rewrite the utility functions again to

$$\begin{aligned} u_1 &= -d_1 + bd_2 p(d_1) \\ u_2 &= -d_2 + bd_1 p(d_2). \end{aligned}$$

Assuming $\alpha = 1$, we get $p(d) = d/(1 + d)$.
Now we want to know if there exists a Nash equilibrium for large enough values of benefits where both peers can get a non-zero utility from their interaction.

We start by finding the Nash equilibria for the system *without* the capacity constraint.
Obviously if one of the players doesn't contribute, the best response from the other would be to do the same, i.e. leave the system. But if we suppose that P_2 decides to make a contribution $d_2 > 0$ to the system, P_1 will respond to that, by resetting his contribution level d_1 so as to maximize u_1:

$$\begin{aligned} u_1 &= -d_1 + bd_2 \frac{d_1}{1+d_1} \\ \frac{du_1}{dd_1} &= -1 + bd_2 \frac{1}{(1+d_1)^2} \\ 0 &= -1 + bd_2 \frac{1}{(1+d_1)^2}. \end{aligned}$$

This gives

$$r_1(d_2) \equiv d_1 = \sqrt{bd_2} - 1.$$

Where $r_1(d_2)$ is the *reaction function* for P_1. Because P_2 knows P_1 will respond, his reaction function will symmetrically be

$$r_2(d_1) \equiv d_2 = \sqrt{bd_1} - 1.$$

We find a Nash Equilibrium if there exists a pair (d_1^*, d_2^*), s.t. they form a fixed point for the reaction functions above:

$$d_1^* = \sqrt{bd_2} - 1$$

$$d_2^* = \sqrt{bd_1} - 1.$$

This gives $d_1^* = d_2^* = d^*$, so we have

$$\begin{aligned} d^* &= \sqrt{bd^*} - 1 \\ (d^* + 1)^2 &= bd^* \\ {d^*}^2 + (2-b)d^* + 1 &= 0. \end{aligned}$$

This gives

$$d^* = \left(\frac{b}{2} - 1\right) \pm \left((\frac{b}{2} - 1)^2 - 1\right)^{1/2}. \tag{3.1}$$

This equation only has a solution for $b \geq 4 \equiv b_c$. So we derive that the critical value for a peer to join the system is $b_c = 4$. This critical value obviously depends on the choice of $p(d)$, but is independent of the number of peers in the system.
This gives us the following result:

Theorem 3.1. *The number of Nash equilibria in the homogeneous two person game depends on b.*
$d^* = 0$ is always a Nash equilibrium. The additional non-zero Nash equilibria are the solutions of (3.1)*, such that:*

- *For $b > 4$ there are 3 Nash equilibria.*
- *For $b = 4$ there are 2 Nash equilibria.*
- *For $b < 4$ there is 1 Nash equilibrium.*

Proof. First we see that for every b, $d_1 = d_2 = 0$ is always a Nash equilibrium. If neither player contributes, neither of them can improve his utility by unilaterally contributing some amount larger than 0.
We find the other Nash equilibria where d_1 is the optimal response to d_2 and vice versa. We found that this is the case for

$$d_1 = d_2 = d^* = \left(\frac{b}{2} - 1\right) \pm \left((\frac{b}{2} - 1)^2 - 1\right)^{1/2}.$$

For $b > 4$, this equation has 2 solutions, so the system has a total of 3 Nash equilibria.
For $b = 4$, the equation has 1 solution, so the system has a total of 2 Nash equilibria.
For $b < 4$, the equation has no solutions, so the only Nash equilibrium is $d_1 = d_2 = 0$. □

For $b = b_c$ the only non-zero solution is $d^* = 1$. For $b > b_c$, there are two non-zero solutions. We define d^*_{hi} to be the bigger of the two and d^*_{lo} to be the smaller.
For general values of α we have

$$\begin{aligned}
u_1 &= -d_1 + bd_2\frac{d_1^\alpha}{1 + d_1^\alpha} \\
\frac{du_1}{dd_1} &= -1 + bd_2\frac{\alpha d_1^{\alpha-1}}{(1 + d_1^\alpha)^2} \\
0 &= -1 + bd_2\frac{\alpha d_1^{\alpha-1}}{(1 + d_1^\alpha)^2} \\
1 &= bd_2\frac{\alpha d_1^{\alpha-1}}{(1 + d_1^\alpha)^2} \\
(1 + d_1^\alpha)^2 &= bd_2\alpha d_1^{\alpha-1}.
\end{aligned}$$

Using $d_1 = d_2 = d$, we get

$$\begin{aligned}
0 &= (d^\alpha)^2 + (2 - b\alpha)d^\alpha + 1 \\
d^\alpha &= \frac{-(2 - \alpha b) \pm \sqrt{(2 - \alpha b)^2 - 4}}{2} \\
d^\alpha &= \left(\frac{\alpha b}{2} - 1\right) \pm \sqrt{\left(\frac{\alpha b}{2} - 1\right)^2 - 1} \\
d^* &= ((b\alpha/2 - 1) \pm ((b\alpha/2 - 1)^2 - 1)^{1/2})^{1/\alpha}. \qquad (3.2)
\end{aligned}$$

Here a solution only exists if $b\alpha \geq 4$. This gives us the following result:

Theorem 3.2. *The number of Nash equilibria in the homogeneous two person game depends on b and α.*
$d^* = 0$ is always a Nash equilibrium. The additional non-zero Nash equilibria are the solutions of (3.2), such that:

- *For $b\alpha > 4$ there are 3 Nash equilibria.*
- *For $b\alpha = 4$ there are 2 Nash equilibria.*
- *For $b\alpha < 4$ there is 1 Nash equilibrium.*

Proof. Substituting $b\alpha$ for b, the proof of this theorem is completely analogous to that of Theorem 3.1. □

Now I'll re-introduce the maximum capacity m_i. Because the system is homogeneous, I'll assume $m_1 = m_2 = m$. I'll characterize the effect on the Nash equilibria in the following theorem.

Theorem 3.3. *The capacity constraint $d \leq m$ has the following effect on the Nash equilibria.*

- *Let $b\alpha > 4$. If $m \geq d^*_{hi}$ the Nash equilibria are unchanged. If $d^*_{lo} \leq m < d^*_{hi}$, the largest Nash equilibrium moves to $d = m$. For $m < d^*_{lo}$, the only remaining Nash equilibrium is $d = 0$.*
- *Let $b\alpha = 4$. If $m \geq d^* = 1$ the Nash equilibria remain unchanged. For $m < d^* = 1$, only the Nash equilibrium $d = 0$ remains.*
- *Let $b\alpha < 4$. The maximum capacity has no effect on the only equilibrium $d = 0$.*

Proof. Because the utilities of player 1 and 2 are symmetric, it is enough to only look at that of player 1.
The utility of player 1 is

$$u_1 = -d_1 + bd_2 \frac{d_1^\alpha}{1 + d_1^\alpha}.$$

The derivative is

$$\frac{du_1}{dd_1} = -1 + bd_2 \frac{\alpha d_1^{\alpha-1}}{(1 + d_1^\alpha)^2}.$$

Using symmetry, $d_1 = d_2 = d$, we get

$$\frac{du_1}{dd_1} = -1 + b\alpha \frac{d^\alpha}{(1 + d^\alpha)^2}.$$

- For $b\alpha > 4$ this means:
  $\frac{du_1}{dd_1} < 0$ for $d < d^*_{lo}$, $\frac{du_1}{dd_1} > 0$ for $d^*_{lo} < d < d^*_{hi}$ and $\frac{du_1}{dd_1} < 0$ for $d > d^*_{hi}$.
  If $m \geq d^*_{hi}$, obviously the 3 original Nash equilibria still exist and because $\frac{du_1}{dd_1} < 0$ for $d > d^*_{hi}$ the capacity constraint doesn't give an extra equilibrium.

For $d^*_{lo} \leq m < d^*_{hi}$, the players can't reach the biggest equilibria d^*_{hi}. However, $\frac{du_1}{dd_1} > 0$ for $d^*_{lo} < d_1 = d_2 = d < d^*_{hi}$. Hence for $d = m$, the players can only improve their utilities by increasing their contribution. Since this is impossible, $d = m$ is a Nash equilibrium.
For $m < d^*_{lo}$. Only the Nash equilibrium $d = 0$ can be reached and because $\frac{du_1}{dd_1} < 0$ for $d < d^*_{lo}$, this is the only equilibrium.

- For $b\alpha = 4$.
  $\frac{du_1}{dd_1} < 0$ for $d \neq 1$.
  Therefore, if $m \geq d^* = 1$, the Nash equilibrium $d^* = 1$ can be reached and $d = m$ isn't an equilibrium. If $m < d^* = 1$, the non-zero Nash equilibrium can't be reached and $d = m$ is not an equilibrium, so the only remaining Nash equilibrium is $d = 0$.

- For $b\alpha < 4$. The only Nash equilibrium is $d = 0$ and the maximum capacity has no effect on that.

□

Price of Stability

We define the systems' utility as the sum of all players utilities. For the 2-person homogeneous system and $\alpha = 1$, this gives

$$u_{sys} = -(d_1 + d_2) + bd_2 \frac{d_1}{1 + d_1} + bd_1 \frac{d_2}{1 + d_2}.$$

Due to homogeneity, we can rewrite this as

$$u_{sys} = -2d + 2b \frac{d^2}{1 + d}.$$

Because the system is homogeneous we'll let $m_1 = m_2 = m$.
We find the optimal system performance by first calculating $u_{sys} \geq 0$. For $b \geq b_c = 4$, this gives

$$\begin{aligned} -2d + \frac{2bd^2}{1 + d} &\geq 0 \\ \frac{bd^2}{1 + d} &\geq d \\ bd^2 &\geq d(1 + d) \\ d((b - 1)d - 1) &\geq 0 \\ d \geq \frac{1}{b - 1}. & \end{aligned} \tag{3.3}$$

Next we find $\frac{du_{sys}}{dd}$

$$\begin{aligned}\frac{du_{sys}}{dd} &= -2+2b\frac{2d+d^2}{(1+d)^2}\\ &= -2+2b-2b\frac{2b}{(1+d)^2}.\end{aligned}$$

Calculating $\frac{du_{sys}}{dd} \geq 0$ gives

$$\begin{aligned}-2+2b-2b\frac{2b}{(1+d)^2} &\geq 0\\ 2b-2 \geq 2b\frac{2b}{(1+d)^2}\\ (1+d)^2 \geq \frac{b}{b-1}\\ 1+d \geq \sqrt{\frac{b}{b-1}}\\ d \geq \sqrt{\frac{b}{b-1}}-1. \qquad (3.4)\end{aligned}$$

Using (3.3) and (3.4) and $b \geq 4$, we find that

$$\begin{aligned}\sqrt{\frac{b}{b-1}}-1 &< \frac{1}{b-1}\\ \sqrt{\frac{b}{b-1}} &< \frac{b}{b-1}.\end{aligned}$$

So u_{sys} is continuously increasing for $u_{sys} > 0$. This means that u_{sys} is optimal for $d_1 = d_2 = m$ and that $u_{sys}(d^*_{hi}) > u_{sys}(d^*_{lo}) \geq 0$. Hence the equilibrium with the best system performance is d^*_{hi}.
For $d_1 = d_2 = m$ the utilities of both players are:

$$u_1 = u_2 = -m + b\frac{m^2}{1+m}.$$

The system has an equilibrium for $d_1 = d_2 = 0$, so the Price of Anarchy only gives us a trivial bound ($POA = \infty$). Therefore, using the Price of Stability seems a logical choice:

$$POS = \frac{-m + b\frac{m^2}{1+m}}{-d^*_{hi} + b\frac{(d^*_{hi})^2}{1+d^*_{hi}}}.$$

Depending on the information available to the system architect, I'll show that the POS can either be 1 or is unbounded as a function of m).
The POS will be 1 if $m = d^*_{hi}$. If M is known to the system architect, he can adjust D_0 so that this is the case.

d^*_{hi} does not depend on D_0, while $m = M/D_0$. By setting $D_0 = \frac{M}{d^*_{hi}}$, we'll get $m = M/D_0 = d^*_{hi}$, which results in a POS of 1.
If the system architect does not have this information, the POS can be arbitrarily large.

Example

Let benefits be $b_{12} = b_{21} = b = 6$ and the maximum capacity $m = 5$, then for $\alpha = 1$ we have

$$d^* = \Big(6/2 - 1\Big) \pm \Big((6/2-1)^2 - 1\Big)^{1/2} = 2 \pm \sqrt{3}.$$

The best attained equilibrium is for $d^*_{hi} = 2 + \sqrt{3}$, resulting in utilities

$$u = -(2+\sqrt{3}) + 6 \cdot \frac{(2+\sqrt{3})^2}{1+2+\sqrt{3}} = 7 + 4\sqrt{3}.$$

$m > d^*_{hi}$, so the utility is optimal for $d = m = 5$:

$$u(5) = -5 + 6 \cdot \frac{5^2}{1+5} = 20.$$

So without interference from the system architect the price of stability is:

$$POS = 20/(7+4\sqrt{3}) \approx 1.436.$$

If the system architect has full information, his best action will be to set D_0 s.t. $m = d^*_{hi}$. Then $POS = 1$.

3.2.2 N person game

Next, we discuss the N person homogeneous game. The fixed point equations (for $\alpha = 1$) are now

$$d^* = \sqrt{b(N-1)d^*} - 1.$$

So

$$d^* = (b(N-1)/2 - 1) \pm ((b(N-1)/2 - 1)^2 - 1)^{1/2}. \tag{3.5}$$

By replacing b with $b(N-1)$ the results of the two player game are transferrable to the N player game. Although we can't transfer this to the heterogeneous system, we'll see that the average properties of the Nash equilibria closely follow the homogeneous case.
Disregarding the capacity constraint we find the following result for the N player game:

Theorem 3.4. *The number of Nash equilibria in the homogenous N person game depends on b and the number of players, N. $d^* = 0$ is always a Nash equilibrium. The additional non-zero Nash equilibria are the solutions of* (3.5), *such that:*

- *For $b(N-1) > 4$, there are 3 Nash equilibria.*
- *For $b(N-1) = 4$, there are 2 Nash equilibria.*
- *For $b(N-1) < 4$, there is 1 Nash equilibrium.*

Proof. By substituting $b(N-1)$ for b, the proof of this theorem is completely analogous to that of Theorem 3.1. □

Just as in the 2 person game, adding the maximum capacity constraint effects the Nash equilibria.

