MECHANISMS TO MANAGE INCENTIVES IN ONLINE SYSTEMS

A DISSERTATION
SUBMITTED TO THE DEPARTMENT OF
MANAGEMENT SCIENCE AND ENGINEERING
AND THE COMMITTEE ON GRADUATE STUDIES
OF STANFORD UNIVERSITY
IN PARTIAL FULFILLMENT OF THE REQUIREMENTS
FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY

Christina Aperjis
July 2009

I certify that I have read this dissertation and that, in my opinion, it is fully adequate in scope and quality as a dissertation for the degree of Doctor of Philosophy.

(Ramesh Johari) Principal Adviser

I certify that I have read this dissertation and that, in my opinion, it is fully adequate in scope and quality as a dissertation for the degree of Doctor of Philosophy.

(Benjamin Van Roy)

I certify that I have read this dissertation and that, in my opinion, it is fully adequate in scope and quality as a dissertation for the degree of Doctor of Philosophy.

(Nimrod Megiddo)

Approved for the University Committee on Graduate Studies.

Abstract

In large scale online systems, such as electronic marketplaces and peer-to-peer systems, users often act in their self interest without considering whether their actions lead to efficient outcomes for the system. This thesis studies two important classes of mechanisms that can be used to incentivize users to act in a way that promotes efficiency. *Aggregation mechanisms* provide aggregate information on the past behavior of a user to other users in the system. When there is such a mechanism in place, a user should expect that bad behavior now affects his future interactions within the system, and may be incentivized to act in a way that is beneficial for the system. *Market mechanisms* can be used to incentivize contribution to the system by using prices to identify value, and associating a budget with each user; the budget increases when the user contributes to the system and decreases when he uses system resources. By requiring that users have non-negative budgets, users can only use the system in return for valuable contributions.

In Chapter 2 we address a basic question: how do we design an aggregation mechanism to encourage trustworthy behavior? Electronic marketplaces, such as eBay, are a natural setting to study this question, since they usually rely on mechanisms that collect ratings of sellers from past transactions, and provide aggregate information to potential buyers. First, we show that weighting all past ratings equally gives sellers an incentive to falsely advertise. We then study aggregation mechanisms that weight recent ratings more heavily, and show that under increasing returns to reputation the optimal strategy of a sufficiently patient and sufficiently high quality seller is to always advertise honestly. We suggest approaches for designing an aggregation mechanism that maximizes the range of parameters for which it is optimal for the seller to be truthful. We show that mechanisms that use information from a larger number of past transactions tend to provide incentives for patient sellers to be

more truthful, but for higher quality sellers to be less truthful.

In Chapter 3 we study the use of market mechanisms for peer-assisted content distribution. We formulate a peer-to-peer filesharing system as an exchange economy: a price is associated with each file, and users exchange files only when they can afford it. The exchange economy approach allows peers to exchange files multilaterally. This formulation solves the free-riding problem, since uploading files is a necessary condition for being able to download. We discuss existence, uniqueness, and dynamic stability of the competitive equilibrium, which is always guaranteed to be Pareto efficient. In addition, a novel aspect of our approach is an allocation mechanism for clearing the market *out of equilibrium*. We analyze this mechanism when users can anticipate how their actions affect the allocation mechanism (price anticipating behavior). For this regime we characterize the Nash equilibria that will occur, and show that as the number of users increases, the Nash equilibrium rates become approximately Pareto efficient. Finally, we consider a system with a general network structure and show that maintaining a single price per peer (even across multiple files) suffices to achieve the benefits of multilateral exchange.

Most prevalent peer-to-peer systems incentivize users to contribute their upload capacity in a bilateral manner: downloading is possible in return for uploading to the same user. In Chapter 4 we provide a formal comparison of peer-to-peer system designs based on bilateral exchange with those that enable multilateral exchange via a price-based market mechanism to match supply and demand. First, we compare the two types of exchange in terms of the equilibria that arise. A competitive equilibrium allocation is Pareto efficient, while we demonstrate that bilateral equilibrium allocations are not Pareto efficient in general. We show that Pareto efficiency represents the "gap" between bilateral and competitive equilibria: a bilateral equilibrium allocation corresponds to a competitive equilibrium allocation if and only if it is Pareto efficient. Second, we compare the two types of exchange through the expected percentage of users that can trade in a large system, assuming a fixed file popularity distribution. Our theoretical results as well as analysis of a BitTorrent dataset provide quantitative insight into regimes where bilateral exchange may perform quite well even though it does not always give rise to Pareto efficient equilibrium allocations.

Acknowledgments

First of all I would like to express my sincere gratitude to my advisor, Professor Ramesh Johari. This work would not have been possible without his support, encouragement, and guidance. I would like to particularly thank Professor Benjamin Van Roy and Dr. Nimrod Megiddo, who served as thesis readers. I am further indebted to the other members of my oral examination committee, Professors Sunil Kumar and Nick Bambos. Special thanks are due to Professor Michael J. Freedman. My discussions with Professors Ashish Goel, Amin Saberi, and Richard Zeckhauser have also helped shape the work in this thesis.

Contents

List of Figures

Chapter 1

Introduction

An important feature of many online systems is that the value that a user derives from the system depends on the behavior of other users. For instance, in electronic marketplaces, such as eBay and the Amazon Marketplace, the value that buyers derive from the market depends on whether sellers commit fraud. In peer-to-peer file-sharing systems the value that a user derives depends on the collection of files and the upload rates that other users make available. In review sites, where users share their opinions on products, users derive more value when they find many high quality reviews.

We consider two types of user behavior that play a critical role in many prominent online systems. First, the performance of the system may depend on the *truthful revelation of private information* by users. This is the case in review sites, where users are expected to accurately describe their experience with a product, and in electronic marketplaces, where sellers are expected to truthfully describe the item they have for sale. Second, *contributions* by users are in many cases essential. The contributions may be in terms of bandwidth or files (peer-assisted content distribution), CPU cycles (grid computing), or time and knowledge (user generated content). In some settings it is important that contributions are valuable; e.g., in the case of peer-to-peer systems a peer should share files that other users desire.

Users, however, may behave selfishly and harm the system. On one hand, users may be better off *misrepresenting private information*. For instance, a seller in an electronic marketplace (such as eBay) may be tempted to exaggerate the value of the item he has for sale in order to increase his expected profit. On the other hand, it is generally costly for users to contribute resources (such as bandwidth and CPU cycles) or time. As a result, users often *free-ride*, i.e., use the resources and information that have been contributed by others without contributing themselves. To ensure that the system works well, it is important to align users' incentives with the overall system objective.

Moral incentives and legal action, which are widespread incentive mechanisms in the offline world, are often hard to use in online systems because of the inherent anonymity of participants. Moral incentives are usually not strong in online systems, because many interactions are between users that do not know each other and will probably never interact again. On the other hand, in many cases a user can easily change his identity in the system without significant effort [30], or legal consequences. Moreover, interactions often do not involve monetary transactions, and as a result legal action may not fit within the system's business plan.

Incentive mechanisms in online systems usually try to *incentivize a user by linking his behavior with the utility he gets from using the system*. With such a mechanism in place, the user needs to trade off the cost of behaving in a way that is beneficial for the system with the value he derives from using the system. As a result, a user that values interactions within the system may be incentivized to contribute or to reveal information truthfully, even though he would not do so in the absence of an incentive mechanism.

A user's behavior can be linked with the utility he gets from using the system in two ways. First, interactions with other users in the system can be affected if information on the user's past behavior is aggregated in a *score* which other users can see. The information that is aggregated in the user's score either consists of ratings that the user has received after past interactions with others in the system, or is a measure of his contribution to the system (e.g., in terms of upload rate or reviews). By convention, a high score indicates good behavior (e.g., contribution of resources or truthful information). In many settings a user with a high score is considered trustworthy, and as a result other users may prefer to interact with him. For instance, in electronic marketplaces buyers are willing to pay more when interacting with a seller that has a high score.