Theorem 3.5. *The capacity constraint $d \leq m$ has the following effect on the Nash equilibria:*

- *Let $b(N-1) > 4$. If $m \geq d^*_{hi}$, the Nash equilibria are unchanged. If $d^*_{lo} \geq m < d^*_{hi}$, the largest Nash equilibrium moves to $d = m$. For $m < d^*_{lo}$, the only remaining Nash equilibrium is $d = 0$.*
- *Let $b(N-1) = 4$. If $m \geq d^* = 1$, the Nash equilibria remain unchanged. For $m < d^* = 1$, only the Nash equilibrium $d = 0$ remains.*
- *Let $b(N-1) < 4$. The maximum capacity has no effect on the only equilibrium $d = 0$*

Proof. By substituting $b(N-1)$ for $b\alpha$, the proof of this theorem is completely analogous to that of Theorem 3.3. □

Example (continued)

For $b = 2$ and $N = 5$ and $m = 10$, we get

$$d^* = (2(5-1)/2-1) \pm ((2(5-1)/2-1)^2-1)^{1/2} = 3 \pm 2\sqrt{2}.$$

The best attained equilibrium is for $d^*_{hi} = 3 + 2\sqrt{2}$, resulting in individual utilities

$$u_i = -(3+2\sqrt{2}) + 2(5-1) \cdot \frac{(3+2\sqrt{2})^2}{1+3+2\sqrt{2}} = 17 + 12\sqrt{2}.$$

$m > d^*_{hi}$, so the utilities are optimal for $d = m = 10$.

$$u_i(10) = -10 + 2(5-1) \cdot \frac{10^2}{1+10} = 62\frac{8}{11}$$

So without interference from the system architect the price of stability is:

$$POS = 62\frac{8}{11}/(17+12\sqrt{2}) \approx 1.847.$$

If the system architect has full information, his best action will be to set D_0 s.t. $m = d^*_{hi}$. Then $POS = 1$.

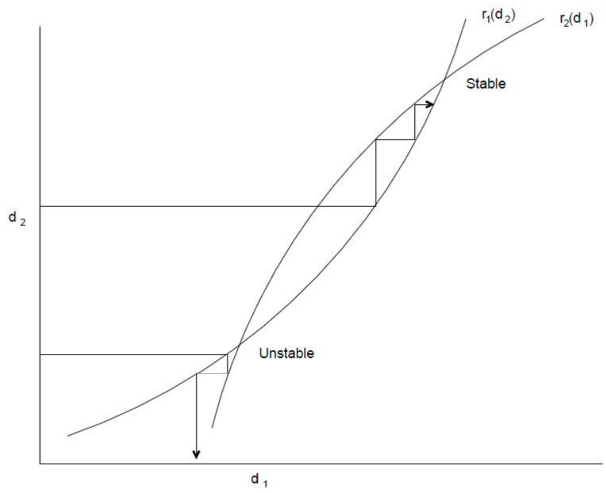

Figure 3.1: The reaction functions of the two-person game

3.3 Stability of equilibria

Because there are two possible equilibria, we must ask ourselves which equilibria will be reached by the system in practice. By plotting the two reaction functions against each other, we can easily deduce that the equilibrium $d^*_{hi} > 1$ is locally stable and the equilibrium $d^*_{lo} < 1$ is locally unstable (as seen in figure 3.1, from [5]).

I'll show this by calculating the Jacobian matrix for the reaction functions $(r_1(d_2), r_2(d_1))$.

$$J(d_1, d_2) = \begin{bmatrix} \frac{\partial r_1(d_2)}{\partial d_1} & \frac{\partial r_1(d_2)}{\partial d_2} \\ \frac{\partial r_2(d_1)}{\partial d_1} & \frac{\partial r_2(d_1)}{\partial d_2} \end{bmatrix} = \begin{bmatrix} 0 & \frac{1}{2}\sqrt{\frac{b}{d_2}} \\ \frac{1}{2}\sqrt{\frac{b}{d_1}} & 0 \end{bmatrix}$$

For $d_1 = d_2 = d$, we find the eigenvalues of this matrix by solving $\lambda^2 - \frac{1}{4}\frac{b}{d} = 0$. This gives us two eigenvalues $\lambda_{1,2} = \pm\frac{1}{2}\sqrt{\frac{b}{d}}$. If the absolute values of these eigenvalues are both strictly less than 1 for an equilibrium, the equilibrium is stable, if one of the absolute values is greater than 1 the equilibrium is unstable. For d^*_{hi}, $|\lambda_{1,2}| = \frac{1}{2}\sqrt{\frac{b}{d^*_{hi}}}$. For $b > 4$, $d^*_{hi} > b/2 - 1 > b/4$, so $|\lambda_{1,2}| < 1$ and d^*_{hi} is a stable equilibrium.

The eigenvalues of $J(d^*_{lo}, d^*_{lo})$ are $\lambda_{1,2} = \pm\frac{1}{2}\sqrt{\frac{b}{d^*_{lo}}}$. $d^*_{lo} < 1$ and $b > 4$, so $|\lambda_{1,2}| > \frac{1}{2}\sqrt{b} > 1$ and d^*_{lo} is an unstable equilibrium.

So for any starting value $d > d^*_{lo}$ the system converges to the stable equilibrium point d^*_{hi} and for $d < d^*_{lo}$ the system ends up with everybody leaving.

3.4 Heterogeneous System

In this system, every peer has his own benefit function.
Just as in the homogeneous system, I'll define the systems' utility as $u_{sys} = \sum_i u_i$, the sum of all the players utilities.
In this section, we set $\alpha = 1$.

3.4.1 Two Person Game

We start with the two person heterogeneous game, without capacity constraints. In this situation, in addition to the trivial equilibrium where none of the players contribute, the Nash equilibria are the fixed points (d_1^*, d_2^*) of the following equations:

$$\begin{aligned} d_1 &= r(d_2) = \sqrt{b_{12}d_2} - 1 \\ d_2 &= r(d_1) = \sqrt{b_{21}d_1} - 1. \end{aligned}$$

Substituting d_2 in the first equation, we get

$$\begin{aligned} d_1 &= \Big(b_{12}((b_{21}d_1)^{1/2} - 1)\Big)^{1/2} - 1 \\ (d_1+1)^2 &= b_{12}(\sqrt{b_{21}d_1} - 1) \\ (d_1+1)^2 + b_{12} &= b_{12}\sqrt{b_{21}d_1} \\ \Big((d_1+1)^2 + b_{12}\Big)^2 &= b_{12}^2 b_{21} d_1. \end{aligned}$$

Resulting in

$$d_1^4 + 4d_1^3 + 2(3+b_{12})d_1^2 + (4 + 4b_{12} - b_{12}^2 b_{21})d_1 + (1+b_{12})^2 = 0.$$

The roots of this equation, provided they exist, give d_1^* and consequently d_2^*.
In Appendix A, I examine the conditions for which this 4th order polynomial equation has solutions. By examining the first and second derivative and then calculating the root for the first derivative, I've found the $d_1 = d_{1,min}$ for which $f(d_1) = d_1^4 + 4d_1^3 + 2(3+b_12)d_1^2 + (4+4b_{12} - b_{12}^2 b_{21})d_1 + (1+b_{12})^2$ is minimal:

$$\begin{aligned} d_{1,min} &= -1 + \frac{1}{6}\Big(27b_{12}^2 b_{21} + 3\sqrt{192b_{12}^3 + 81b_{12}^4 b_{21}^2}\Big)^{\frac{1}{3}} \\ &\quad - \frac{2b_{12}}{\Big(27b_{12}^2 b_{21} + 3\sqrt{192b_{12}^3 + 81b_{12}^4 b_{21}^2}\Big)^{\frac{1}{3}}}. \end{aligned} \tag{3.6}$$

It follows that if $f(d_{1,min}) \leq 0$, the equation has solutions.
Substituting $d_1 = d_{1,min}$ in the equation doesn't get us much closer to the constraints, but doing so, does give us an equation with only 2 unknown values, b_{12} and b_{21}. This means we are able to plot the equation. On the next page we see the result, where the vertical axis is $f(d_{1,min})$. As expected we see that

for $b_{12}, b_{21} \geq 4$ we always have solutions for the equation. Furthermore we see that for $b_{12} \cdot b_{21} < 16$ the equation has no solutions. So we need *at least* $b_{12} \cdot b_{21} \geq 16$. However, for $b_{12} \neq b_{21}$ this restriction is not enough. Further examination shows that the bigger the difference between b_{12} and b_{21}, the bigger the product $b_{12} \cdot b_{21}$ has to be. For example $b_{12} = 5$ and $b_{21} = 3.3$ give two roots while $5 \cdot 3.3 = 16.5$, meanwhile if $b_{12} = 16$, b_{21} has to be bigger than 1.2 in order for $f(d_1)$ to have roots, although $16 \cdot 1.2 > 19$.

Now, because we don't have explicit values of the equilibria, it is hard prove something conclusively about the effect of introducing the capacity constraints m_1, m_2. However, I do claim the following.

Claim 3.1. *Assume the heterogeneous two person game without capacity constraints has two non-zero equilibria. Let $(d^*_{1,hi}, d^*_{2,hi})$ and $(d^*_{1,lo}, d^*_{2,lo})$ respectively be the larger and the smaller non-zero equilibrium. Then introducing capacity constraints, m_1, m_2, effect the equilibria in the following ways:*

- *For $m_1 \geq d^*_{1,hi}$ and $m_2 \geq d^*_{2,hi}$. The capacity constraints have no effect on the equilibria.*
- *For $d^*_{1,lo} < m_1 < d^*_{1,hi}$ and $m_2 \geq d^*_{2,hi}$. The larger equilibrium $(d^*_{1,hi}, d^*_{2,hi})$ moves to $(m_1, r(m_1))$. The other two equilibria remain unchanged.*
- *For $d^*_{1,lo} < m_1 < d^*_{1,hi}$ and $d^*_{2,lo} < m_2 < d^*_{2,hi}$, the larger equilibrium $(d^*_{1,hi}, d^*_{2,hi})$ moves to either (m_1, m_2), $(m_1, r(m_1))$ or $(r(m_2), m_2)$. The other two equilibria remain unchanged.*
- *For $m_1 < d^*_{1,lo}$, only the trivial equilibrium $(0, 0)$ will remain.*

For the rest of the chapter, we'll assume $m_1 \geq d^*_{1,hi}$ and $m_2 \geq d^*_{2,hi}$.

Price of Stability

We defined the systems' utility as the sum of all individual utilities. For the 2 person game, this gives:

$$u_{sys} = -(d_1 + d_2) + b_{12} d_2 \frac{d_1}{1 + d_1} + b_{21} d_1 \frac{d_2}{1 + d_2}.$$

Differentiating over d_1 and d_2 gives

$$\begin{aligned} \frac{\partial u_{sys}}{\partial d_1} &= -1 + \frac{b_{12} d_2}{(1 + d_1)^2} + \frac{b_{21} d_2}{1 + d_2} \\ \frac{\partial u_{sys}}{\partial d_2} &= -1 + \frac{b_{12} d_1}{1 + d_1} + \frac{b_{21} d_1}{(1 + d_2)^2}. \end{aligned}$$

Unlike the homogeneous version, even for b_{12} and b_{21} large enough such that an equilibrium exist, and $d_1, d_2 \geq 1$, it is still possible that $\frac{\partial u_{sys}}{\partial d_i} < 0$, for b_{ij} small, d_i close to 1 and d_j big.

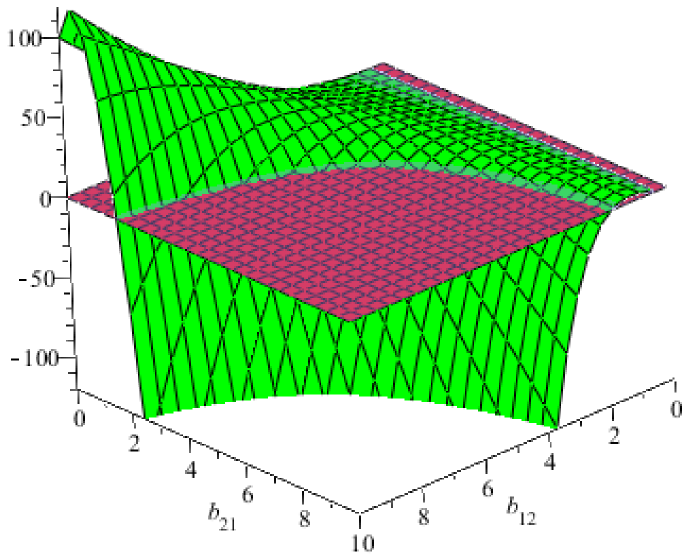

Figure 3.2: If $f(d_{1,min}) \leq 0$, the original equation has solutions

Therefore we can't find a general solution for the optimal performance.
However, for the system where $b_{12}, b_{21} \geq 2$, we see that $\frac{\partial u_{sys}}{\partial d_1} > 0$ and $\frac{\partial u_{sys}}{\partial d_2} > 0$ for $d_1, d_2 \geq 1$. Hence u_{sys} is increasing in d_1 and d_2.
For the heterogeneous system, we assume the maximal capacities are heterogeneous, so $m_1, m_2 \geq 1$ are independent. The optimal outcome for the system performance is reached for $d_1 = m_1, d_2 = m_2$. So the optimal system performance is

$$u_{sys}(m_1, m_2) = -(m_1 + m_2) + b_{12}m_2\frac{m_1}{1+m_1} + b_{21}m_1\frac{m_2}{1+m_2}.$$

The system performance in the best Nash equilibrium (d_1^*, d_2^*) is

$$u_{sys}(d_1^*, d_2^*) = -(d_1^* + d_2^*) + b_{12}d_2^*\frac{d_1^*}{1+d_1^*} + b_{21}d_1^*\frac{d_2^*}{1+d_2^*}.$$

Again we can define the Price of Stability:

$$POS = \frac{-(m_1 + m_2) + b_{12}\frac{m_1 m_2}{1+m_1} + b_{21}\frac{m_1 m_2}{1+m_2}}{-(d_1^* + d_2^*) + b_{12}\frac{d_1^* d_2^*}{1+d_1^*} + b_{21}\frac{d_1^* d_2^*}{1+d_2^*}}.$$

Just as in the homogenous model the bounds of the POS depends on the knowledge of the system architect. However, unlike the homogenous model, even with full information, the system architect can't force an optimal outcome if $\frac{m_1}{d_1^*} \neq \frac{m_2}{d_2^*}$.
$m_i = M_i/D_0$, so if the architect knows M_1, M_2, by resetting $D_0 = M_i/d_i^*$ he can force player i to contribute to maximum capacity. However, he can only do this for one player.
For $\frac{m_1}{d_1^*} < \frac{m_2}{d_2^*}$, in order to not upset the equilibrium, the architects best action would be to reset $D_0 = M_1/d_1^*$.