Second, *by design* the system may reward good behavior by allowing the user to enjoy certain privileges, and punish bad behavior by banning the user from the system. These devised rules of the system often depend on aggregate information on users' past behavior (i.e., users' scores). For instance, private peer-to-peer trackers reward peers that meet specific ratio requirements by giving them certain privileges such as granting permission to a new user to register at the site. In the BitTorrent protocol a peer splits its upload capacity between the peers from which he gets the highest download rates. On the other hand, electronic marketplaces often ban sellers with very bad history from participating in the market.

In order to design an effective mechanism, it is necessary to understand the effect of an incentive mechanism on the behavior of users. In particular, an important modeling decision concerns the interaction of a user with the incentive mechanism and other users in the system. A strategic user chooses a strategy that is a best response to the mechanism and the strategies of others in the system. If all users are strategic, then the outcome will correspond to a Nash equilibrium of the game between all users with respect to the incentive mechanism. However, this level of sophistication may be unrealistic in the context of a large scale online system, since it postulates very strong knowledge and rationality assumptions. A user may not know the strategies of others, especially if he has not interacted with them in the past as is usually the case in a large system. Users' behavior may often depend on simple heuristics, Bayesian techniques, or some combination of the two [61]. In general, different users may exhibit different levels of sophistication.

1.1 Related Literature

This thesis studies the design of mechanisms that incentivize users to provide valuable contributions and reveal private information truthfully. We study the former for peer-assisted content distribution, and the latter in the setting of electronic marketplaces. In this section we provide a brief overview of related work. A more detailed and focused literature review is contained in each chapter.

Mechanisms that associate a score with each user are usually called *reputation mechanisms* in the literature. The score aggregates information on past performance and represents the user's reputation within the system. As mentioned above, such a score may serve two purposes. First, it may affect the behavior of others in the system towards the user.

Second, there may be specific rules imposed or suggested by the system (e.g., through a protocol) on how much a user with a specific score can benefit from using the system.

In a market, buyers tend to trust reputed sellers more. As a result, reputed sellers enjoy a reputation *premium*, i.e., on average they sell at higher prices than sellers with lower scores. Shapiro studies the emergence of a premium for sellers of high quality products [64]. A number of empirical studies quantify the reputation premium in electronic marketplaces, i.e., consider the effect of a seller's score on the payment he receives [53, 47, 48, 37, 16, 31]. For example, a controlled experiment on vintage postcard auctions on eBay showed that an experienced seller enjoyed an 8% premium [61].

In electronic marketplaces, a reputation mechanism collects ratings from buyers and aggregates them in the seller's reputation score [59]. The two main design issues are (1) to elicit informative feedback [52, 14, 45], and (2) to aggregate ratings in a way that provides the right incentives to the seller and promotes trust (aggregation mechanism) [25, 22, 26]. We study the design of aggregation mechanisms in Chapter 2.

Reputation mechanisms have also been considered in the context of peer-assisted content distribution, where the goal is to incentivize contribution of upload capacity. The score of a user is usually a function of his contribution. Kazaa associates a participation level (low, medium, or high) with each peer depending on whether the peer is downloading more megabytes from other users than other users download from him. Other peer-to-peer systems such as eDonkey and eMule use similar metrics. A differential service-based incentive scheme is proposed in [34].

Market mechanisms also associate a score with each user: the user's budget. The budget (score) increases when the user contributes to the system and decreases when he uses system resources. Prices are associated with different types of contributions to identify their values. As a result, the budget of a user increases more for valuable contributions. The monetary incentives that have been considered in the literature [34, 39, 70] do not associate different prices with different types of contributions and therefore do not incentivize valuable contributions. We study market mechanisms for peer-assisted content distribution in Chapter 3.

BitTorrent's "tit for tat" is one of the most celebrated incentive mechanisms for peer-to-peer systems. This is an incentive mechanism with no naturally defined score for each user;

however, by design a user can increase his download rate by increasing his contribution. According to the BitTorrent protocol, each peer splits its available upload rate among peers from which it gets the highest download rates. As a result, the total download rate of a peer is a nondecreasing function of his total upload rate, and peers are incentivized to contribute [19]. We observe that the BitTorrent protocol incentivizes users in a bilateral basis: an increase in the upload rate to one peer may increase the download rate from that particular peer. On the other hand, with a system-wide score or a market mechanism in place, by increasing the upload rate to one peer, a user may increase his download rate from other peers in the system. This distinction is important in a system with multiple files and is a central theme of Chapter 4.

1.2 Contribution of this Thesis

1.2.1 Aggregation Mechanisms for Electronic Marketplaces

An aggregation mechanism that maps past ratings of the seller into reputation scores is a key component of a reputation system. In Chapter 2, we address a basic question: how should we aggregate information into the user's reputation score to best incentivize him? To address this question, it is important to understand the tradeoff that a user faces. We study this in the context of an electronic marketplace: a seller should expect that bad behavior now may increase the payment he receives in the current period, but will decrease future reputation scores, and thus future expected payments.

We model the seller as a long-lived player who believes that the expected payment he receives is a function of his reputation score and the description he posts. We assume that the expected payment is increasing in the seller's score, an assumption that is supported by empirical studies. We formulate the seller's decision problem, and study its properties using results from dynamic programming [11].

We take a non-equilibrium approach to study aggregation mechanisms. We model the buyer population by an exogenously defined *premium function* that captures the premium high-reputation sellers can command for goods declared to have high quality. Importantly, we do not assume that buyers play a best response to the seller's strategy. Our motivation comes from the fact that buyers may not know the strategy of the seller. Buyers often interact with sellers that they never interacted with before, so even if buyers can learn the strategy of a seller with repeated transactions, this effect will be limited.

Our goal is to incentivize the seller to always describe his items truthfully, a goal that is closely related to efficient trade. Our first result is that weighting all past ratings equally gives sellers an incentive to falsely advertise. This result supports eBay's decision in May 2008 to base the Positive Feedback percentage on the past 12 months of feedback, rather than the entire lifetime of the seller.[1]

We then study aggregation mechanisms that weight recent ratings more heavily. In particular, we assume that the seller's score s is a weighted sum of his past ratings, where the weights are given by some vector $\vec{w}$. We assume that the expected premium to the seller is $b_{\vec{w}}(s)$, i.e., the premium may explicitly depend on the weights of the aggregation mechanism. Special cases of this class of mechanisms include showing the mean of the seller's ratings in the last T transactions (the window mechanism), and exponential smoothing. We show that when recent ratings are weighted more, it is possible to incentive the seller to be always truthful. We identify conditions under which the optimal strategy of a sufficiently patient and sufficiently high quality seller is to always advertise honestly.

We suggest approaches for designing a weighted aggregation mechanism that maximizes the range of parameters for which it is optimal for the seller to be truthful. We show that *mechanisms that use information from a larger number of past transactions tend to provide incentives for patient sellers to be more truthful, but for higher quality sellers to be less truthful*. We show this tradeoff both for the general weighted aggregation mechanism where the decision is to choose the vector of weights $\vec{w}$ and for the window aggregation mechanism where the design parameter is the window size T.

We also show the tradeoff between incentivizing patient and high quality sellers for a broad range of settings. First, we consider the case of *perfect monitoring*, where buyers rate sellers accurately. Then, we consider two types of *imperfect monitoring*: the seller may not receive ratings after some transactions, and the rating that the seller receives may not accurately reflect his action. Our analysis employs monotone comparative statics [68, 51].

[1] http://pages.ebay.com/help/feedback/scores-reputation.html

1.2.2 Market Mechanisms for Peer-to-Peer Systems

Market mechanisms can be used to incentivize user contribution. These mechanisms associate a budget with each user; the budget increases when the user contributes to the system and decreases when he uses system resources. Prices are associated with different types of contributions to identify their values. As a result, the budget of a user increases more for valuable contributions. We note that market mechanisms can be interpreted as a subclass of reputation mechanisms. In particular, we can interpret the budget of a user as his reputation score: a user that has contributed a lot of valuable content to the system has a high reputation score, which he can use to download valuable content.

We study market mechanisms in the context of peer-assisted content distribution. In peer-to-peer systems, users share files or resources with each other. By sharing, a peer incurs a cost (because uploading a file consumes network resources), but no direct benefit. Thus, if there is no mechanism that stimulates sharing, a peer has a strong incentive to free ride, i.e., use the resources of other peers without contributing his own. Such behavior is observed in existing peer-to-peer systems; for instance, early data showed that nearly 70 percent of peers of Gnutella were sharing no files, and nearly 50 percent of all responses were returned by the top 1 percent of sharing hosts [1]. A more recent study shows that 85 percent of Gnutella peers share no files [38]. Even worse, according to [1], there were peers in Gnutella who were free riding on the system *despite* sharing files: the files that they were sharing were unpopular, and hence not widely uploaded.

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

In Chapter 3 we show how a peer-to-peer system can be formulated as an exchange economy. Even though monetary incentives have been previously proposed to incentivize uploading in peer-to-peer systems [34, 70, 66], these approaches neither study the efficiency benefits nor the optimal way to set prices. We show that the most suitable prices are the ones that align supply and demand when all peers have optimized. This is the concept of competitive equilibrium in economics [50]. We refer to this exchange as *multilateral*, because it allows peers to trade multilaterally. In Chapter 4 we compare a system that allows multilateral exchange with a system that restricts exchange to being bilateral.

We discuss existence, uniqueness, and dynamic stability of the competitive equilibrium, which is always guaranteed to be Pareto efficient. That is, it is not possible to reallocate rates in a way that makes some user better off without making someone else worse off. An important aspect of competitive equilibrium is that users are assumed to be price takers [50]. That is, users optimize with respect to given prices, and not with respect to the strategies of others. This is a reasonable assumption for a large peer-to-peer system, where we do not expect users to be able to optimize with respect to the strategies of others.

A novel aspect of our approach is an allocation mechanism for clearing the market *out of equilibrium*. If users exhibit price taking behavior, then prices will converge to equilibrium. Moreover, we analyze our allocation mechanism when users can anticipate how their actions affect the allocation mechanism (price anticipating behavior). For this regime we characterize the Nash equilibria that will occur, and show that *as the number of users increases, the Nash equilibrium rates become approximately Pareto efficient*. In other words, in a large system, even if users are strategic the allocation will be efficient. Finally, we consider a system with a general network structure and show that maintaining a single price per peer (even across multiple files) suffices to achieve the benefits of multilateral exchange.

1.2.3 Bilateral and Multilateral Exchange

Users of the BitTorrent file sharing protocol and its variants are incentivized to contribute their upload capacity in a bilateral manner: downloading is possible in return for uploading to the same user. A market mechanism on the other hand uses *multilateral exchange* to match user demand for content to available supply at other peers in the system. The fact that BitTorrent accounts for the majority of peer-to-peer traffic raises the following question:

how do bilateral and multilateral exchange compare? In Chapter 4, we provide a formal comparison of peer-to-peer system designs based on bilateral exchange with those that enable multilateral exchange via a price-based market mechanism to match supply and demand.

We start with a fundamental abstraction of content exchange in systems like BitTorrent: *exchange ratios*. The exchange ratio from one user to another gives the download rate received per unit upload rate. Exchange ratios are a useful formal tool because they allow us to define and study the equilibria of bilateral exchange. In a bilateral equilibrium each user optimizes with respect to exchange ratios. By contrast, in a *multilateral price-based exchange*, the system maintains one price per peer, and peers optimize with respect to these prices. For bilateral (resp., multilateral) exchange, an *equilibrium* is a combination of a rate allocation vector and an exchange ratio vector (resp., price vector) such that all peers have solved their corresponding optimization problems. In this case, the exchange ratios (resp., prices) have exactly aligned supply and demand. An equilibrium that corresponds to multilateral exchange is called a *competitive equilibrium*, and is studied in Chapter 3.

We compare bilateral and multilateral peer-to-peer systems through the allocations that arise at equilibria. A competitive equilibrium allocation is always Pareto efficient, while bilateral equilibria may be inefficient. Our main result is that *a bilateral equilibrium allocation is Pareto efficient if and only if it is a competitive equilibrium allocation*. This result provides formal justification of the efficiency benefits of competitive equilibria. The proof exploits an interesting connection between equilibria and Markov chains: an important step of the proof is to show that Pareto efficiency of a bilateral equilibrium rate allocation implies reversibility of an appropriately defined Markov chain, and that this chain has an invariant distribution that corresponds to a price vector of a competitive equilibrium.

We also perform a quantitative comparison of bilateral and multilateral exchange. As discussed in [40], "there may be many people wanting, and many possessing those things wanted; but to allow of an act of barter, there must be a double coincidence, which will rarely happen." We quantify how rarely this double coincidence of wants occurs under different assumptions on the popularity of files in the system. We first perform an asymptotic analysis assuming that file popularity follows a power law and study two extreme regimes. We find that *asymptotically all users are able to trade bilaterally when the file popularity is* *very concentrated.* On the other hand, multilateral exchange performs significantly better than bilateral when the file popularity is not concentrated. We complement our theoretical analysis by studying file popularity from a BitTorrent dataset. Although bilateral equilibria may in general be inefficient, the gap between bilateral and multilateral exchange can be narrowed significantly if each user shares a sufficient number of files (in practice as small as ten).

Chapter 2

Aggregation Mechanisms for Electronic Marketplaces

In online trading communities, sellers have a temptation to dishonesty, because potential buyers have to decide how much to pay for an item without being able to observe it firsthand. In particular, the buyer typically cannot know in advance whether the seller is describing the item honestly, and hence may be afraid that he might be exploited if he trusts the seller. This effect is exacerbated because a buyer will often interact with sellers with whom he has never interacted before and may only seldom interact in the future. The absence of trust created by this information asymmetry may result in market failure [2].

Aggregation mechanisms can be used to encourage buyers to trust sellers. Such mechanisms provide buyers with aggregate statistics on the past behavior of sellers; as a result, dishonest advertising by the seller involves a greater immediate payoff at the expense of a lower long-term payoff. Aggregation mechanisms typically operate as follows: after a transaction, the buyer rates the seller depending on how satisfied he was with the transaction. Then all ratings are aggregated into the seller's *score*. We note that electronic marketplaces typically also allow users to search for more detailed information on the seller's past ratings; nevertheless, aggregated statistics play an important role in buyers' decisions, because of the time and cognitive cost required to go through all available information.

Empirical studies have shown that sellers with high scores enjoy a *price premium*: on average they sell at higher prices than sellers with lower scores. For example, a controlled experiment on vintage postcard auctions showed that an experienced seller enjoyed an 8% price premium [61], and data from eBay auctions revealed that a one point increase in the percentage of negative ratings led to a 9% decline in sale price [16]. These studies suggest that aggregation mechanisms can effectively link current actions of the seller with future payoffs, and thus may incentivize the seller to act in a way that benefits buyers and promote trust.

With the goal of incentivizing truthful advertising, we study how the marketplace is affected by (i) the *aggregation mechanism*, i.e., the way ratings are aggregated into the seller's score, and (ii) the *premium function*, i.e., the way buyers interpret the seller's score. Our main contributions are the following.

1. We demonstrate shortcomings of a system where all ratings are weighted equally in the score of the seller. Such a mechanism will typically be unable to guarantee that a seller is truthful after completing a large number of transactions.

2. We define an aggregation mechanism which weights ratings on recent transactions more. We characterize conditions under which it is optimal for the seller to advertise truthfully, and relate seller truthfulness to returns to reputation.