Example

Let $b_{12} = 10$ and $b_{21} = 2$.
Using the equation above, we find $d_1^* \approx 2.5427, d_2^* \approx 1.2551$ to be the best Nash equilibrium in the system, resulting in

$$\begin{aligned} u_1 &= -d_1^* + b_{12}\frac{d_1^* d_2^*}{1+d_1^*} \approx 6.4655 \\ u_2 &= -d_2^* + b_{21}\frac{d_1^* d_2^*}{1+d_2^*} \approx 1.5752 \\ u_{sys} &= u_1 + u_2 \approx 8.0408. \end{aligned}$$

Let $m_1 = 5$ and $m_2 = 3$.
Without interference from the system architect, we find that u_{sys} is optimal for $d_1 = m_1 = 5$ and $d_2 = m_2 = 3$.

$$u_{sys}(5, 3) = -(5 + 3) + 10 \cdot \frac{5 \cdot 3}{1+5} + 2 \cdot \frac{5 \cdot 3}{1+3} = 24.5.$$

So $POS = 24.5/8.0408 \approx 3.047$.
If the architect did have full information, his best action would be to set D_0 s.t. $m_1 = d_1^*$. Now $m_1 = d_1^* = 2.5427$ and $m_2 = 3/5 \cdot 2.5427 = 1.5256$.

$$\begin{aligned} u_{sys}(m_1, m_2) &= -(2.5427 + 1.5256) + 10 \cdot \frac{2.5427 \cdot 1.5256}{1 + 2.5427} + 2 \cdot \frac{2.5427 \cdot 1.5256}{1 + 1.5256} \\ &\approx 9.9534. \end{aligned}$$

This gives $POS = 9.9534/8.0408 \approx 1.238$, significantly better than without interference.

3.4.2 N person game

For $\alpha = 1$ the fixed point equations can be derived from the two player game as

$$d_i^* = \Big[\sum_{j \neq i} b_{ij} d_j^*\Big]^{1/2} - 1.$$

It is not possible to solve these equations analytically. In [5] an iterative learning model is used to solve the system of equations.
To mimic the interaction of users in a real P2P system, in this learning algorithm any peer P_i only interacts with a small numbers of peers, who contribute files that are of interest to P_i. When he interacts with these peers, P_i learns of their contributions and can adjust his own contribution accordingly. Obviously this is not globally optimal, because P_i only has information concerning the peers he interacts with. However, after P_i has set his contribution, this information is passed to the peers he interacts with, who will adjust their contributions again, and so on. Assuming there are no subsets with less than N players that only interact with each other, the actions of any peer P_i will eventually be known to all other peers.
To start, all peers have some random set of contributions and in every single iteration of the algorithm, every peer P_i determines his optimal contribution d_i given the contributions of d_{-i} of other peers and the values of b_{ij}. When this iterative process converges to a stable point, a Nash equilibrium is reached.
By simulating this in a numerical experiment Buragohain et al. were able to demonstrate that for sufficient benefits b_{ij} the iterative learning process does converge to a Nash Equilibrium for a heterogeneous system, even if b_{ij} is only non-zero for a small number of j's for each peer. In this simulation, the values b_{ij} are constant and are chosen from a Gamma distribution. Initial values of d_i are chosen from Gaussian distribution (excluding possible negative values).

3.4.3 Results

For $N = 1000$ and $b_{av} = 6.0$, the iterative learning process in [5] found a Nash equilibrium were the equilibrium values d_i^* are distributed in a bell shaped distribution with a mean of $d_{av}^* = 3.68$. In comparison, the corresponding homogeneous system, where $b_{ij} = b = b_{av}/(N-1)$ for all $i \neq j$ gives an equilibrium

with values $d^*_{hi} = 3.73$, only $1,5\%$ away from d^*_{av}.
As one would expect, for average benefits $b_{av} < b_c$, like the homogenous system, the algorithm converges to $d_i = 0$. From their series of experiments it shows that these results are independent of system size. It also shows that the higher the average benefit, the faster the convergence to equilibrium occurs. The closer the average gets to b_c, the slower the convergence. The convergence does depend on the initial values of d_i. For small values, the iterations converge to zero, i.e. everybody stops contributing and the system collapses.
Another point of interest is the effect that peers leaving the system have on the remaining peers. Again the simulations show what one would expect, namely that while the fraction of peers in the system lowers, the benefits of the remaining peers lowers and therefore their contribution levels will decrease. If this continues, at some point the benefits will become too low for the peers and the whole system will collapse.
For high benefits, the systems can be quite robust, as they've shown that for $(b_{av} - b_c)/b_c = 2.0$, the system can survive until 2/3 of the peers leave the system.
As the system grows bigger in size, it becomes more robust to random fluctuations.

3.5 Conclusion

In this chapter we examined a model with a reputation based incentive mechanism implemented. From the start it is clear that without such a differential service mechanism no player would ever contribute to the system.
For the homogeneous 2-person game, we found that for $b > 4$ there is a stable equilibrium where $d^* > 1$. Depending on the knowledge of the system architect, this equilibrium either gives a *Price of Stability* depending on m or, if the architect has complete information, $POS = 1$.
The results we found in the 2-person game were transferrable to the N-person homogeneous game, so that for $b(N-1) > 4$ we again find a stable equilibrium with $d^* \geq 1$.
For the 2-person heterogeneous system, we found that for the right constraints, the system again gives an equilibria s.t. $d^*_1 d^*_2 \geq 1$. Again the POS was dependent on the knowledge of the system architect, but even with complete information available, he can't enforce the optimal outcome.
For the N-person heterogeneous game it seemed that for reasonable values of b_{ij}, an equilibrium was found that approaches the values of the corresponding homogeneous system, e.g. for $b_{av} = 6$, they found an equilibrium s.t. $d^*_{av} = 3.68$, just 1.5% smaller than the corresponding homogeneous system where $b_{ij} = b_{av}/(N-1)$ gives an equilibrium with $d^*_{hi} = 3.73$ [5].

Chapter 4

Burden Sharing

In this chapter, we will shift our focus to the *sharing* in file sharing. Instead of focussing on the availability of files, we'll look at the transfer of these files, i.e. uploading and downloading.
Again this creates a tradeoff between benefit and cost. Here, downloading files results in a benefit for the peer (why else would he download the files), while the uploading comes with a cost in the form of bandwidth and CPU usage.
Naturally, rational peers will want to download as much as they want, while minimizing their uploads (their costs). In this model, due to Feldman et al. (2004) [7], we assume rational peers.
However, in practice we find that most users of P2P-systems don't mind incurring *some* (usually low) cost. Because most internet service providers (ISP's) reserve some of the bandwidth exclusively for uploading, this is bandwidth they would generally have no use for, so don't mind contributing. Here we model this form of altruism with a *generosity level* θ_i for each peer i. We say that θ_i is the maximum cost peer i is willing to contribute to the system.
The whole load of the system is divided evenly among the peers who do contribute. All peers have two options: they can either contribute to the system or free-ride.
We let x be the fraction of peers contributing at a particular moment, which we'll call the *contribution level*. We assume homogeneity in the amount each peer wants to receive, and normalize this to 1.
We also assume homogeneity for the benefit peers receive from the system and let the benefits depend on the contribution level, so that the benefit of peer i is $Q = \alpha x^{\beta}$, for some $\alpha \geq 1$ and $\beta \leq 1$.
We assume a sufficiently large population in the system, so that instead of talking about the number of people following a certain strategy, we'll use the fraction of the population who are following that strategy.

4.1 Standard model

We start with the basic model. Here, we have no incentive mechanism, so as long as the system doesn't collapse, everyone will receive full 'service'. The *burden* R that is placed on all contributing peers, is the inverse of the total percentage of contributors, so $R = 1/x$.
We assume a peer will contribute if his generosity level θ_i is greater than the burden R. Therefore the decision of all peers will be:

$$\text{Contribute, if } \theta_i \geq 1/x$$
$$\text{Free-ride, otherwise.}$$

Every moment in time, peers make a new decision whether to contribute or to free-ride. Those decisions lead to a new contribution level x, which in turn will lead to a new burden R, which in turn will lead to new decisions, and so on.
In this system, the contribution level x is in equilibrium at the intersection of the type distribution $x = P(\theta_i \geq \theta)$ with the curve $\theta = 1/x$. That is, if the present contribution level x creates a burden $R = 1/x$ such that the percentage of players deciding to contribute based on that burden, $P(\theta_i \geq 1/x)$, is again equal to x.
We find these equilibria by solving the fixed point equation $x = P(\theta \geq 1/x)$. For the rest of this chapter, we'll assume a population with generosity types uniformly distributed between 0 and θ_m.
So $\theta_i \sim U(0, \theta_m)$, $P(\theta_i \geq 1/x) = 1 - \frac{1}{x\theta_m}$.
For the Nash equilibria, this gives the equation $x = P(\theta \geq 1/x) = 1 - \frac{1}{x\theta_m}$, which has the solutions

$$x_{1,2} = \frac{\theta_m \pm \sqrt{\theta_m^2 - 4\theta_m}}{2\theta_m}. \tag{4.1}$$

This leads us to the following result.

Theorem 4.1. *The number of Nash equilibria depends on θ_m. $x = 0$ is always a Nash equilibrium. The additional non-zero Nash equilibria are the solutions of* (4.1)*, such that:*

- *For $\theta_m > 4$, there are 3 Nash equilibria.*
- *For $\theta_m = 4$, there are 2 Nash equilibria.*
- *For $\theta_m < 4$, there is 1 Nash equilibrium.*

Proof. Independently of θ_m, for $x = 0$, the best response for all peers, will be clearly be to not contribute, so the contribution level $x = 0$ is a Nash equilibrium. The other Nash equilibria are found when $x = 1 - \frac{1}{x\theta_m}$. This gives

$$x_{1,2} = \frac{\theta_m \pm \sqrt{\theta_m^2 - 4\theta_m}}{2\theta_m}.$$

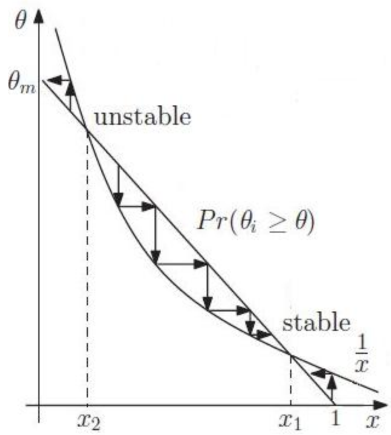

Figure 4.1: Without Incentives

For $\theta_m > 4$, this equation has 2 solutions, so there exist 3 Nash equilibria.
For $\theta_m = 4$, the equation has only 1 solution ($x = \frac{1}{2}$), so there exist 2 Nash equilibria.
For $\theta_m < 4$, the equation has no solutions, so the only Nash equilibrium is $x = 0$. $\square$

In the figure above the equilibria are shown for $\theta_m > 4$, at the intersections between $P(\theta_i \geq \theta)$ and $1/x$.

Stability

In a dynamical system $x = f(x)$, such as this, an equilibrium x_i is stable, if $|f'(x)| < 1$ for $x = x_i$. If $|f'(x)| > 1$ for $x = x_i$, x_i is unstable [14].
Here, we have $x = f(x) = 1 - \frac{1}{x\theta_m}$, so $f'(x) = \frac{1}{x^2\theta_m}$.
For $\theta_m > 4$, we have two distinct non-zero equilibria. First we'll look at $x_1 = \frac{\theta_m + \sqrt{\theta_m^2 - 4\theta}}{2\theta}$.

$$
\begin{aligned}
f'(x_1) &= \frac{1}{\left(\frac{\theta_m + \sqrt{\theta_m^2 - 4\theta_m}}{2\theta_m}\right)^2 \theta_m} \\
&= \frac{4\theta_m}{\left(\theta_m + \sqrt{\theta_m^2 - 4\theta_m}\right)^2}.
\end{aligned}
$$

$\theta_m > 4$, so $f'(x) > 0$. To show that x_1 is stable, we need to show that $f'(x) < 1$, i.e.

$$\left(\theta_m + \sqrt{\theta_m^2 - 4\theta_m}\right)^2 > 4\theta_m.$$

Calculating, we find

$$\begin{aligned}
\left(\theta_m + \sqrt{\theta_m^2 - 4\theta_m}\right)^2 &> 4\theta_m \\
\theta_m^2 + 2\theta_m\sqrt{\theta_m^2 - 4\theta_m} + \theta_m^2 - 4\theta_m &> 4\theta_m \\
2(\theta_m^2 - 4\theta_m) + 2\theta_m\sqrt{\theta_m^2 - 4\theta_m} &> 0.
\end{aligned}$$

Because $\theta_m > 4$, this clearly holds, so x_1 is a stable equilibrium.
Secondly, we'll look at $x_2 = \frac{\theta_m - \sqrt{\theta_m^2 - 4\theta_m}}{2\theta_m}$.

$$\begin{aligned}
f'(x_2) &= \frac{1}{\left(\frac{\theta_m - \sqrt{\theta_m^2 - 4\theta_m}}{2\theta_m}\right)^2 \theta_m} \\
&= \frac{4\theta_m}{\left(\theta_m - \sqrt{\theta_m^2 - 4\theta_m}\right)^2}.
\end{aligned}$$

x_2 is unstable if $f'(x_2) > 1$.

$$\begin{aligned}
\frac{4\theta_m}{\left(\theta_m^2 - \sqrt{\theta_m^2 - 4\theta_m}\right)^2} &> 1 \\
4\theta_m &> \left(\theta_m - \sqrt{\theta_m^2 - 4\theta_m}\right)^2 \\
4\theta_m &> \theta_m^2 - 2\theta_m\sqrt{\theta_m^2 - 4\theta_m} - \theta_m^2 + 4\theta_m \\
0 &> -2\theta_m\sqrt{\theta_m^2 - 4\theta_m}.
\end{aligned}$$

Again, this is clearly the case, so x_2 is an unstable equilibrium.
For $\theta_m = 4$, the only equilibrium is $x = 1/2$. This gives $f'(x) = 1$, which is inconclusive. Further inspection finds that x converges to $x = 1/2$ if we start at $x > 1/2$, but for $x < 1/2$, the system diverges from $x = 1/2$ and will converge to 0.

Claim 4.1. *The stable nonzero equilibrium contribution level* (x_1) *increases in* θ_m *and converges to 1 as* $\theta_m \to \infty$*, but converges to zero when* $\theta_m < 4$.

Now that we have found the equilibria, we will look at the performance of the system W_{sys}. We will define this as the difference between the total benefits received by all peers and the total contribution cost incurred by all peers. By normalizing the network size to 1, for $x > 0$ we have

$$W_{sys} = \alpha x^\beta - (1/x)x = \alpha x^\beta - 1.$$

Price of Stability

To calculate the quality of equilibria, we can choose between the *Price of Stability* and the *Price of Anarchy*. Of the three equilibria, the bigger non-zero

equilibrium x_1 yields the best system performance. x_1 is also a stable equilibrium, therefore it seems logical to choose the *Price of Stability* as our measuring method.
Furthermore, because $x = 0$ is also an equilibrium, the Price of Anarchy only gives us a trivial bound ($POA = \infty$). The average burden is 1, no matter the level of participation, so it is obvious that W_{sys} is increasing in x and optimal for $x = 1$. For $\beta = 1$, this gives:

$$\begin{aligned} POS &= \frac{\alpha - 1}{\alpha^{\frac{\theta_m + \sqrt{\theta_m^2 - 4\theta_m}}{2\theta_m}} - 1} \\ &= \frac{2\theta_m(\alpha - 1)}{(\alpha - 2)\theta_m + \alpha\sqrt{\theta_m^2 - 4\theta_m}}. \end{aligned}$$

Example

For $\alpha = 4$ and $\theta_m = 6$, the equilibria of this system become

$$x_{1,2} = \frac{6 \pm \sqrt{12}}{12}$$

Where x_1 is the better equilibrium:

$$x_1 = \frac{6 + \sqrt{12}}{12} \approx 0,789.$$

The system performance in x_1 is

$$W_{sys} = 4\frac{6 + \sqrt{12}}{12}1 \approx 2,155.$$

The price of stability now becomes:

$$POS = \frac{4 - 1}{4^{\frac{6+\sqrt{12}}{12}} - 1} \approx 1,392.$$

4.1.1 Alternative distribution

We again consider the P2P models

Contribute, if $\theta_i > 1/x$
Free-ride, otherwise.