3. We propose approaches to choose the aggregation mechanism that achieves seller truthfulness over the largest range of parameters. We show that mechanisms that use information from a larger number of past transactions tend to provide incentives for patient sellers to be *more* truthful, but for higher quality sellers to be *less* truthful.

In Section 2.1 we introduce the model. We assume that in each period the seller has a high or low value item for sale. At the beginning of each period the seller observes the value of the item and decides how to advertise it. Potential buyers observe the seller's advertisement and his *score*, i.e., an aggregate of the seller's past ratings. Using this information, we postulate that buyers employ simple heuristics [e.g., 69], Bayesian techniques, or some combination of the two, in order to decide how much to bid [61]. The paper, however, does not require specific models for the buyers' behavior.

The expected payment to the seller is a function of his advertisement and his score, called the *premium function*. Motivated by empirical studies [e.g., 31, 16], we assume that the premium function is increasing in the seller's score. Moreover, we assume that the payment is increasing in the seller's advertisement, that is the expected payment that the seller receives is higher when he advertises a high value item. Under these assumptions, a rational seller advertises high value items truthfully, but may exaggerate the value of a low value item in his advertisement. The seller faces the following trade-off: falsely advertising may result in a higher payment now, but also in a bad rating which implies lower scores and payments in the future.

After a transaction the buyer reports whether the seller advertised the item truthfully. Then, the aggregation mechanism uses this information to compute the seller's new score. For most of this chapter we assume that buyers' ratings accurately reflect the sellers' actions; however, we also demonstrate how some of the results can be extended in the presence of imperfect monitoring. Note that here we do not concern ourselves with incentives for buyers to rate truthfully; our focus is entirely on the optimal strategy of the *seller*.

In Section 2.2 we study an aggregation mechanism that weights all ratings equally (the Unweighted Aggregation Mechanism). We show that with this mechanism the seller is incentivized to falsely advertise at some rating histories for a large class of premium functions, and thus demonstrate that this aggregation mechanism is not effective. In particular, the temptation for the seller to become dishonest increases as the total number of ratings increases, because each additional rating has a smaller and smaller effect on the seller's overall rating.

Our result on the Unweighted Aggregation Mechanism supports eBay's recent decision to base the Positive Feedback percentage on the past 12 months of feedback, rather than the entire lifetime of the seller[1]. We refer to the previous mechanism employed by eBay, which was in use from March 2003 until May 2008, as *EBAY03-08*. In *EBAY03-08* all ratings were weighted equally in the information shown next to the item's description, since lifetime statistics were given. Even though detailed information on recent ratings was available in that system (by clicking through to view the seller's profile), buyers did not always spend the time and effort to search for detailed information. For instance, data collected before and after the change in eBay's aggregation mechanism in March 2003 showed that buyers respond to the reputation information that is easiest to access [16]. This suggests that it is

[1] http://pages.ebay.com/help/feedback/scores-reputation.html

better to primarily show information that weights recent ratings more, as validated by our analysis.

In Section 2.3 we study the class of Weighted Aggregation Mechanisms, which weight ratings from recent transactions more heavily. In particular, the score that is shown to the buyers is a weighted average of past ratings of the seller. We study this class of mechanisms because this is a natural way to weight recent ratings more, and because it includes as special cases several well known aggregation mechanisms. The expected payment that the seller receives is a function of his score. Depending on how buyers interpret the seller's score, the payment to the seller may take different forms.

In contrast to the Unweighted Aggregation Mechanism, under the Weighted Aggregation Mechanism it becomes possible, under some additional conditions, to ensure the seller is always truthful. The main contribution of Section 2.3 is to identify for which premium functions it is possible or impossible to have a seller that always advertises truthfully. We show that under increasing returns to reputation (i.e., the premium function is convex) the optimal strategy of a sufficiently patient and sufficiently high quality seller is to advertise honestly. On the other hand, we show that if returns to reputation are decreasing (i.e., the premium function is concave) and there is no premium for low value items, then for any seller it is optimal to falsely advertise at some scores. The intuition behind this result is that when returns to reputation are increasing, the seller suffers a large reduction in future payment due to even a single deviation from truthful advertising. By contrast, this reduction is relatively small when returns to reputation are decreasing. We also study the effect of step premium functions on the seller's optimization, and give conditions under which it is optimal for the seller to be truthful. In this setting, our main finding is that when there is no premium for low value items to make a sufficiently patient and sufficiently high quality seller truthful, the mechanism must give sufficient weight to the two most recent ratings.

In Section 2.4 we address the design question of choosing the right Weighted Aggregation Mechanism. We note that online marketplaces often change their aggregation mechanisms, which indicates that it is hard to design an effective aggregation mechanism. We first study the Window Aggregation Mechanism, a widely used Weighted Aggregation Mechanism which only reveals the percentage of positive ratings within some fixed window. We define an optimal window size as one which maximizes the range of parameters for which the seller is truthful, and we study the dependence of the optimal window size on the parameters of the model. The optimal sufficient statistic (i.e., window size) neither has very short-term memory (since then current actions do not affect future scores), nor has very long term memory (since then the seller is not incentivized to be truthful). The best choice of averaging window lies between these extremes. We also note an interesting qualitative tradeoff in the choice of window size: informally, *increasing the window size is more likely to make patient sellers truthful, while it is less likely to make high quality sellers truthful.* We show the tradeoff between incentivizing patient and high quality sellers for a range of monitoring settings. First, we consider the case of *perfect monitoring*, where the buyers rate the seller accurately. Then, we consider two types of *imperfect monitoring*: the seller may not receive ratings after some transactions, and the rating that the seller receives may not accurately reflect his action.

We then formulate the design problem for the general class of Weighted Aggregation Mechanisms; here the goal is to choose the optimal vector of weights applied to past ratings to compute a score. Our results here match the insights obtained in our study of the Window Aggregation Mechanism: mechanisms that use information from a larger number of past transactions tend to provide incentives for patient sellers to be *more* truthful, but for higher quality sellers to be *less* truthful. Moreover, our insight suggests that the Window Aggregation Mechanism (which weights all ratings within some fixed window equally) is not optimal and a more robust mechanism might be one that weights more recent ratings even more heavily than older ratings. This is particularly interesting when one considers that nearly all online marketplaces that use mechanisms which only weight recent ratings, tend to use a Window Aggregation Mechanism to do so.

In Section 2.5 we demonstrate that the results also hold for a setting with multiple possible values for the item and the ratings, where a buyer's rating depends on the difference between the advertised and true value of the item. This setting shows our results are robust even if the seller has potentially many "levels" of dishonesty possible.

Related Work

In this section we discuss related literature on seller reputation that can be applied to an online market setting. There are two main approaches. In the adverse selection approach the seller has a hidden type which buyers are trying to learn. Alternatively, reputation can only be used to incentivize the seller to behave well. A standard assumption is that a long-lived seller interacts with short-lived buyers.

In the adverse selection approach, the seller's reputation is studied in a Bayesian setting. The seller is assumed to have a type which buyers are trying to learn from his past ratings. In this setting, the seller's reputation is the belief that buyers have about his type after observing the available information on his past behavior.

A common assumption in the adverse selection literature on reputation is that with some (small) probability the seller's type is such that he always plays an action that promotes trust [49]. If the buyers have access to all past ratings of the seller, they are eventually going to discover his type [21], which means that reputations are not sustainable; this is similar to our insight regarding the Unweighted Aggregation Mechanism (cf. Section 2.2). In this setting, information censoring can result in sustainable reputation [25]. Also, if buyers have to pay to discover the seller's past behavior, then equilibrium behavior is cyclical: the seller builds his reputation up only to exploit it [46].