But now we assume that the generosity of the peers is distributed as follows: a fraction ϕ of the peers have their type θ_i uniformly distributed between 0 and θ_m, a fraction $(1 - \phi)/2$ are of type $\theta_i = 0$, so will never contribute, and the remaining $(1 - \phi)/2$ are of type $\theta_i = \theta_m$.
We would now like to know how this different distribution will affect the resulting equilibria.

We find the equilibrium at the point where the curve $x = 1/\theta$ intersects with $x = P(\theta_i \geq \theta)$, so we again need to solve the equation $x = P(\theta_i \geq 1/x)$.
For $1/x \leq \theta_m$ this gives

$$\begin{aligned} x = P(\theta_i \geq 1/x) &= \Big(1 - \frac{1}{x\theta_m}\Big) \cdot \phi + \frac{1-\phi}{2} \\ &= \phi - \frac{\phi}{x\theta_m} + \frac{1-\phi}{2} \\ &= \frac{1+\phi}{2} - \frac{\phi}{x\theta_m} \\ x^2 &= \Big(\frac{1+\phi}{2}\Big)x - \frac{\phi}{\theta_m} \\ 0 &= \theta_m x^2 - \Big(\frac{1+\phi}{2}\Big)\theta_m x + \phi. \end{aligned}$$

This gives the solutions

$$\begin{aligned} x_{1,2} &= \frac{(\frac{1+\phi}{2})\theta_m \pm \sqrt{\frac{(1+\phi)^2}{4}\theta_m^2 - 4\phi\theta_m}}{2\theta_m} \\ &= \frac{(\frac{1+\phi}{2})\theta_m \pm \sqrt{\frac{(1+\phi)^2}{4}\theta_m^2 - 4\phi\theta_m}}{2\theta_m} \\ &= \frac{1+\phi}{4} \pm \sqrt{\frac{(1+\phi)^2}{16} - \frac{\phi}{\theta_m}}. \end{aligned} \tag{4.2}$$

The bigger equilibrium, with contribution level (x_1) now increases in θ_m and converges to $\frac{1+\phi}{2}$ as θ_m tends to ∞.
The fraction of peers with generosity type $\theta_i = \theta_m$ (size $\frac{1-\phi}{2}$) either all contribute, or none of them contribute. Therefore, equilibria can only exist for $x \geq (1-\phi)/2$. This gives us the following results.

Theorem 4.2. *The number of Nash equilibria depends on ϕ and θ_m.*
$x = 0$ is always a Nash equilibrium. The additional non-zero Nash equilibria are solutions of (4.2), for $x \geq (1-\phi)/2$, such that:

- *For $\theta_m \geq \frac{2}{1-\phi}$, there are 2 Nash equilibria for all $\phi \in (0, 1)$.*
- *For $\frac{2}{1-\phi} > \theta_m > \frac{16\phi}{(1+\phi)^2}$, there are 3 Nash equilibria for $\phi \geq 1/3$ and 1 Nash equilibrium for $\phi < 1/3$.*
- *For $\theta_m = \frac{16\phi}{(1+\phi)^2}$ there are 2 Nash equilibria for $\phi \geq 1/3$ and 1 Nash equilibrium for $\phi < 1/3$.*
- *For $\theta_m < \frac{16\phi}{(1+\phi)^2}$ there is 1 Nash equilibrium for all $\phi \in (0, 1)$.*

Proof. Independently of θ_m and ϕ, $x = 0$ is always a Nash equilibrium.
The other Nash equilibria are found when $x = (1 - \frac{1}{x\theta_m}) \cdot \phi + \frac{1-\phi}{2}$, where $x \geq \frac{1-\phi}{2}$. For $\theta_m > \frac{16\phi}{(1+\phi)^2}$, (4.2) has two solutions.

$$\begin{aligned} x_1 &= \frac{1+\phi}{4} + \sqrt{\frac{(1+\phi)^2}{16} - \frac{\phi}{\theta_m}} \\ x_2 &= \frac{1+\phi}{4} - \sqrt{\frac{(1+\phi)^2}{16} - \frac{\phi}{\theta_m}}. \end{aligned}$$

If $\theta_m \geq \frac{2}{1-\phi}$, then $\theta_m \geq \frac{16\phi}{(1+\phi)^2}$:

$$\begin{aligned} \frac{2}{1-\phi} &\geq \frac{16\phi}{(1+\phi)^2} \\ 2(1+\phi)^2 &\geq 16\phi(1-\phi) \\ 18\phi^2 - 12\phi + 2 &\geq 0 \\ 2(3\phi - 1)^2 &\geq 0. \end{aligned}$$

So (4.2) has two solutions.

$$\begin{aligned} x_1 &= \frac{1+\phi}{4} + \sqrt{\frac{(1+\phi)^2}{16} - \frac{\phi}{\theta_m}} \\ &\geq \frac{1+\phi}{4} + \sqrt{\frac{(1+\phi)^2}{16} - \frac{\phi(1-\phi)}{2}} \\ &= \frac{1+\phi}{4} + \sqrt{\frac{(1-3\theta)^2}{16}} \\ &= \frac{1-\phi}{2}. \end{aligned}$$

This means $x_1 \geq \frac{1-\phi}{2}$, hence a Nash equilibrium.

$$\begin{aligned} x_2 &= \frac{1+\phi}{4} - \sqrt{\frac{(1+\phi)^2}{16} - \frac{\phi}{\theta_m}} \\ &\leq \frac{1+\phi}{4} + \sqrt{\frac{(1+\phi)^2}{16} - \frac{\phi(1-\phi)}{2}} \\ &= \frac{1-\phi}{2}. \end{aligned}$$

$x_2 \leq \frac{1-\phi}{2}$, so no equilibrium, unless $x_2 = \frac{1-\phi}{2}$, in which case $x_1 = x_2$.
$x_2 \leq \frac{1-\phi}{2} \leq x_1$, so there are 2 Nash equilibria for $\theta_m \geq \frac{2}{1-\phi}$.

If $\frac{2}{1-\phi} > \theta_m > \frac{16\phi}{(1+\phi)^2}$, (4.2) has again 2 solutions.

For $\phi \geq 1/3$ we have

$$\begin{aligned} x_1 &= \frac{1+\phi}{4} + \sqrt{\frac{(1+\phi)^2}{16} - \frac{\phi}{\theta_m}} \\ &> \frac{1+\phi}{4}. \end{aligned}$$

this means $x_1 > \frac{1-\phi}{2}$, so this is a Nash equilibrium.
Using $\theta_m < \frac{2}{1-\phi}$

$$\begin{aligned} x_2 &= \frac{1+\phi}{4} - \sqrt{\frac{(1+\phi)^2}{16} - \frac{\phi}{\theta_m}} \\ &> \frac{1+\phi}{4} - \sqrt{\frac{(3\phi-1)^2}{16}} = \frac{1+\phi}{4} - \frac{3\phi-1}{4} \\ &= \frac{1-\phi}{2}. \end{aligned}$$

So x_2 is also a Nash equilibrium.
We now have $\frac{1-\phi}{2} < x_2 < x_1$.
Hence for $\frac{2}{1-\phi} > \theta_m > \frac{16\phi}{(1+\phi)^2}$ and $\phi \geq 1/3$ there are 3 Nash equilibria.

For $\frac{2}{1-\phi} > \theta_m > \frac{16\phi}{(1+\phi)^2}$ and $\phi < 1/3$, using $\theta_m < \frac{2}{1-\phi}$, we get

$$\begin{aligned} x_1 < \frac{1+\phi}{4} + \sqrt{\frac{(1-3\phi)^2}{16}} &= \frac{1+\phi}{4} + \frac{1-3\phi}{4} \\ &= \frac{1-\phi}{2}. \end{aligned}$$

This means x_1 is no equilibrium.

$$\begin{aligned} x_2 &= \frac{1+\phi}{4} - \sqrt{\frac{(1+\phi)^2}{16} - \frac{\phi}{\theta_m}} \\ &< \frac{1+\phi}{4}. \end{aligned}$$

This means x_2 is also no equilibrium. Hence the only equilibrium for $\frac{2}{1-\phi} > \theta_m > \frac{16\phi}{(1+\phi)^2}$ and $\phi < 1/3$ is $x = 0$.

For $\theta_m = \frac{16\phi}{(1+\phi)^2}$, (4.2) has 1 solution.

$$x = \frac{1+\phi}{4}.$$

If $\phi \geq 1/3$, $x \geq \frac{1-\phi}{2}$, so this is a Nash equilibrium. Hence there are 2 Nash equilibria for $\theta_m = \frac{16\phi}{(1+\phi)^2}$, $\phi \geq 1/3$.

If $\phi < 1/3$, $x < \frac{1-\phi}{2}$, so this is no equilibrium and the only Nash equilibrium is $x = 0$.
Finally, for $\theta_m < \frac{16\phi}{(1+\phi)^2}$, (4.2) has no solutions. Hence for all $\phi \in (0,1)$, the only Nash equilibrium is $x = 0$. □

We see that for this system, in order to find non-zero equilibria we require lower maximum generosity levels θ_m than in the previous section, however the best equilibria we do find are strictly worse than those found in the previous section.

$$\frac{1+\phi}{4} - \sqrt{\frac{(1+\phi)^2}{16} - \frac{\phi}{\theta_m}} < \frac{1}{2} + \sqrt{\frac{1}{4} - \frac{1}{\theta_m}}, \text{ for } \phi < 1. \tag{4.3}$$

Stability

We show the (in)stability of the equilibria for the case where we have two distinct non-zero equilibria the same way as in the previous section.
Here we have $x = f(x) = \left(1 - \frac{1}{x\theta_m}\right)\phi + \frac{1-\phi}{2}$, so

$$f'(x) = \frac{1}{x^2\theta_m} \cdot \phi.$$

First we examine the equilibrium $x_1 = \frac{1+\phi}{4} + \sqrt{\frac{(1+\phi)^2}{16} - \frac{\phi}{\theta_m}}$.

$$f'(x_1) = \frac{1}{\left(\frac{1+\phi}{4} + \sqrt{\frac{(1+\phi)^2}{16} - \frac{\phi}{\theta_m}}\right)^2 \theta_m}\phi.$$

To show that x_1 is stable, we show that

$$\begin{aligned}
\Big(\frac{1+\phi}{4} + \sqrt{\frac{(1+\phi)^2}{16} - \frac{\phi}{\theta_m}}\Big)^2\theta_m &> \phi \\
\Big(\frac{(1+\phi)^2}{16} + 2\big(\frac{1+\phi}{4}\big)\sqrt{\frac{(1+\phi)^2}{16} - \frac{\phi}{\theta_m}} + \frac{(1+\phi)^2}{16} - \frac{\phi}{\theta_m}\Big) &> \phi \\
\frac{(1+\phi)^2}{8}\theta_m - \phi + 2\big(\frac{1+\phi}{4}\big)\sqrt{\frac{(1+\phi)^2}{16} - \frac{\phi}{\theta_m}} &> \phi \\
(1+\phi)^2\theta_m + 4\theta_m(1+\phi)\sqrt{\frac{(1+\phi)^2}{16} - \frac{\phi}{\theta_m}} &> 16\phi.
\end{aligned}$$

Because $\theta_m(1+\phi)^2 > 16\phi$, this clearly holds, so x_1 is a stable equilibrium.
Next, we examine the equilibrium $x_2 = \frac{1+\phi}{4} - \sqrt{\frac{(1+\phi)^2}{16} - \frac{\phi}{\theta_m}}$. We have

$$f'(x_2) = \frac{1}{\left(\frac{1+\phi}{4} - \sqrt{\frac{(1+\phi)^2}{16} - \frac{\phi}{\theta_m}}\right)^2 \theta_m}\phi.$$

We show that x_2 is unstable, by showing $f'(x_2) > 1$, i.e.

$$\begin{aligned}
\Big(\frac{1+\phi}{4} - \sqrt{\frac{(1+\phi)^2}{16} - \frac{\phi}{\theta_m}}\Big)^2\theta_m &< \phi \\
\Big(\frac{(1+\phi)^2}{16} - 2\big(\frac{1+\phi}{4}\big)\sqrt{\frac{(1+\phi)^2}{16} - \frac{\phi}{\theta_m}} - \frac{(1+\phi)^2}{16} + \frac{\phi}{\theta_m}\Big)\theta_m &< \phi \\
-2\theta_m\big(\frac{1+\phi}{4}\big)\sqrt{\frac{(1+\phi)^2}{16} - \frac{\phi}{\theta_m}} + \phi &< \phi \\
-2\theta_m\big(\frac{1+\phi}{4}\big)\sqrt{\frac{(1+\phi)^2}{16} - \frac{\phi}{\theta_m}} &< 0.
\end{aligned}$$

Again, this holds, so x_2 is an unstable equilibrium.

4.2 Penalty Mechanism

To improve upon our previous model, in this section we'll add an incentive mechanism. We use a reputation based mechanism where a penalty is posed, based on the actions of the players. The assumption here, is that while user type is not observable, user behaviour is.
We introduce a penalty p for free-riders. I will treat this as a service differentiation policy, where free-riders will receive a reduced benefit of $(1-p)Q$. If the system can only find free-riders with a low probability, another interpretation is that p represents the probability that a free-rider gets caught and subsequently excluded from the system.
This incentive mechanism will increase user contribution in two ways. Firstly, it reduces the system load and thus the burden R imposed on contributors because free-riders only receive $(1-p)$ of service. The load of the system decreases to $x + (1-x)(1-p)$, where $(1-x)$ is the percentage of peers free-riding. The contribution cost, therefore lowers to $\frac{x+(1-x)(1-p)}{x}$. Secondly, it introduces a *threat* T, because peers know they will receive reduced service (or expulsion) if caught free-riding.
This threat is the potential loss of service, so $T = p\alpha x^\beta$. In this altered model, we get the following performances for the peers. The performance of a contributor will be $Q - R$ while the performance of a free-rider will be $Q - T$:

$$\begin{aligned}
W_{contributor} &= Q - R = & \alpha x^\beta - \frac{x + (1-x)(1-p)}{x} \\
W_{free-rider} &= Q - T = & \alpha x^\beta - p\alpha x^\beta.
\end{aligned}$$

Players want to optimize their performance, so now they are willing to contribute as long as the extra costs of contributing (the burden, R) are smaller than the

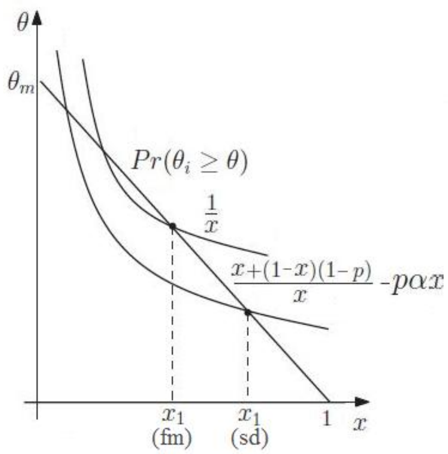

Figure 4.2: With Service Differentiation

potential loss in performance by the reduced service (the threat, T). This means the contribution level can now be derived from

$$x = P(\theta_i \geq R - T)$$

$$x = P(\theta_i \geq \frac{x + (1-x)(1-p)}{x} - p\alpha x^{\beta}).$$

For $\beta = 1$, this gives

$$\begin{aligned} x &= 1 - \Big(\frac{x + (1-x)(1-p)}{x} - p\alpha x\Big)\frac{1}{\theta_m} \\ (1-x)\theta_m &= \frac{x + (1-x)(1-p)}{x} - p\alpha x \\ (1-x)x\theta_m &= x + (1-x)(1-p) - p\alpha x^2 \\ 0 &= (p\alpha - \theta_m)x^2 + (\theta_m - p)x + (p-1). \end{aligned}$$

Resulting in equilibria for the contribution levels:

$$x_{1,2}(p) = \frac{(\theta_m - p) \pm \sqrt{p^2 + 2\theta_m p + \theta_m^2 - 4\theta_m + 4p\alpha - 4p^2\alpha}}{2(\theta_m - p\alpha)}.$$

These equilibria exist if $\theta_m^2 + (2p-4)\theta_m + p(p + 4\alpha - 4p\alpha) \geq 0$.
This gives

$$\begin{aligned} \theta_m &\geq \frac{(4-2p) + \sqrt{(4-2p)^2 - 4p(p + 4\alpha(1-p))}}{2} \\ \theta_m &\geq 2 - p + 2\sqrt{1 - p - \alpha p + \alpha p^2} \end{aligned}$$

and subsequently, $1 - p - \alpha p + \alpha p^2 \geq 0$ gives $p \leq 1/\alpha$.
In figure 4.2 these equilibria are shown together with the equilibria from the standard model. Here $x_1(fm)$ represents the best equilibrium in the standard, free market model, while $x_1(sd)$ is the best equilibrium in the system with a penalty mechanism and service differentiation.
We find that if the benefits of participating in the system are high (i.e. α is large), a service differentiation policy that imposes a small performance penalty on free riders is sufficient to induce a high level of cooperation (independent of the maximal generosity level θ_m). This is especially important to our alternate interpretation. Even if we can only catch and exclude free-riders with a small probability, a high contribution level will be attained.
This penalty yields a social benefit due to the higher contribution level it achieves, but incurs some social cost in the form of reduced benefits to free-riders. However, if we set p high enough such that it achieves full cooperation ($x = 1$) the mechanism incurs no social cost. The penalty would merely function as a threat, because no one would free-ride. In that case the optimal system performance is achieved ($Q_{max} = \alpha$).
We find the required p by substituting $x = 1$ in $x = 1 - \left(\frac{x+(1-x)(1-p)}{x} - p\alpha x\right)\frac{1}{\theta_m}$ we find

$$\begin{aligned} 1 &= 1 - (1 - p\alpha)\frac{1}{\theta_m} \\ p\alpha &= 1 \\ p &= \frac{1}{\alpha}. \end{aligned} \tag{4.4}$$

So, with a penalty of $p = 1/\alpha$ we realize full cooperation.
Now that we know the equilibria of the system, we find the system performance.

$$W_{sys}(p) = x(Q - R) + (1 - x)(Q - T) = (\alpha x - 1)(x + (1 - x)(1 - p)).$$

For $p = 1/\alpha$, $x = 1$, so no one free-rides and subsequently, no one receives reduced benefits. This results in a system performance in $x = 1$ of $W_{sys} = \alpha - 1$, which is the optimal result. This means the *Price of Stability* is 1, $POS = 1$.
For $p < 1/\alpha$, a situation we logically could only find ourselves in if we use the alternative interpretation of our penalty mechanism, we get the general result:

$$POS = \frac{\alpha - 1}{(\alpha x(p) - 1)(x(p) + (1 - x(p))(1 - p))}. \tag{4.5}$$

Example

We use the same parameters as in the standard model, $\theta_m = 6$, $\alpha = 4$ and we assume the chance of catching a free-rider is $p = 0.1$.
First we calculate the contribution level in our best equilibrium:

$$x(0.1) = \frac{(6 - 0.1) + \sqrt{0.1^2 + 1.2 + 36 - 24 + 1.6 - 0.16}}{2(6 - 0.4)} \approx 0.869.$$

Knowing this, we can calculate the system performance in the equilibrium:

$$W_{sys} = (4 \cdot x(0.1) - 1)(x(0.1) + (1 - x(0.1))(1 - 0.1)) \approx 2.442.$$

Finally, we find the Price of Stability:

$$POS = \frac{4 - 1}{(4 \cdot x(0.1) - 1)(x(0.1) + (1 - x(0.1))(1 - 0.1))} \approx 1.229.$$

This is still better than the POS in the standard model (which gave us a POS of 1.392).

4.2.1 Basin of Attraction

Although the right penalty level ($p = 1/\alpha$), can give an equilibrium where full cooperation is attained ($x = 1$), this all has little value if in practice that equilibrium is rarely reached. Therefore we look at the *basin of attraction*, defined as the set of initial conditions for which the system will in time converge to the stable equilibrium, $x = 1$.
For this system, that means the set $[1 - \epsilon, 1]$ of initial contribution levels, for which x converge to $x = 1$. The basin of attraction is dependent on θ_m. A low θ_m value may lead to an extremely small ϵ, which means the system will never converge to $x = 1$, unless it actually starts there.
The threshold value of θ_m, above which the system converges to $x = 1$ is a function of α and ϵ.
To find this value we'll first rewrite $x_{1,2}(p)$ for $p = \frac{1}{\alpha}$.

$$\begin{aligned} x_{1,2}(1/\alpha) &= \frac{(\theta_m - \frac{1}{\alpha}) \pm \sqrt{(\theta_m + \frac{1}{\alpha})^2 - 4(\theta_m + \frac{1}{\alpha}) + 4}}{2(\theta_m - 1)} \\ &= \frac{(\theta_m - \frac{1}{\alpha}) \pm \sqrt{((\theta_m + \frac{1}{\alpha}) - 2)^2}}{2(\theta_m - 1)} \\ &= \frac{(\theta_m - \frac{1}{\alpha}) \pm ((\theta_m + \frac{1}{\alpha}) - 2)}{2(\theta_m - 1)}. \end{aligned}$$

We find

$$\begin{aligned} x_1(1/\alpha) &= \frac{2(\theta_m - 1)}{2(\theta_m - 1)} = 1 \\ x_2(1/\alpha) &= \frac{2 - \frac{2}{\alpha}}{2(\theta_m - 1)} = \frac{1 - \frac{1}{\alpha}}{\theta_m - 1}. \end{aligned}$$

x converges to 1 for $x > x_2$, so we find the threshold value by solving

$$1 - \epsilon = \frac{1 - \frac{1}{\alpha}}{\theta_m - 1}.$$

Doing this, gives

$$\begin{aligned} 1-\epsilon &= \frac{1-\frac{1}{\alpha}}{\theta_m - 1} \\ (1-\epsilon)(\theta_m - 1) &= \frac{\alpha - 1}{\alpha} \\ \theta_m - 1 &= \frac{\alpha - 1}{\alpha(1-\epsilon)} \\ \theta_m &= \frac{\alpha - 1}{\alpha(1-\epsilon)} - \frac{\alpha(1-\epsilon)}{\alpha(1-\epsilon)} \\ \theta_m &= \frac{2\alpha - \alpha\epsilon - 1}{\alpha(1-\epsilon)}. \end{aligned}$$

So the threshold value, for which x converges to 1 for the initial $x \in [1-\epsilon, 1]$ is

$$\theta_m^{\text{threshold}} = \frac{2\alpha - \alpha\epsilon - 1}{\alpha(1-\epsilon)}. \tag{4.6}$$

As expected the threshold value increases in ϵ and α.

4.3 Turnover

Until now we've worked under the assumption of a closed network. No new peers entered the system, and no existing peers left.
In this section we'll adjust our model as to allow for turnover. This opens us up for another problem, *whitewashing*. While we've seen that the penalty mechanism is effective in discouraging free-riding, the availability of cheap pseudonyms can undermine it's effectiveness. As discussed, by whitewashing, one leaves the network and rejoins with a new identity, so avoiding the penalty imposed on free-riders. Players that enter the system after whitewashing are indistinguishable from genuine newcomers, so it's not possible to single them out for penalty.
For this dynamic model a turnover rate of d is assumed and the population is divided into four different groups, namely *Existing Contributors* (EC), *Existing Free-riders / Whitewashers* (EF/WW), *New Contributors* (NC) and *New Free-riders* (NF).
Two scenarios are considered: *permanent identities* (PI) and *free identities* (FI). Under PI, identity costs (the costs to create a new digital identity) are taken to be infinity, under FI there are no costs.
While a fraction of d users depart and are replaced by the same number of newcomers, a fraction $(1-d)(1-x_s)$ of the users whitewashes under FI, where x_s is the contribution level of peers who will stay in the system.
The significant difference in the PI and FI model is now the second group (EF/WW). If identities are permanent, free-riders have the choice to either stay in the system or leave it (but never come back), if identities are free, they will normally choose a whitewashing strategy. However, with a penalty of the right size imposed on newcomers, free-riders become indifferent between staying and

whitewashing.
If no penalty mechanism would be in place, the results of this model, would be identical to those of the standard model.

Because some users will leave the system at the end of each period, they are not affected by threat of a penalty they would have paid had they stayed in the system. Therefore we divide the peers in two groups: *leavers* and *stayers*.
This creates two separate contribution levels:

$$x_l : \textit{the contribution level of the 'leavers'}$$

$$x_s : \textit{the contribution level of the 'stayers'}.$$

These contribution levels follow the following pair of equations:

$$x_l = P(\theta_i \geq R)$$

$$x_s = P(\theta_i \geq R - T).$$

The average contribution level (x_a) will be:

$$x_a = dx_l + (1 - d)x_s.$$

The stayers are affected by the penalty mechanism in place, while the leavers are not, therefore $x_s \geq x_l$ and the average contribution level, x_a, is always greater than or equal to that of the contribution level of the leavers. Unlike the static system, where $x = 1$ can be achieved for a sufficiently high p, in dynamic scenarios it is not possible to achieve $x_a = 1$, due to these 'leavers'.

4.3.1 Permanent Identities

Under PI, we don't have to worry about whitewashing because identities are permanent (i.e. the cost of a new identity is ∞). Therefore it is unnecessary to penalize newcomers.
This means that the only group of peers, receiving reduced service are the existing free-riders (EF). This means that the burden under PI, becomes:

$$R_{PI} = \frac{(1-d)x_s + d + (1-d)(1-x_s)(1-p)}{x_a}.$$

A fraction $(1 - d)x_s$ are players who contributed last period and stayed in the system (EC) and a fraction d are newcomers (NC+NF), both these groups receive full service. A fraction $(1 - d)(1 - x_s)$ are the players who stayed in the system, but didn't contribute nonetheless (EF), they receive reduced benefits.
Based on our previous model, it seems plausible that with a high enough threat T, we can maximize the contribution level of the stayers to $x_s = 1$.
Because all stayers will contribute, the only free-riders will be those who are about to leave the system and can do so without repercussion.
This means that everyone in the system, receives full service. We use this to

find the equilibria. With that information we find the p for which all stayers will contribute:

$$\begin{aligned} x_s &= P(\theta_i \geq R - T) \\ 1 &= 1 - (\frac{1}{x_a} - p\alpha x_a)\frac{1}{\theta_m} \\ p\alpha x_a &= \frac{1}{x_a} \\ p &= \frac{1}{\alpha x_a^2}. \end{aligned} \tag{4.7}$$

We see that for large enough α, this p does in fact exist. With this information we find the contribution level for the leavers in equilibrium:

$$\begin{aligned} x_l = P(\theta_i \geq R_{PI}) &= P(\theta_i \geq \frac{1}{x_a}) \\ &= 1 - \frac{1}{x_a \theta_m}. \end{aligned} \tag{4.8}$$

Substituting $x_s = 1$ in $x_a = dx_l + (1-d)x_s$ gives, $x_a = dx_l + (1-d) = 1 - d(1-x_l)$ and we find:

$$\begin{aligned} x_l &= 1 - \frac{1}{(1-d(1-x_l))\theta_m} \\ (1-x_l)\theta_m &= \frac{1}{1-d(1-x_l)} \\ 0 &= (1-x_l)(1-d(1-x_l))\theta_m - 1 \\ 0 &= d\theta_m(1-x_l)^2 - \theta_m(1-x_l) + 1. \end{aligned}$$

Resulting in

$$\begin{aligned} (1-x_l) &= \frac{\theta_m \pm \sqrt{\theta_m^2 - 4d\theta_m}}{2d\theta_m} \\ x_l &= 1 - \frac{1 \pm \sqrt{1 - 4d/\theta_m}}{2d}. \end{aligned} \tag{4.9}$$

So the best equilibrium of the system is for contribution levels

$$\begin{aligned} x_s &= 1 \\ x_l &= 1 - \frac{1 - \sqrt{1 - 4d/\theta_m}}{2d}. \end{aligned}$$

This gives an average contribution of

$$\begin{aligned} x_a &= dx_l + (1-d)x_s \\ &= d\Big(1 - \frac{1 - \sqrt{1 - 4d/\theta_m}}{2d}\Big) + (1-d) \end{aligned}$$

$$= \quad d - \frac{1}{2} + \frac{1 - \sqrt{1 - 4d/\theta_m}}{2} + (1 - d)$$

$$= \quad \frac{1 + \sqrt{1 - 4d/\theta_m}}{2}. \qquad (4.10)$$

This equilibrium exists as long as $4d/\theta_m < 1$, i.e. $\theta_m > 4d$.
Substituting $x_a = (1 + \sqrt{1 - 4d/\theta_m})/2$ in $p = \frac{1}{\alpha x_a^2}$ we find

$$p = \frac{4}{\alpha(1 + \sqrt{1 - 4d/\theta_m})^2}. \qquad (4.11)$$

Now that we got all the information on the equilibrium, we refocus on the system performance. Because only some leavers free-ride, in the equilibrium everybody gets full benefit and thus:

$$W_{sys} \quad = \quad \alpha x_a - 1$$

$$W_{sys} \quad = \quad \alpha\Big(\frac{1 + \sqrt{1 - 4d/\theta_m}}{2}\Big) - 1. \qquad (4.12)$$

The system is optimal if everyone contributes, so

$$POS = \frac{\alpha - 1}{\alpha\Big(\frac{1+\sqrt{1-4d/\theta_m}}{2}\Big) - 1}. \qquad (4.13)$$

Example (continued)

We continue with the parameter $\theta_m = 6$ and $\alpha = 4$. Let $d = 0.7$.
Calculating p gives:

$$p = \frac{4}{4(1 + \sqrt{1 - 4 \cdot 0.7/6})^2} \approx 0.334.$$

This gives a contribution level of $x_a \approx 0.865$. The system performance now becomes

$$W_{sys} = 4\Big(\frac{1 + \sqrt{1 - 4 \cdot 0.7/6}}{2}\Big) - 1 \approx 2.461.$$

The system is optimal for $x_a = 1$, so

$$POS = \frac{4 - 1}{W_{sys}} = 1.219.$$

4.3.2 Free Identities

Under FI, identities are free. Players who want to free-ride can just do so and don't worry about the consequences: they just get a new identity and start all over. To combat this whitewashing, we will have to reduce service, not only for the existing free-riders, but for all newcomers. This results in a social loss (all

newcomers get reduced benefits, even when contributing).
The burden for the contributors in this system is lower than that under PI. Under FI, we always penalize newcomers, so the demand on the system is lower. Only the existing contributors get full service; all newcomers and existing free-riders get reduced service. This means that the burden R_{FI} becomes

$$R_{FI} = \frac{(1-d)x_s + d(1-p) + (1-d)(1-x_s)(1-p)}{x_a}.$$

Here $(1-d)x_s$ are the existing contributors, the only group to receive full service in this model. d is the fraction of newcomers, who all get reduced service and $(1-d)(1-x_s)$ are the existing free-riders, who also receive reduced services.
Under PI, if the penalty p is set sufficiently high, it's possible to get a scenario where the penalty is only used as a threat, without actually reducing service for anyone, as I showed in the previous section. In FI, imposing a penalty always results in a social loss because newcomers, whether contributing or not, are always penalized.