In the setting we study, the seller does not have a hidden type and the seller's score is only used to incentivize good behavior by the seller. The objective of the reputation mechanism is to induce sellers to behave in a way that promotes trust. This approach has also been taken by [22] and [26].

Dellarocas studies a setting where the seller has two possible effort levels which buyers observe imperfectly [22]. He shows that there is no equilibrium where the seller always exerts high effort, and that eBay's simple mechanism is capable of inducing the maximum theoretical efficiency. In this paper we take a non-equilibrium approach. We consider the best response of the seller to a *fixed* Markov strategy of the buyers; that is, a fixed payment function which only depends on the information available to buyers. Our work takes the point of view that aggregation mechanisms calculate sufficient statistics of the past, and act as the "state" in the interaction between buyers and sellers. We believe that this nonequilibrium approach is reasonable in practice, because of the coordination required for "convergence" to equilibrium. The large and dynamic set of participants in the major online markets makes the rationality, knowledge, and coordination required for equilibrium improbable.

[26] considers a similar setting, where the seller exerts high or low effort, but take a non-equilibrium approach and assume that the seller has a belief over the average bidder behavior. For a specific behavior of bidders it is shown that two simple aggregation mechanisms are inadequate and propose exponential smoothing instead. We consider a general class of payment functions and aggregation mechanisms. Moreover, in our model, the seller decides whether to advertise truthfully and not whether to choose a high or low cost action. Thus, in this paper the goal of an aggregation mechanism is to incentivize truthful behavior by the seller.

Although we do not consider incentives for buyers to leave honest feedback in this paper, another line of research considers how truthful feedback can be elicited. In online markets agents may undertake fake transactions in order to enhance their reputation. This can be avoided if a specific relation between the reputation premium and the transaction cost holds [12]. Alternatively, even if fake transactions can not be undertaken, buyers may not leave honest feedback after a transaction. [52] devises a scoring system that induces honest reporting of feedback. In this paper, we assume that buyers always leave truthful feedback in order to focus on the seller's decisions.

We conclude by noting that a number of papers have empirically studied the effect that the seller's score has on the average payment he receives (a survey is provided in [60]). Some sample studies include data about eBay auctions for coins [48]; Palm Pilots [43]; Pentium III processors [37]; collectible coins, Thinkpads and Beanie babies [16]; and postcards [61]. The last of these papers [61] conducts a controlled experiment on eBay to study the price premium of an experienced seller. Finally, [31] and [56] study the effect of different dimensions of a seller's reputation on pricing power by considering text comments on Amazon and eBay respectively.

2.1 Model

We consider a single seller who is a long-lived player with discount factor δ. The seller interacts with short-lived potential buyers, i.e., buyers who are interested in the seller's item for exactly one round and then depart. We do not explicitly model the market mechanism at each time period, and instead abstract the aggregate behavior of all the buyers in a single time period via a single premium function, described further below.

In every period the seller has an item for sale whose value is high (v_H) with probability q_H and low (v_L) with probability $1 - q_H$. We assume that $0 \leq v_L < v_H \leq 1$. The seller observes the value of the item at the beginning of a period and decides what advertisement to post. Potential buyers observe the seller's advertisement and his *score*, i.e., an aggregate of the seller's past ratings. The expected payment to the seller is a function of his advertisement and his score.

After purchase, the buyer accurately reports whether the seller described the item correctly.[2] The rating is good if the seller described the item truthfully and bad otherwise (i.e., a binary rating system). For simplicity, for most of this paper we assume that both the item's value and the rating can only have two possible values; however, these assumptions can be weakened. In Section 2.5 we demonstrate that our results can be generalized for a setting with multiple possible values for the item and ratings, where the value of the rating depends on how much the seller exaggerates the item's value in his description.

An *aggregation mechanism* specifies the rules for calculating the *score* (i.e., the information shown to potential buyers) from past ratings of the seller. Let r_i be the i-th most recent rating of the seller. In particular $r_i = 1$ or $r_i = 0$ depending on whether the corresponding rating was good or bad. Let $\vec{r} = (r_0, r_1, \ldots)$ be the vector of ratings that the seller has received up to now. We denote the aggregation mechanism by $s(\vec{r})$. In particular, $s(\vec{r})$ is a function that maps rating vectors to scores. Both the seller and the mechanism know $\vec{r}$, but potential buyers only observe $s(\vec{r})$. The score may either be a scalar or a vector. In the next sections we consider both cases: in the Unweighted Aggregation Mechanism (Section 2.2) it is a vector consisting of the total number of ratings and the number of good ratings, while in the Weighted Aggregation Mechanism (Section 2.3) it is a scalar.

[2]In Section 2.4.2 we also consider imperfect monitoring, where buyers' ratings do not always accurately reflect the seller's action.

We assume that the expected payment that the seller receives when his score is s and he chooses advertisement a, where a is either high or low is $a \cdot b(s)$; we call $b(s)$ the *premium function*. We assume that the premium function is non-negative and non-decreasing.[3]

Two realistic assumptions are incorporated in the form of the premium function we chose. First, better ratings yield higher payments to the seller, as has been shown by empirical studies [e.g., 31, 16]. Second, the premium function is increasing in the advertisement a.

The seller chooses a policy that is a best response to the premium function $b(\cdot)$. In our model, we emphasize that the seller is not intrinsically honest or dishonest; he is rational and chooses the advertisement that maximizes his payoff. This is in contrast to the adverse selection approach, where the seller has an intrinsic type (see Section 2). Under the assumptions we made on the payment, it is optimal for the seller to advertise a high value item truthfully.

Let $V(\vec{r})$ be the maximum infinite horizon discounted payoff of the seller when his current vector of ratings is $\vec{r}$. The seller's optimal policy is given by solving the following dynamic program.

$$V(\vec{r}) = q_H(v_H \cdot b(s(\vec{r})) + \delta \cdot V(1,\vec{r})) + (1-q_H)\max\{v_H \cdot b(s(\vec{r})) + \delta \cdot V(0,\vec{r}), v_L \cdot b(s(\vec{r})) + \delta \cdot V(1,\vec{r})\} \tag{2.1}$$

In particular, with probability q_H the seller has a high value item for sale, which he advertises truthfully. The immediate payment he receives is $v_H \cdot b(s(\vec{r}))$ and his ratings "increase" to $(1,\vec{r})$. With probability $1-q_H$ the seller has a low value item for sale. If he advertises it as a high value item, his payoff is $v_H \cdot b(s(\vec{r})) + \delta \cdot V(0,\vec{r})$, since he receives $v_H \cdot b(s(\vec{r}))$ now, but his ratings "decrease" to $(0,\vec{r})$. If he advertises truthfully, he receives a low payment now ($v_L \cdot b(s(\vec{r}))$), but his ratings "increase" to $(1,\vec{r})$. The seller will choose the advertisement with the maximum payoff.

We say that the seller is *truthful at* $\vec{r}$ if it is optimal for him to advertise a low value item truthfully when his vector of ratings is $\vec{r}$. By (2.1), it is optimal for the seller to be truthful

[3] In particular, we assume that for any $n,m \geq 0$, we have $b(s(\vec{x};1;\vec{y})) \geq b(s(\vec{x};0;\vec{y}))$ for all $\vec{x} \in \{0,1\}^n, \vec{y} \in \{0,1\}^m$.

at $\vec{r}$ if and only if

$$(v_H - v_L)b(s(\vec{r})) \leq \delta(V(1,\vec{r}) - V(0,\vec{r})). \tag{2.2}$$

In particular, if the seller is untruthful, his current payoff will increase by $(v_H - v_L)b(s(\vec{r}))$ but his expected payoff starting from the next period will decrease by $V(1,\vec{r}) - V(0,\vec{r})$ (relative to being truthful).

We use this model to study the seller's optimal strategy under various aggregation mechanisms. In particular, we are interested in which aggregation mechanisms and which premium functions induce truthful behavior.

2.2 Unweighted Aggregation Mechanism

In this section, we consider an aggregation mechanism that weights all ratings equally, and show that the seller has an incentive to falsely advertise after a sufficiently large number of ratings. This suggests that weighting recent ratings more is a necessary condition for efficiency.

In the *Unweighted Aggregation Mechanism*, the seller's reputation score consists of the total number of ratings (s_T) and the number of positive ratings (s_P). In particular, given the vector of past ratings $\vec{r}$, we have $s_T(\vec{r}) = |\vec{r}|$ and $s_P(\vec{r}) = |\{r_i : r_i = 1\}|$. (We use the notation $|\cdot|$ to denote both the number of components of a vector and the cardinality of a set.) At score (s_P, s_T), we assume that the seller receives payment $a \cdot b(s_P, s_T)$ for advertising an item of value a.

Proposition 1 shows that under some assumptions on the premium function, the seller will eventually be better off falsely advertising a low value item. The intuition is that when all ratings are weighted equally, after a large number of ratings one more positive rating does not make an appreciable difference to the seller's payoff. This was partially the case with eBay's aggregation mechanism until May 2008 (*EBAY03-08*), where the information that was shown to potential buyers on the item description page weighted all ratings equally. For clarity, all proofs are in the appendix.

Proposition 1 *If $b(s_T, s_T)$ is bounded away from zero as $s_T \to \infty$ and*

$$b(s_T, s_T) - b(s_T - 1, s_T) \to 0 \text{ as } s_T \to \infty, \tag{2.3}$$

then there exists an s_T at which it is optimal for the seller to falsely advertise a low value item.

Proposition 1 relies on two assumptions. The assumption on $b(s_T, s_T)$ is not particularly restrictive: we expect the payment to a seller with maximum reputation to be bounded away from zero after a large number of ratings.

Condition (2.3) is a regularity condition on the slope of the premium functions; this formalizes the idea that as the total number of ratings increases, the marginal effect of each additional positive rating on the seller's expected payment eventually becomes negligible. This condition is not particularly restrictive; e.g., see Example 2. On the other hand, if Condition (2.3) does not hold, then it may be optimal for a sufficiently patient seller to be always truthful. In particular, this is the case if a single bad rating causes a discrete drop in premium regardless of how many positive ratings have been received, i.e., if there exist T and $\alpha > 0$ such that $b(s_T, s_T) - b(s_T - 1, s_T) > \alpha$ for all $s_T \geq T$, and $(v_H - v_L)b(T', T') \leq (q_H v_H + (1 - q_H)v_L)\alpha\delta/(1 - \delta)$ for all $T' \geq T$.

We briefly discuss related results in the literature. A similar result is shown by Fan et al. for a specific premium function and two specific unweighted reputation mechanisms [26]. Moreover, we observe that Proposition 1 is similar to the result Cripps et al. that reputations are not sustainable with imperfect monitoring and incomplete information [21]. This is a different setting than ours, since buyers are optimizing with respect to the seller's strategy, while we assume that the buyers' behavior only depends on the aggregation mechanism that is being used, and the seller's score and advertisement. The result in [21] depends on imperfect monitoring: the seller chooses low effort with some small probability, because it will not significantly affect the buyers' beliefs in the near future, and eventually buyers learn the seller's type. In our model, after a large number of transactions it becomes profitable for the seller to be untruthful, even though monitoring is perfect.

Example 1 ***EBAY03-08***

Figure 2.1: Sample eBay auction listing. 99% of the ratings that user satmr2 received in the last twelve months were positive.

The seller's score shown by the EBAY03-08 mechanism next to the description of the item consisted of: (1) the difference between positive and negative ratings (i.e., $2s_P - s_T$), and (2) the ratio of positive ratings over the total number of ratings (i.e., s_P/s_T). Users could access more information on the seller's past ratings by clicking on the seller's pseudonym. However, if many users did not spend the time to search for more information when bidding, this aggregation mechanism would not be effective in the long run for a large class of premium functions. In particular, it is possible that a seller would exaggerate the value of the item once he had a large number of ratings. This intuition is supported by data analyzed by [60], which show a decline in performance once the seller completed a large number of transactions. This observation may have influenced eBay's decision to change the initially shown information; the new system includes statistics about the ratings the seller received in the last twelve months (cf. Figure 2.1). □

Proposition 1 shows that the seller will eventually be better off falsely advertising a low value item. How many transactions pass before the seller is tempted to be dishonest? We can compute a bound on this time by upper bounding the increase in payment due to one more positive rating with a function of the total number of ratings. In particular, suppose there exists a function $f(s_T)$ such that

$$\max_{0 \le s_P \le s_T - 1} \{b(s_P + 1, s_T) - b(s_P, s_T)\} \le f(s_T).$$

Let s_T^* denote the first time a seller would choose to falsely advertise a low value item; it follows from the proof of Proposition 1 that an upper bound for s_T^* is:

$$\hat{s}_T = \min\left\{ s_T : (v_H - v_L) b(s_T, s_T) \geq \frac{\delta}{1-\delta} f(s_T) \right\}. \tag{2.4}$$

By computing $\hat{s}_T$ we know that the seller will not advertise low value items truthfully for more than $\hat{s}_T$ consecutive transactions. The following example upper bounds s_T^* for a specific premium function.

Example 2 *Let $b(s_P, s_T) = \alpha/(\alpha + 1 - s_P/s_T)$ for some $\alpha \in (0,1)$; this model arises when buyers use a certain maximum likelihood estimate based on the rating history, [cf. Appendix 2, 7]. Note that this function only depends on the proportion of positive ratings. Further $b(s_T, s_T) = 1$, and $b(s_P + 1, s_T) - b(s_P, s_T) \leq 1/(s_T \cdot \alpha)$. Let $f(s_T) = 1/(s_T \cdot \alpha)$. By (2.4),*

$$\hat{s}_T = \left\lceil \frac{\delta}{1-\delta} \frac{1}{\alpha(v_H - v_L)} \right\rceil.$$

This upper bound on the number of consecutive truthful transactions of the seller is increasing in δ, the seller's discount factor. This is something we expect: the more the seller values the future, the greater the effectiveness of the aggregation mechanism, and thus the greater his incentive to tell the truth. □

2.3 Weighted Aggregation Mechanism: Characterization

In Section 2.2 we showed that weighting all ratings equally does not incentivize truthfulness. In this section we analyze mechanisms which put more weight on recent ratings, and show that it is possible to incentivize truthfulness and promote trust under these mechanisms.

We first introduce the *Weighted Aggregation Mechanism*; as we will discuss, special cases of this mechanism are widely used in practice. Let r_i be the value of the i-th most recent rating. We consider an aggregation mechanism where buyers only see the following score:

$$s(\vec{r}) = \sum_{i=0}^{\infty} w_i \cdot r_i,$$

where $w_i \geq 0$, and $w_0 \geq w_1 \geq w_2 \geq \cdots$, with $\sum_i w_i = 1$. We assume that the weights w_i are non-increasing in i so that recent ratings are weighted more. We make the seller's score a scalar, because aggregation mechanisms typically rely primarily on a summary of ratings.

According to the general model introduced in Section 2.1, at score s, the seller receives expected payment equal to the product of his advertisement and the premium function $b_{\vec{w}}(s)$. We use the subscript $\vec{w}$ on the premium function to denote that the payment may explicitly depend on the aggregation mechanism. *We assume that $b_{\vec{w}}(s)$ is increasing in s for each vector $\vec{w}$.*

Both the seller and the mechanism have access to $\vec{r}$; the seller remembers his past actions and the mechanism keeps this information in order to update the seller's score. However, as we discuss in the following examples of Weighted Aggregation Mechanisms, it may not be necessary to keep the whole vector of ratings.