We find the Nash equilibria the same way as we did for the PI-system. Because the burden is lower than in PI, we know that is is again possible to find a p such that all stayers will contribute ($x_s = 1$). Because all stayers will contribute to the system, only some peers leaving the system will free-ride. This means that all peers except newcomers receive full service. Using that, we can calculate the contribution level of the leavers:

$$\begin{aligned}
x_l = P(\theta_i \geq R_{FI}) &= P(\theta_i \geq \frac{(1-d)+d(1-p)}{x_a}) \\
&= 1 - \frac{1-dp}{x_a\theta_m} \\
&= 1 - \frac{1-dp}{(dx_l + (1-d))\theta_m} \\
(1-x_l)\theta_m &= \frac{1-dp}{1-d(1-x_l)} \\
(1-x_l)(1-d(1-x_l))\theta_m &= 1-dp \\
0 &= d\theta_m(1-x_l)^2 - \theta_m(1-x_l) + (1-dp).
\end{aligned}$$

This results in

$$\begin{aligned}
1 - x_l &= \frac{\theta_m \pm \sqrt{\theta_m^2 - 4d\theta_m(1-dp)}}{2d\theta_m} \\
x_l &= 1 - \frac{1-\sqrt{1-4d(1-dp)/\theta_m}}{2d}.
\end{aligned}$$

We now find that the average contribution level in the best equilibrium in the system is:

$$x_a = (1-d)x_s + dx_l$$

$$
\begin{aligned}
&= (1-d) + d\Big(1 - \frac{1-\sqrt{1-4d(1-dp)/\theta_m}}{2d}\Big) \\
&= 1 - d + d - \frac{1}{2} + \frac{1-\sqrt{1-4d(1-dp)/\theta_m}}{2d} \\
&= \frac{1+\sqrt{1-4d(1-dp)/\theta_m}}{2}. \qquad (4.14)
\end{aligned}
$$

This equilibrium exist as long as $4d(1-dp)/\theta_m \leq 1$, i.e. $\theta_m \geq 4d(1-dp)$. Again, we find the penalty level p that is required to attain this equilibrium, by substituting $x_s = 1$ in

$$
\begin{aligned}
x_s &= 1 - \Big(\frac{1-dp}{x_a} - p\alpha x_a\Big)\frac{1}{\theta_m} \\
1 &= 1 - \Big(\frac{1-dp}{x_a} - p\alpha x_a\Big)\frac{1}{\theta_m} \\
0 &= \frac{1-dp}{x_a} - p\alpha x_a \\
0 &= p\alpha x_a^2 - (1-dp) \\
0 &= p\alpha\Big(\frac{1+\sqrt{1-4d(1-dp)/\theta_m}}{2}\Big)^2 - (1-dp). \qquad (4.15)
\end{aligned}
$$

The desired $p = p_{eq}$ is the root of the equation above.
Next, we find the system performance in equilibrium.
In this system, only the existing contributors $(1-d)x_s$ receive full service $Q = \alpha x_a$. The other peers (the existing free-riders and all newcomers) receive reduced services $(1-p)Q = (1-p)\alpha x_a$. This group has the size $(1-d)(1-x_s)+d$. All contributors x_a contribute the burden $R_{FI} = \frac{1-p(1-(1-d)x_s)}{x_a}$.
This gives the following system performance:

$$
\begin{aligned}
W_{sys} &= (1-d)x_s Q + ((1-d)(1-x_s)+d)(1-p)Q - x_a R_{FI} \\
&= Q - p(1-(1-d)x_s)Q - x_a R_{FI} \\
&= \alpha x_a - p\alpha x_a(1-(1-d)x_s) - (1-p(1-(1-d)x_s)) \\
&= (\alpha x_a - 1)(1-p(1-(1-d)x_s)). \qquad (4.16)
\end{aligned}
$$

For the equilibrium we found, everybody except the newcomers received full service, so

$$
\begin{aligned}
W_{sys} &= \alpha x_a(1-d) + (1-p)\alpha x_a d - (1-dp) \\
&= \alpha x_a(1-dp) - (1-dp) = (\alpha x_a - 1)(1-dp). \qquad (4.17)
\end{aligned}
$$

The optimal system performance is attained when everyone contributes, so,

$$W_{sys}(x_a = 1) = (\alpha - 1)(1-dp).$$

Now we can calculate the Price of Stability:

$$POS = \frac{(\alpha-1)(1-dp)}{(\alpha x_a - 1)(1-dp)}$$

$$= \quad \frac{\alpha - 1}{\alpha x_a - 1}. \qquad (4.18)$$

Example (continued)

For the parameters $\theta_m = 6$, $\alpha = 4$ and $d = 0.7$, we find p_{eq} by solving

$$0 = 4p\Big(\frac{1 + \sqrt{1 - 4 \cdot 0.7(1 - 0.7p)/6}}{2}\Big)^2 - (1 - 0.7p).$$

This gives $p_{eq} \approx 0.257$. The penalty here is lower than in the PI system. The reason for this is that in the FI system, all newcomers are penalized as opposed to the PI system. This leads to a lower overall burden and therefore the players need less incentive to contribute.
Now, using p_{eq}, we find the average contribution level x_a:

$$x_a = \frac{1 + \sqrt{1 - 4 \cdot 0.7(1 - 0.7p_{eq})/6}}{2} \approx 0.893.$$

Subsequently, the system performance in the equilibrium becomes:

$$W_{sys} = (4x_a - 1)(1 - 0.7p_{eq}) \approx 2.109.$$

and the Price of Stability:

$$POS = \frac{4 - 1}{4x_a - 1} \approx 1.166.$$

We see that the Price of Stability here is slightly better than for the permanent identities (the POS for PI is 1.219). However, the actual system performance has actually dropped significantly from $W_{PI} = 2.461$ tot $W_{FI} = 2.109$.

4.4 Conclusion

In this chapter we tried to model human behaviour. While theoretically peers should want their costs minimized at all times, in practice we see that people don't mind a little cost (in the form of uploading data), as long as it doesn't bother them. We tried to capture this willingness by assuming generosity levels for all players, the maximum costs they are willing to selflessly incur.
In the closed system I showed that adding an incentive mechanism results in a clear improvement of the system. For the right penalty level ($p = 1/\alpha$), all players are forced to contribute to the system, creating a stable equilibrium at $x = 1$ that optimizes the system performance, resulting in a *Price of Stability* of 1. Even for a lower penalty level ($p < 1/\alpha$), the stable equilibrium is a clear improvement over the basic system, in terms of contribution level, system performance and consequently POS.
After that we have examined the dynamical system, where we assumed a constant turnover. We distinguished between two systems. One with permanent

identities, where newcomers are given full service and the other with free identities, where all newcomers are penalized.
In the system with permanent identities, we again were able to manipulate the system so that we reach an equilibrium $x_a = \frac{1+\sqrt{1-4d/\theta_m}}{2}$ where we only use the penalty p as a threat. Everyone received full service and only a small fraction of the 'leavers' did not contribute (the threat does not affect them).
When assuming free identities, we addressed whitewashing by penalizing all newcomers, so there was no difference between whitewashing and free-riding. While we only used p as a threat in the PI system, here we were forced to penalize a percentage of users (all newcomers), resulting in a social loss.
However, because all newcomers are penalized, the burden is lower and in order to ensure all stayers to contribute, a smaller penalty than in PI is required. We found that the acquired equilibrium $x_a = \frac{1+\sqrt{1-4d(1-dp)/\theta_m}}{2}$ gives a lower system performance than the PI system, but a higher contribution level such that the price of stability is actually smaller than for the PI system.

Chapter 5

Request Forwarding in P2P Systems

Strategic behavior in P2P systems goes far beyond free-riding in file-sharing networks. Peers may make strategic decisions on the timing of their arrivals and departures from the network, in selecting which peers to connect to, on what private information such as cost and valuations they want to report truthfully to the system, or engage in other ways of manipulating the system protocol or mechanism.

In this section we take a look at how peers may behave strategically when their actions are hidden and how currency-based incentive mechanisms could help the system to overcome these issues.

We once again look at P2P file-sharing. Besides sharing files, peers of networks such as KaZaA are also expected to forward protocol messages to and from their neighbours. When a peer receives a query from one of his neighbours he is expected to forward this to his other neighbours in addition to responding to the query if he is able to. However, the peer can strategically choose to ignore the message as to reduce his message forwarding costs. Those actions are usually not easy to notice and often it is not possible to quickly find such a defective node. This is because messages are forwarded on a 'best-effort' basis and the topology is continually changing as peers enter and leave the network [12]. Obviously, if everyone would decide to stop forwarding the messages, the system would collapse because nobody would find the files he is looking for. This leaves us with the question how the requesting peer can provide incentives for the others to forward his messages.

This problem of incomplete information or *hidden action* is hardly unique to networks, and has in fact long been studied by economists as the problem of *moral hazard* in the contexts ranging from insurances to labor contracts [12].

In the next section, we'll apply the principal-agent framework to analyze the efficiency loss due to hidden action, and the design of optimal contracts to induce effort by the agents, this model is due to Babaioff et al. (2006) [3].

In P2P models, the upcoming principal is the peer requesting a file and the agents represent the peers who can decide whether to forward the request.

5.1 The Principal-Agent Model

In the principal-agent problem, one or more agents act on behalf of the principal. The problem arises when the agents and the principal have different objectives. The two parties also have asymmetric information: the principal can't fully monitor the actions of the agents.
Here we assume a principal who employs a set of n agents, N. Each agent $i \in N$ can either exert low effort or high effort, denoted by the set of possible actions $A_i = \{0,1\}$. The agents' cost, related to his effort is denoted by $c(a_i) \geq 0$ for each possible action $a_i \in A_i$. Low effort incurs a cost of zero while the cost of high effort is $c > 0$, i.e., $c(0) = 0$ and $c(1) = c$.
The actions of the agents collectively and probabilistically determine a 'contractible' outcome, $o \in \{0,1\}$, where the outcome 0 denotes the failure and 1 the success of the project. The principal's valuation of a successful project is given by a scalar $v > 0$, while he gains nothing if the project fails. The outcome is determined according to the project *technology*, a success function $t : A_1 \times ... \times A_n \rightarrow [0,1]$, where $t(a_1, ..., a_n)$ denotes the probability of project success when agents adopt the action profile $a = (a_1, ..., a_n) \in A_1 \times ... \times A_n = A$.
We will examine a subclass of technologies that can be represented by *read-once networks*. Read-once networks are given by a graph with two special nodes, a *source* and a *sink*, where each agent i controls a single edge. If an agent exerts low effort, he succeeds with probability γ_i and if he exerts high effort, the success probability increases to $\delta_i > \gamma_i$. The projects succeeds if there is a successful path, from source to sink, where the technology maps the individual successes and failures of agents (denoted by $x_i = 1$ and $x_i = 0$ respectively) into the probability of project success.
Two most common examples are the *sequence technology* and the *parallel technology*. We'll discuss the case where the technology is anonymous, i.e. symmetric with respect to the agents. This means that $t(a_1, ...a_n)$ only depends on $\sum_i a_i$. Furthermore the technology is determined by a single parameter $\gamma \in (0, 1/2)$ that satisfies $1 - \delta_i = \gamma_i = \gamma$ for all i.

The sequence technology Here $f(x_1, .., x_n) = \prod_{i \in N} x_i$. So the projects is a success if and only if all individual agents are successful with their tasks. So if m agents exert high effort, then $t(a) = \gamma^{n-m}(1-\gamma)^m$.
This is the situation where a message is sent over a graph consisting of sequential arrows from source to sink and everyone has to forward the message in order for it to arrive at its goal.

The parallel technology In this situation $f(x_1, .., x_n) = 1 - \prod_{i \in N}(1 - x_i)$. This means that the project is a success if and only if at least one of the agents is successful in his task. In this case, with m agents exerting effort,

$t(a) = 1 - \gamma^m(1-\gamma)^{n-m}$.
This is represented by the graph with n parallel arrows from source to sink. Here we can think of a situation where a message is sent through a number of different channels. As long as one of these delivers, the message will reach its goal.

In this model, the principal may design contracts based on the observable outcome. We don't allow for fines, so negative payments to the agents are not allowed, they are either rewarded or receive nothing. The contract is a commitment to pay agent i an amount $p_i \geq 0$ upon project success regardless of the success of the individual agent. If the project fails the agents receive nothing. Given this, the agents are in a game, where the utility of agent i with the profile of actions $a = (a_1, ..., a_n)$ is $u_i(a) = p_i \cdot t(a) - c(a_i)$.
The principal's challenge is to design a contract such as to maximize his own expected utility $u(a, v) = t(a) \cdot (v - \sum_{i \in N} p_i)$, where the actions $a_1, ..., a_n$ are at Nash equilibrium, so that the agents will honor the contract. If there are multiple equilibria, the principal can choose a desired one and 'suggest' it to the agents. Here we also assume that agents can not coordinate.
The goal is to motivate the agents, therefore we assume that more effort by an agent always results in a higher probability of success. Let a_{-i} be the action profile of all peers except i, from the set of all possible actions A_{-i}, then:

$$\forall i \in N, \forall a_{-i} \in A_{-i}, t(1, a_{-i}) > t(0, a_{-i})$$

In addition we assume that $t(a) > 0$ for any $a \in A$.