Example 3 *In the* Window Aggregation Mechanism, *$w_i = 1/T$ for $i = 0, 1, ..., T-1$, and $w_i = 0$ for $i \geq T$, for some $T \geq 1$, which we call the window size. Thus the score is the percentage of good ratings that the seller received in the last T transactions. To compute the seller's score, the aggregation mechanism needs to keep information on the T most recent ratings of the seller. This score is widely used in online marketplaces, such as eBay and Amazon, and has been studied in various settings [22, 7].* □

Example 4 *In the* Exponential Aggregation Mechanism *$w_i = (1-\alpha)\alpha^i$, for some $\alpha \in (0,1)$. To compute the seller's score, the aggregation mechanism only needs to know the score of the seller in the previous period ($\hat{s}$) and his most recent rating (r_0). Then the new score is $\alpha\hat{s} + (1-\alpha)r_0$. We note that the Exponential Aggregation Mechanism has been previously suggested as a good design [26]. Moreover, even though to the best of our knowledge exponential smoothing is not being used by any electronic marketplace to promote trust, it is a good model of how people update their impressions without an aggregation mechanism in place [e.g., 4, 36, 44].* □

We study the Weighted Aggregation Mechanism because (1) it is a natural way to weight recent ratings more, and (2) it is a generalization of both the Window Aggregation Mechanism and the Exponential Aggregation Mechanism. However, there are other ways to summarize feedback by incorporating text comments (which are more descriptive than numerical ratings) and weighting recency. Such a summary score could offer richer information without increasing the buyers' search costs [56].

In Section 2.3.1 we derive necessary and sufficient conditions for the truthful policy to be optimal for the seller. We use these conditions in subsequent sections to show that efficiency can be achieved under strictly increasing returns to reputation and may be possible under decreasing returns to reputation (Section 2.3.2), and under a step premium function (Section 2.3.3). In Section 2.3.4 we discuss empirical studies on eBay and Amazon. In Section 2.3.5 we show that under an additional assumption on the premium function, we can simplify the sufficient condition for the seller to be always truthful; this assumption is valuable in part because it is satisfied by all logarithmically concave premium functions. We use this assumption in Section 2.3.6, where we identify a dominance relation between premium functions such that if b_1 dominates b_2, then b_1 better incentivizes truthful advertisement.

2.3.1 Conditions for Optimality of Truthfulness

We are interested in whether it is optimal for the seller to be always truthful. It is optimal for the seller to be always truthful if and only if any one step deviation from the truthful policy (i.e., the policy of always advertising items truthfully) does not yield a higher payoff. Let $\hat{V}(\vec{r})$ be the infinite horizon discounted expected value when the seller's current ratings are $\vec{r}$ and the seller is always truthful. Let $s_i(\vec{r})$ be the seller's score after i periods if his current ratings are $\vec{r}$ and receives positive ratings in the following i periods. Note that $s_0(\vec{r}) = s(\vec{r})$. This implies that $\hat{V}(1,\vec{r}) - \hat{V}(0,\vec{r}) = (q_H v_H + (1-q_H)v_L)\sum_{i=0}^{\infty}\delta^i(b_{\vec{w}}(s_i(1,\vec{r})) - b_{\vec{w}}(s_i(0,\vec{r})))$ and $s_i(1,\vec{r}) - s_i(0,\vec{r}) = w_i$. By (2.2), the seller is not better off deviating from the truthful policy when his ratings are $\vec{r}$ if and only if $(v_H - v_L)b_{\vec{w}}(s(\vec{r})) \leq \delta(\hat{V}(1,\vec{r}) - \hat{V}(0,\vec{r}))$. We conclude that it is optimal for the seller to be always truthful if and only if the previous condition holds for all $\vec{r}$. Substituting $\hat{V}(1,\vec{r}) - \hat{V}(0,\vec{r})$ in the previous condition, the following lemma is proved.

Lemma 1 *It is optimal for the seller to be truthful at all $\vec{r}$ if and only if*

$$b_{\vec{w}}(s(\vec{r})) \leq \frac{q}{v_H - v_L} \sum_{i=0}^{\infty} \delta^{i+1} (b_{\vec{w}}(s_i(1,\vec{r})) - b_{\vec{w}}(s_i(0,\vec{r}))) \tag{2.5}$$

for all rating vectors $\vec{r}$.

Thus, if we know the premium function $b_{\vec{w}}(\cdot)$ and the seller's parameters δ and q, we can check whether it is optimal for the seller to be truthful by checking whether condition (2.5) is satisfied for all $\vec{r}$.

2.3.2 Increasing and Decreasing Returns to Reputation

Depending on how buyers interpret the information on the seller available to them, i.e., the score $\sum_i w_i \cdot r_i$, the payment to the seller may exhibit increasing or decreasing returns to reputation. Increasing returns to reputation correspond to a convex premium function, i.e., there are increasing marginal benefits to higher reputation. On the other hand, decreasing returns to reputation are associated with a premium function that is concave in the seller's score.

The following lemma indicates that it is easier to incentivize truthfulness under increasing returns to reputation, since it is more likely that (2.5) holds.

Lemma 2 *(i) If the premium function $b_{\vec{w}}$ is strictly convex, $w_0 < 1$ and $b_{\vec{w}}(0) = 0$, then*

$$b_{\vec{w}}(s(\vec{r})) < \sum_{i=0}^{\infty} (b_{\vec{w}}(s_i(1,\vec{r})) - b_{\vec{w}}(s_i(0,\vec{r})))$$

for every $\vec{r}$.

(ii) If the premium function $b_{\vec{w}}$ is concave and $b_{\vec{w}}(0) \geq 0$, then

$$b_{\vec{w}}(s(\vec{r})) \geq \sum_{i=0}^{\infty} (b_{\vec{w}}(s_i(1,\vec{r})) - b_{\vec{w}}(s_i(0,\vec{r})))$$

for $\vec{r} = \vec{1}$.

The following proposition demonstrates a dichotomy between convex and concave premium functions when $v_L = 0$ and $b_{\vec{w}}(0) = 0$. If the premium function exhibits increasing returns to reputation, and if δ and q_H are sufficiently large, then it is possible to achieve our goal of making the seller truthful regardless of his rating history. On the other hand, if the payment is concave in the score, then it is not optimal for the seller to be truthful at all rating histories; in particular, (2.5) can not be satisfied since if $v_L = 0$, then $q/(v_H - v_L) = q_H < 1$.

Proposition 2 *(i) If $v_L = 0$, the payment function $b_{\vec{w}}$ is strictly convex, $w_0 < 1$ and $b_{\vec{w}}(0) = 0$, then it is optimal for the seller to be always truthful for all sufficiently large $\delta < 1$ and $p_H < 1$.*

(ii) If $v_L = 0$, the payment function $b_{\vec{w}}$ is concave and $b_{\vec{w}}(0) \geq 0$, then there do not exist values of $\delta < 1$ and $p_H < 1$ for which it is optimal for the seller to always advertise truthfully.

If $v_L = 0$, the condition $w_0 < 1$ is necessary for truthfulness under convex premia. In particular, if only the most recent rating affects the seller's score, then it will be optimal for the seller to exaggerate the value of a low value item in his advertisement when his score is high. This happens because the seller discounts future payments and prefers a high payment now to a high payment in the next period. We conclude that if the payment is convex and there is no premium for low value items, then the weights on the two most recent ratings must both be strictly positive to induce truthfulness.

If v_L is strictly greater than zero, then it may be possible to have a truthful seller under a concave premium. For any concave strictly increasing premium, there exists a sufficiently large v_L such that it is optimal for the seller to be always truthful for a given δ and q_H. In particular, this is the case if

$$\frac{v_L}{v_H - v_L} \geq \max_{\vec{r}} \left\{ \frac{b_{\vec{w}}(s(\vec{r}))}{\sum_{i=0}^{\infty} \delta^{i+1}(b_{\vec{w}}(s_i(1,\vec{r})) - b_{\vec{w}}(s_i(0,\vec{r})))} \right\} - q_H.$$

We conclude that it is possible to have a truthful seller under both convex and concave premium functions. However, increasing returns to reputation tend to better incentivize truthfulness.