Definition 5.1. *The marginal contribution of agent i, given $a_{-i} \in A_{-i}$ is*

$$\Delta_i(a_{-i}) = t(1, a_{-i}) - t(0, a_{-i}) \in (0, 1).$$

The marginal contribution is the increase in the probability of success of the project due to the agents increase in effort. The best strategy of agent i can now be determined by $a_{-i} \in A_{-i}$, and his contract p_i.

Lemma 5.1. *Given a profile of actions a_{-i}, agent i's best strategy is $a_i = 1$ if $p_i \geq \frac{c}{\Delta_i(a_{-i})}$ and $a_i = 0$ if $p_i \leq \frac{c}{\Delta_i(a_{-i})}$.*

Proof. Given a_{-i}, the expected utility of player i, for $a_i = 1$ is

$$u_i(1) = p_i t(1, a_{-i}) - c.$$

For $a_i = 0$ this is

$$u_i(0) = p_i t(0, a_{-1}).$$

For $p_i = \frac{c}{\Delta_i(a_{-i})}$:

$$\begin{aligned} u_i(1) &= p_i t(1, a_{-i}) - c \\ &= \frac{c}{\Delta_i(a_{-i})} t(1, a_{-i}) - c \end{aligned}$$

$$
\begin{aligned}
&= \frac{c}{\Delta_i(a_{-i})}\Big(\Delta_i(a_{-i}) + t(0, a_{-i})\Big) - c \\
&= c + \frac{c \cdot t(0, a_{-i})}{\Delta_i(a_{-i})} - c \\
&= \frac{c}{\Delta_i(a_{-i})} t(0, a_{-i}) \\
&= u_i(0).
\end{aligned}
$$

Because $t(1, a_{-i}) > t(0, a_{-i})$, for $p_i > \frac{c}{\Delta_i(a_{-i})}$, $u_i(1) > u_i(0)$. So the best strategy for player i is $a_i = 1$.
For $p_i < \frac{c}{\Delta_i(a_{-i})}$, $u_i(1) < u_i(0)$. So the best strategy is $a_i = 0$. For $p_i = \frac{c}{\Delta_i(a_{-i})}$ player i is indifferent between strategies $a_i = 1$ and $a_i = 0$. □

Now we can specify the principal's optimal contracts for inducing a given equilibrium.

Lemma 5.2. *The best contracts for the principal that induce $a \in A$ as an equilibrium are $p_i = 0$ for agent i who exerts no effort ($a_i = 0$) and $p_i = \frac{c}{\Delta_i(a_{-i})}$ for agent i who exerts effort ($a_i = 1$).*
In this case, the expected utility of agent i who exerts high effort is $c \cdot (\frac{t(1,a_{-i})}{\Delta_i(a_{-i})} - 1)$, and 0 for an agent who exerts low effort. The principal's expected utility is given by $u(a, v) = (v - \sum_{i|a_i=1} \frac{c}{\Delta_i(a_{-i})}) \cdot t(a)$.

Proof. As I've shown in the proof of Lemma 5.1, $p_i = \frac{c_i}{\Delta_i(a_{-i})}$ is the lowest payment for which agent i doesn't prefer $a_i = 0$ over $a_i = 1$. Not paying agents who only exert low effort is clearly optimal. These payments are therefore the best contracts the principal can draft. □

If, in the induced equilibrium, agent i exerts high effort, we say that the principal contracts with this agent. Because the payment to each agent is higher than the agents cost, the utility of the principal is lower than in the observable-actions case, where it is enough to pay the agents cost in order to let him exert high effort.
The goal of the principal is to determine the profile of actions $a^* \in A$ which gives him the highest utility $u(a, v)$ in equilibrium, given his valuation v. This means choosing a set S of agents that exert high effort ($S = i|a_i = 1$). The set of agents $S^*(v)$ that the principal contracts with in a^* is an *optimal contract* for the principal at value v.
To measure these choices, we can compare the decisions with the observable-actions case. There, the principal can induce effort with a payment $p_i = c$ to agent i, resulting in the principal's utility being

$$u(a, v) = t(a) \cdot v - \sum_{i|a_i=1} c,$$

where all the agents work for cost.
We will call the worst case ratio between the optimal utility in the observable-actions case and in the hidden-actions case the *Price of Unaccountability* [12].

Let S^*_{oa} denote an optimal contract in the observable-actions case, when the principal's valuation is v.

Definition 5.2. *The* ***Price of Uunaccountability*** *$POU(t)$ of a technology t is defined as the worst ratio (over v) between the principal's utility in the observable-actions case and in the hidden actions case:*

$$POU(t) = \sup_{v>0} \frac{t(S^*_{oa}(v)) \cdot v - \sum_{i \in S^*_{oa}(v)} c}{t(S^*(v))(v - \sum_{i \in S^*(v)} \frac{c}{t(S^*(v)) - t(S^*(v) \setminus \{i\})})}. \tag{5.1}$$

5.2 Results

We want to be able to construct a optimal set of contracted agents, as a function of the principal‘s valuation of project success.

Lemma 5.3 (Monotonicity lemma)**.** *For any technology t, in both the hidden-actions and the observable-actions cases, the expected utility of the principal at the optimal contracts, the success probability of the optimal contracts, and the expected payment of the optimal contract, are all monotonically nondecreasing with the valuation v.*

The proof of this lemma can be found in [3].
For the anonymous sequence and parallel technologies, this implies that the number of contracted agents is a non-decreasing function of the valuation.
We now take another look at these two specific technologies.
First for the sequence technology:

Theorem 5.1. *Assume an anonymous sequence technology with n agents and with $\gamma = \gamma_i = 1 - \delta_i \in (0, \frac{1}{2})$ for all i.*
There exists a valuation $v_* < \infty$ such that for any $v < v_*$ it is optimal to contract with zero agents, for $v > v_*$ it is optimal to contract with all n agents, and for $v = v_*$ either contracting none or all agents is optimal.
The Price of Unaccountability is obtained at v^*, such that $POU = (\frac{1}{\gamma} - 1)^{n-1} + (1 - \frac{\gamma}{1-\gamma})$.

Proof. For any fixed number of contracted agents, k, the principal's utility is a linear function in v, where the slope equals the success probability under k contracted agents. Thus the optimal contract corresponds to the maximum over a set of linear functions. Let v_* denote the point at which the principal is indifferent between contracting 0 and n agents and let $v_{0,k}$ be the value where the principal is indifferent between contracting 0 agents and k agents.
For $v = 0$ contracting 0 agents is obviously optimal. The first transition point (where the principal changes the number of agents he contracts with) is reached at the lowest v for which $u(0, v) = u(a, v)$ with $\sum_i a_i = k$ for some $k = 1, .., n$.
For the hidden action case of the sequence technology, this is the equation

$$\gamma^n \cdot v = \gamma^{n-k}(1-\gamma)^k \cdot \Big(v - \frac{ck}{\gamma^{n-k}(1-\gamma)^k - \gamma^{n-k+1}(1-\gamma)^{k-1}}\Big). \tag{5.2}$$

Calculating we find:

$$\begin{aligned}
\gamma^n \cdot v &= \gamma^{n-k}(1-\gamma)^k \cdot v - \frac{ck \cdot \gamma^{n-k}(1-\gamma)^k}{\gamma^{n-k}(1-\gamma)^{k-1}(1-2\gamma)} \\
&= \gamma^{n-k}(1-\gamma)^k \cdot v - ck\frac{1-\gamma}{1-2\gamma} \\
\gamma^{n-k}((1-\gamma)^k - \gamma^k)v &= ck\frac{1-\gamma}{1-2\gamma} \\
v_{0,k} &= \frac{1-\gamma}{1-2\gamma} \cdot \frac{ck}{\gamma^{n-k}((1-\gamma)^k - \gamma^k)}.
\end{aligned}$$

Here $\frac{1-\gamma}{1-2\gamma} > 0$ for $\gamma \in (0, \frac{1}{2})$.
Taking the derivative over k, I'll show that $v_{0,k}$ is decreasing in $k \geq 1$.

$$\begin{aligned}
\frac{\partial v_{0,k}}{\partial k} < 0 \quad &\Leftrightarrow \quad \frac{\partial}{\partial k}\Big(\frac{k}{(1-\gamma)^k\gamma^{-k} - 1}\Big) < 0 \\
&\Leftrightarrow \quad \frac{((1-\gamma)^k\gamma^{-k} - 1) - k(1-\gamma)^k\gamma^{-k}(\ln(1-\gamma) - \ln(\gamma))}{((1-\gamma)^k\gamma^{-k} - 1)^2} < 0 \\
&\Leftrightarrow \quad (1-\gamma)^k\gamma^{-k}(1 - k\ln(\frac{1-\gamma}{\gamma})) - 1 < 0 \\
&\Leftrightarrow \quad 1 + k\ln(\frac{\gamma}{1-\gamma}) < \Big(\frac{\gamma}{1-\gamma}\Big)^k.
\end{aligned}$$

substituting $\frac{\gamma}{1-\gamma} = x$, we get the inequality

$$x^k > 1 + k \cdot \ln(x).$$

This inequality holds for all $x \neq 1$. Which means the original inequality holds for all $\gamma \neq 1/2$. $\gamma \in (0, 1/2)$, so $v_{0,k}$ is always decreasing in $k \geq 1$ and $v_{0,k}$ is smallest for $k = n$. This means the first transition point is at v_*, such that $u(0, v_*) = u(1, v_*)$, where $a_i = 1$ for all $i \in N$. Because the slope of the utility function is increasing in k, this is the only transition point.
As the number of contracted agents is monotonically non-decreasing in the value, for any $v < v_*$, contracting 0 agents is optimal and for $v > v_*$, contracting n agents is optimal. The *POU* is obtained at a transition point [3]. I'll denote t_k as the probability of project success with k agents exerting high effort. We now find

$$POU = \frac{v_* \cdot t_n - c \cdot n}{v_* \cdot t_0}.$$

v_* is the value for which the utility with zero contracted agents is equal to the utility with n contracted agents in the hidden-actions case. So

$$\begin{aligned}
t_0 \cdot v_* &= t_n \cdot \Big(v_* - \frac{c \cdot n}{t_n - t_{n-1}}\Big) \\
v_* &= \frac{c \cdot n}{t_n - t_o} \cdot \frac{t_n}{t_n - t_{n-1}}.
\end{aligned}$$

Also we know that

$$t_i = \gamma^{n-i}(1-\gamma)^i.$$

Now replacing v_*, t_0 and t_n in the POU equation above gives

$$POU = (\frac{1}{\gamma} - 1)^{n-1} + (1 - \frac{\gamma}{1-\gamma}).$$

□

Notice that the POU is not bounded for various n, γ, as $POU \to \infty$ for either $\gamma \to 0$ or $n \to \infty$ (for any fixed $\gamma \in (0, \frac{1}{2})$).

For the parallel technology.

Theorem 5.2. *Assume an anonymous parallel technology with n agents and with $\gamma = \gamma_i = 1 - \delta_i \in (0, \frac{1}{2})$ for all i.*
There exist finite positive values $v_1 < v_2 < ... < v_n$ such that any v where $v_k < v < v_{k+1}$, contracting with exactly k agents is optimal. For $v < v_1$, no agent is contracted, for $v > v_n$, all n agents are contracted, and for $v = v_k$ the principal is indifferent between contracting with $k-1$ or k agents.
For any parallel technology with any n, c and $\gamma \in (0, \frac{1}{2})$, $POU \leq 5/2$

Proof. Define v_k as the value where the principal is indifferent between contracting with k or $k-1$ agents. As the number of contracted agents is monotonic non-decreasing in the value, due to the Monotonicity Lemma, it is sufficient to show that $v_1 < v_2 < ... < v_n$.
For that we need to find v_k. This is the value for which the utilities with k and $k+1$ contracted agents are equal.
For the hidden action case of the parallel technology, this is the equation

$$\begin{aligned} t_k \cdot \left(v - \frac{ck}{t_k - t_{k-1}}\right) &= t_{k+1} \cdot \left(v - \frac{c(k+1)}{t_{k+1} - t_k}\right) \\ (t_k - t_{k+1}) \cdot v &= \frac{t_k \cdot ck}{t_k - t_{k-1}} - \frac{t_{k+1} \cdot c(k+1)}{t_{k+1} - t_k}, \end{aligned}$$

with $t_k = (1 - \gamma^k(1-\gamma)^{n-k})$ this gives

$$\begin{aligned} -\gamma^k(1-\gamma)^{n-k-1}(1-2\gamma) \cdot v &= \frac{(1-\gamma^k(1-\gamma)^{n-k})ck}{\gamma^{k-1}(1-\gamma)^{n-k}(1-2\gamma)} \\ &\quad - \frac{(1-\gamma^{k+1}(1-\gamma)^{n-k-1})c(k+1)}{\gamma^k(1-\gamma)^{n-k-1}(1-2\gamma)} \\ &= \frac{\gamma(1-\gamma^k(1-\gamma)^{n-k})ck}{\gamma^k(1-\gamma)^{n-k}(1-2\gamma)} \\ &\quad - \frac{(1-\gamma)(1-\gamma^{k+1}(1-\gamma)^{n-k-1})c(k+1)}{\gamma^k(1-\gamma)^{n-k}(1-2\gamma)} \\ &= \frac{(2\gamma-1)ck - (1-\gamma)c + \gamma^{k+1}(1-\gamma)^{n-k}c}{\gamma^k(1-\gamma)^{n-k}(1-2\gamma)}. \end{aligned}$$

This gives

$$\begin{aligned} v_k &= -c \cdot \frac{(2\gamma-1)k-(1-\gamma)+\gamma^{k+1}(1-\gamma)^{n-k}}{\gamma^{2k}(1-\gamma)^{2n-2k-1}(1-2\gamma)} \\ &= \frac{c}{(1-2\gamma)(1-\gamma)^{2n-1}} \cdot \frac{(1-2\gamma)k+(1-\gamma)-\gamma^{k+1}(1-\gamma)^{n-k}}{\gamma^{2k}(1-\gamma)^{-2k}}. \end{aligned}$$

Now

$$\frac{\partial v_k}{\partial k} > 0 \Leftrightarrow \frac{\partial}{\partial k}\Big(\frac{(1-2\gamma)k+(1-\gamma)-\gamma^{k+1}(1-\gamma)^{n-k}}{\gamma^{2k}(1-\gamma)^{-2k}}\Big) > 0. \qquad (5.3)$$

We show that this is the case.

$$\begin{aligned} \frac{\partial}{\partial k}\Big(...\Big) &= \frac{(1-2\gamma)-\gamma^{k+1}(1-\gamma)^{n-k}\ln(\frac{\gamma}{1-\gamma})}{\gamma^{2k}(1-\gamma)^{-2k}} \\ &\quad -\frac{2((1-2\gamma)k+(1-\gamma)-\gamma^{k+1}(1-\gamma)^{n-k})\ln(\frac{\gamma}{1-\gamma})}{\gamma^{2k}(1-\gamma)^{-2k}} \\ &= \frac{(1-2\gamma)+(2(1-2\gamma)k+2(1-\gamma)-\gamma^{k+1}(1-\gamma)^{n-k})\ln(\frac{1-\gamma}{\gamma})}{\gamma^{2k}(1-\gamma)^{-2k}} \end{aligned}$$

Because $\gamma \in (0, \frac{1}{2})$, $(1-2\gamma) > 0$, $2(1-\gamma) > 1$, $\gamma^{k+1}(1-\gamma)^{n-k} < 1$ and $\ln(\frac{1-\gamma}{\gamma}) > 0$, so $\frac{\partial v_k}{\partial k} > 0$ for $k \geq 1$ and we find that v_k is increasing in k. So $v_1 < v_2 < ... < v_n$. □

The proof for the upper bound of the POU for the parallel technology can be found in [4]. Finally we'll disprove the claim that in general the number of transitions in the hidden actions case is always equal to the number of transitions in the observable-actions case.

Lemma 5.4. *The number of transitions in the hidden action case is not always equal to the number of transitions in the observable actions case.*

Proof. We show this with an example.
We consider an anonymous technology with two agents, such that $c = 1, t_0 = 0, t_1 = 0.2$ and $t_2 = 0.5$.
$p_1 = 1/(t_1 - t_0) = 1/0.2$ if only player 1 is contracted. $p_1 = p_2 = 1/(t_2 - t_1) = 1/0.3$ if both players are contracted.
For the hidden-actions case we now have

$$u((0,0), v) = t_0 \cdot (v - 0) = 0$$

$$u((1,0), v) = u((0,1), v) = t_1 \cdot (v - p_1) = 0.2(v - \frac{1}{0.2}) = 0.2 \cdot v - 1$$

$$u((1,1), v) = t_2 \cdot (v - (p_1 + p_2)) = 0.5 \cdot (v - \frac{1}{0.3} \cdot 2) = 0.5 \cdot v - \frac{1}{0.3}$$

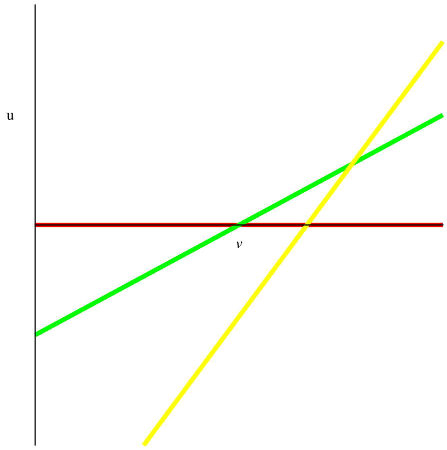

Figure 5.1: Hidden actions

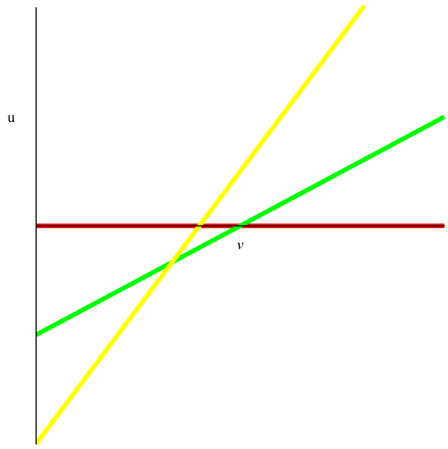

Figure 5.2: Observable actions

To determine the transition points we solve the linear equations $u((0,0),v) = u((0,1),v)$, $u((0,0),v) = u((1,1),v)$ and $u((0,1),v) = u((1,1),v)$. This gives:

$$\begin{array}{ccc} u((0,0),v) = u((0,1),v), & u((0,0),v) = u((1,1),v), & u((0,1),v) = u((1,1),v) \\ 0 = 0.2v - 1 & 0 = 0.5v - \frac{1}{0.3} & 0.2v - 1 = 0.5v - \dfrac{1}{0.3} \\ v = 5, & v = 6\frac{2}{3}, & v = 7\dfrac{7}{9}. \end{array}$$

We find two transition points: $v_1 = 5$ and $v_2 = 7\frac{7}{9}$ For an optimal contract, for $0 \leq v \leq 5$ the principal will contract no agents, for $5 \leq v \leq 7\frac{7}{9}$ he will contract 1 agent and for $v \geq 7\frac{7}{9}$ 2 agents will be contracted.

Now for the observable-actions case, we have

$$\begin{aligned} u((0,0),v) &= t_0 \cdot v = 0 \\ u((1,0),v) = u((0,1),v) &= t_1 \cdot v - c = 0.2v - 1 \\ u((1,1),v) &= t_2 \cdot v - 2c = 0.5v - 2. \end{aligned}$$

Again we solve the linear equations $u((0,0),v) = u((0,1),v)$, $u((0,0),v) = u((1,1),v)$ and $u((0,1),v) = u((1,1),v)$:

$$\begin{array}{ccc} u((0,0),v) = u((0,1),v), & u((0,0),v) = u((1,1),v), & u((0,1),v) = u((1,1),v) \\ 0 = 0.2 \cdot v - 1 & 0 = 0.5 \cdot v - 2 & 0.2v - 1 = 0.5v - 2 \\ v = 5, & v = 4, & v = 3\dfrac{1}{3}. \end{array}$$

Here we find only one transition point, $v_1 = 4$. Hence it is optimal for the principal to contract 0 agents for $v \leq 4$ and 2 agents for $v \geq 4$. It is never optimal to contract just one agent. □

In figure 5.1 above we clearly see that the hidden action case has 2 transition points, while the observable actions case has just one transition point.

5.3 Conclusion

With the principal-agent model, we examined the problem of the hidden actions in message forwarding in P2P systems. By introducing a micro payment incentive scheme, we encouraged the agents (i.e. peers) to forward messages (requests). Without these payments no agent would be willing to exert any effort to forward these messages. Due to the unaccountability of the intermediate agents, the principal does have to overpay for them to participate. Because the principal can only see the outcome of the process (success or failure), he only pays the agents if the process succeeds. We found for the two most common anonymous 'technologies' (sequence and parallel) that this does lead to overpaying. Comparing the observable-actions case with the hidden-actions case, we found that for sequence technology, the ratio can be infinitely bad. For the parallel technology we also have to overpay, but the principal's utility here is at least 40% of that of the observable-actions case.

Appendix A

3rd and 4th order polynomial equations

In section 3.4, we encountered the 4th order polynomial equation

$$d_1^4 + 4d_1^3 + (6 + 2b_{12})d_1^2 + (4 + 4b_{12} - b_{12}^2 b_{21})d_1 + (1 + b_{12})^2 = 0.$$

To know under which conditions this equation will have solutions $d_1 > 0$, we'll have to look at the derivatives. Let

$$f(d_1) = d_1^4 + 4d_1^3 + (6 + 2b_{12})d_1^2 + (4 + 4b_{12} - b_{12}^2 b_{21})d_1 + (1 + b_{12})^2$$

then

$$\frac{df(d_1)}{dd_1} = 4d_1^3 + 12d_1^2 + 2(6 + 2b_{12})d_1 + (4 + 4b_{12} - b_{12}^2 b_{21})$$

and

$$\frac{d^2 f(d_1)}{dd_1^2} = 12d_1^2 + 24d_1 + 2(6 + 2b_{12}).$$

We see that for $b_{12} > 0$, $\frac{d^2 f(d_1)}{dd_1^2}$ is always positive. This means that $\frac{df(d_1)}{dd_1}$ is continuously increasing. Because $\frac{df(d_1)}{dd_1}$ is a cubic function this means that $\frac{df(d_1)}{dd_1} = 0$ has only one solution, which in turn means that $f(d_1)$ has only one minimum. Next we see that for $d_1 = 0$, $f(0) = (1 + b_{12})^2 > 0$. In order for $f(d_1) = 0$ to have solutions for $d_1 > 0$, this means that the only minimum of $f(d_1)$ has to be attained for $d_1 > 0$.
To ensure that, $\frac{df(0)}{dd_1}$ has to be negative. We can easily see that this is the case if $4 + 4b_{12} - b_{12}^2 b_{21} < 0$. This gives us our first constraint:

$$\begin{aligned} 4 + 4b_{12} - b_{12}^2 b_{21} &< 0 \\ b_{12}^2 b_{21} &> 4 + 4b_{12} \\ b_{21} &> \frac{4 + 4b_{12}}{b_{12}^2} \end{aligned}$$

Now that we've ensured that the minimum of $f(d_1)$ is reached for $d_1 > 0$, we need to find out when this minimum is smaller or equal to zero.
To do that, we'll first solve $\frac{df(d_1)}{dd_1} = 0$.
In general third order polynomial equations have the form $ax^3+bx^2+cx+d = 0$.
If this equation has a single solution, the algebraic solution, with

$$\begin{aligned} Q &= \sqrt{(2b^3 - 9abc)^2 - 4(b^2 - 3ac)^3} \\ C &= \sqrt[3]{\frac{1}{2}(Q + 2b^3 - 9abc + 27a^2 d)}, \end{aligned}$$

is

$$x = -\frac{b}{3a} - \frac{C}{3a} - \frac{b^2 - 3ac}{3aC}. \tag{A.1}$$

for $b^2 - 3ac \neq 0$.
$\frac{df(d_1)}{dd_1} = 4d_1^3 + 12d_1^2 + 2(6 + 2b_{12})d_1 + (4 - 4b_{12} - b_{12}^2 b_{21})$, so $a = 4$, $b = 12$, $c = 2(6 + 2b_{12})$ and $d = (4 + 4b_{12} - b_{12}^2 b_{21})$. Substituting this, we find

$$\begin{aligned} 2b^3 - 9abc + 27a^2 d &= 2 \cdot 12^3 - 9 \cdot 4 \cdot 12 \cdot 2(6 + 2b_{12}) + 27 \cdot 4^2 \cdot (4 + 4b_{12} - b_{12}^2 b_{21}) \\ &= -432 b_{12}^2 b_{21} = -3^3 \cdot 4^2 b_{12}^2 b_{21} \end{aligned}$$

and

$$\begin{aligned} 4(b^2 - 3ac)^3 &= 4(12^2 - 3 \cdot 4 \cdot 2(6 + 2b_{12}))^3 \\ &= 4(-48 b_{12})^3 = 4(-3 \cdot 4^2 b_{12})^3. \end{aligned}$$

This gives

$$\begin{aligned} Q &= \sqrt{(-3^3 \cdot 4^2 b_{12}^2 b_{21})^2 - 4(-3 \cdot 4^2 b_{12})^3} \\ &= \sqrt{3^6 \cdot 4^4 b_{12}^4 b_{21}^2 + 4^7 \cdot 3^3 b_{12}^3} \\ &= \sqrt{3^3 \cdot (-4)^4 (3^3 b_{12}^4 b_{21}^2 + 4^3 b_{12}^3)} \\ &= -3 \cdot 4^2 \sqrt{3^4 b_{12}^4 b_{21}^2 + 3 \cdot 4^3 b_{12}^3}. \end{aligned}$$

Consequently

$$\begin{aligned} C &= \left(\frac{1}{2}(-3 \cdot 4^2 \sqrt{3^4 b_{12}^4 b_{21}^2 + 3 \cdot 4^3 b_{12}^3} - 3^3 \cdot 4^2 b_{12}^2 b_{21})\right)^{\frac{1}{3}} \\ &= -2\left(3^3 b_{12}^2 b_{21} + 3\sqrt{3 b_{12}^2 (3^3 b_{12}^2 b_{21}^2 + 4^3 b_{12}^3)}\right)^{\frac{1}{3}}. \end{aligned}$$

We now get

$$\begin{aligned} d_1 &= -\frac{b}{3a} - \frac{C}{3a} - \frac{b^2 - 3ac}{3aC} \\ &= -1 + \frac{1}{6}\left(27 b_{12}^2 b_{21} + 3\sqrt{192 b_{12}^3 + 81 b_{12}^4 b_{21}^2}\right)^{\frac{1}{3}} \\ &\quad - \frac{2 b_{12}}{\left(27 b_{12}^2 b_{21} + 3\sqrt{192 b_{12}^3 + 81 b_{12}^4 b_{21}^2}\right)^{\frac{1}{3}}}. \end{aligned}$$

Bibliography

[1] E. Adar and B.A. Huberman, *Free Riding on Gnutella*, 2000.

[2] E. Anshelevich, A. Dasgupta, J. Kleinberg, É. Tardos, T. Wexler and Tim Roughgarden, *The Price of Stability for Network Design with Fair Cost Allocation*, 2004.

[3] M. Babaioff, M. Feldman and N. Nisan, *Combinatorial Agency*, ACM Conference on Electronic Commerce, 2006.

[4] M. Babaioff, M. Feldman and N. Nisan *Combinatorial Agency*, Full Version, 2006.

[5] C. Buragohain, D. Agrawal and S. Suri, *A Game Theoretic Framework for Incentives in P2P Systems*, 2003.

[6] M. Feldman, K. Lai, I. Stoica and J. Chuang, *Robust Incentive Techniques for Peer-to-Peer Networks*, 2004.

[7] M. Feldman, C.H. Papadimitriou, J. Chuang and I. Stoica, *Free-riding and Whitewashing in Peer-to-Peer Systems* 3rd Annual Workshop on Economics and Information Security, 2004.

[8] E.J. Friedman, J.Y. Halpern, I. Kash, *Efficiency and Nash Equilibria in a Scrip System for P2P Networks*, 2006.

[9] E.J. Friedman and P. Resnick, *The Social Cost of Cheap Pseudonyms*, 1999.

[10] P. Golle, K. Leyton-Brown, I. Mironov and M. Lillibridge *Incentives for Sharing in Peer-to-Peer Networks*, 2001.

[11] R. Johari and J.N. Tsitsiklis, *Efficiency Loss in a Network Resource Allocation Game*, Mathematics of Operations Research, Vol 29, No. 3, 2004.

[12] N. Nisan, T. Roughgarden, E. Tardos and V.V. Vazirani, *Algoritmic Game Theory*, Camebridge Univerity Press, New York, 2007.

[13] C.H. Papadimitriou, *Algorithms, Games, and the Internet*, STOC 2001.

[14] M.W. Hirsch, S. Smale, R.L. Devaney, *Differential Equations, Dynamical Systems & An Introduction to Chaos*, Academic Press, Elsevier, 2004.

[15] A.C. Pigou, *The Economics of Welfare*, Macmillan, 1920.

[16] L. Ramaswamy and L. Liu, *Free Riding: A New Challenge to Peer-to-Peer File Sharing Systems*, 2003.

[17] T. Roughgarden and E. Tardos, *How Bad is Selfish Routing?*, J.ACM 2002.

[18] A.S. Schultz and N.S. Moses, *On The Performance of User Equilibria in Traffic Networks*, 2003.

[19] B. Yu and M.P. Singh, *Incentive Mechanisms for Peer-to-Peer Systems*, 2003.

[20] Jupiter Research, *Charting Piracy*, 2008.

[21] Solution Research Group, *Movie File-Sharing Booming: Study*, 2006.

[22] BitTorrent, *http://www.bittorrent.com/*.

[23] Encyclopaedia Britannica, *http://www.britannica.com/*.