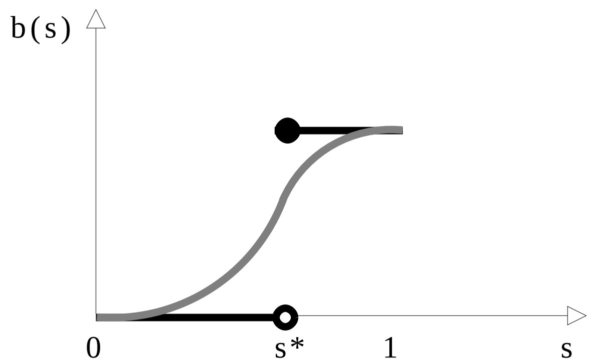

Figure 2.2: Sigmoid and step payments. The grey line represents a premium function that is convex at low scores and concave at high scores. The black line represents a step function that approximates it.

2.3.3 Step Functions

In Section 2.3.2 we considered premium functions that exhibit either increasing or decreasing returns to reputation throughout their domains. Another natural class of payments are step premium functions, which are functions that are constant throughout most of their domains. Our motivation for studying this class of functions is that it may provide a good approximation for the expected payment in electronic marketplaces. These premium functions capture the following intuition: a potential buyer is not willing to buy an item from a seller that has a very low score; however, there is some threshold on the seller's score, above which the buyer trusts the seller and is willing to pay up to his valuation for the item. Since different buyers may have different thresholds and different valuations for the item, the premium function may be a smoothed step function, e.g., a function that is convex at low scores and concave at large scores (see Figure 2.2). In this section we study step premium functions, and note that similar results can be obtained for smooth approximations.

In this section we show that with step premium functions, optimality of the truthful policy heavily depends on the aggregation mechanism, and in particular on the magnitude of w_1 (i.e., the second largest weight) relative to the threshold of the payment. We then apply the result to the Window Aggregation Mechanism, which was introduced in Example 3.

We formally define a step premium function in the following definition.

Definition 1 *The premium function $b_{\vec{w}}$ is a* step function *with threshold $s^* \in (0,1]$ if $b_{\vec{w}}(s) = b_{\vec{w}}(1) > 0$ for $s \geq s^*$, and $b_{\vec{w}}(s) = 0$ otherwise.*

We are interested in understanding the conditions under which a Weighted Aggregation Mechanism can make a seller truthful when the payment is a step function. A necessary condition is $s^* > 1 - w_0$; otherwise the seller will not be truthful when his score is equal to 1. If $v_L/(v_H - v_L)$ is large then it may be possible to incentivize the seller to be truthful as long as $s^* < 1 - w_0$. In this section we focus on the case that $v_L = 0$. Our main result is the following proposition.

Proposition 3 *Suppose the premium function $b_{\vec{w}}$ is a step function with threshold s^* and $v_L = 0$.*

(i) If $s^* \leq 1 - w_1$, then there are no values of $q_H < 1$ and $\delta < 1$ for which it is optimal for the seller to be always truthful;

(ii) If $s^* > 1 - w_1$, then it is optimal for the seller to be truthful for all sufficiently large $\delta < 1$ and $q_H < 1$.

We note that w_1 is the weight on the second most recent rating and that $w_0 \geq w_1$. The result provides a useful insight: recent ratings, in particular the one and two period old ratings, must be weighted sufficiently to ensure that sellers can be made truthful.

As a specific application, we consider the Window Aggregation Mechanism, which was introduced in Example 3, and a premium function that yields a positive expected payoff for at least some score less than 1. These premium functions are reasonable in many applications: buyers often trust a seller with an almost maximum score. In the following example, we show that even though for any fixed window size it is not optimal for the seller to be always truthful, truthfulness may be optimal when information from multiple window sizes is aggregated. This suggests that the use of multiple window sizes, a common practice in online marketplaces, may better incentivize truthfulness.

Example 5 *Assume that a Window Aggregation Mechanism with window size T is used and $v_L = 0$. The payment is a step function with threshold s^*. We consider premium functions that are positive for at least some possible scores other than 1 (the maximum possible* *score). This implies that $s^* \leq 1 - 1/T$. Thus, $s^* \leq 1 - w_0$. Then, there do not exist $\delta < 1$ and $q < 1$ for which it is optimal for the seller to be always truthful.*

Surprisingly, there is a way to aggregate information from multiple window sizes, so that it is optimal for the seller to always be truthful under step functions that are not only positive at the maximum possible score. For example, suppose the aggregation mechanism is using information on k windows with sizes $T_k \geq T_{k-1} \geq ... \geq T_1 \geq 2$. Let s_i be the score on window i and assume that the aggregate score of the seller is the average of the scores for each window, i.e., $s = (\sum_{i=1}^{k} s_i)/k$. Thus $w_1 = (1/k)\sum_{i=1}^{k} 1/T_i$, while the smallest positive weight is $w_{T_k-1} = 1/(kT_k)$. In this case, by Proposition 3, the seller can be made truthful if $1/(kT_k) \leq 1 - s^* < (1/k)\sum_{i=1}^{k} 1/T_i$; depending on the value of s^*, these inequalities can be satisfied if more than one window is used. □

2.3.4 Empirical Insights

In Sections 2.3.2 and 2.3.3, we have studied the Weighted Aggregation Mechanism under various premium function forms. A natural question arises: what is the form of the premium function in electronic marketplaces? In this section we attempt to answer this question using results from empirical studies on eBay and Amazon. In particular, we wish to understand the dependence of the expected payment on the percentage of positive ratings (since this is the information shown in these markets that best matches the seller's score in a Weighted Aggregation Mechanism). We note however, that eBay's mechanism has changed multiple times within the last years, and thus different studies collect data on different versions of the mechanism.

Two empirical studies explore whether negative feedback is associated with increasing or decreasing returns, and suggest increasing returns in the percentage of positive ratings on eBay.[4]

[4]There have also been studies that find decreasing returns in reputation; however, these studies look at different metrics than the one we are interested in here. [47] studies the effect of the total number of positive ratings on the expected payment to the seller, and finds severely decreasing marginal returns. [53] finds decreasing marginal returns in the difference between the total number of positive and negative ratings. However, these results do not apply here, because we are interested in the effect of the percentage of positive ratings.

1. Kalyanam and McIntyre collect data from Palm Pilot auctions on eBay and conjecture diminishing impact of negative feedback [43]. In particular, they suggest that the dependence of price on percentage of negative feedback is $-f(\text{Negative }\%)$ for some concave function f, which implies increasing returns in the percentage of positive ratings.

2. A panel data set from eBay indicates that the first negative is very harmful for the seller's growth rate, while subsequent negatives have lower impact [16]. This implies increasing returns in the percentage of positive feedback for the growth, and suggests a similar effect for the expected payment to the seller.

Other studies suggest a specific dependence based on the regression that is used, without explicitly stating it. For instance, an OLS regression on data from Amazon's retail platform suggests that the price premium is convex in the number of stars, a measure similar to the percentage of positive ratings [31]. On the other hand, there are also studies which find that the percentage of positive feedback has no effect on the selling price [e.g., 61]. We note that eBay was only showing the difference between the number of positive and negative ratings (and not the percentage of positive ratings) next to the description of the item when this experiment was conducted. Nevertheless, [43] and others collected data under the same mechanism and found that the percentage of negative ratings had an effect on the expected price.

To interpret these contradictory results, we need to consider that most eBay sellers have an extremely high percentage of positive ratings [60]. Low variability in the percentage of positive ratings may be one reason that, for some data sets [e.g., 61], the price does not significantly depend on the percentage of positive ratings. Another possible interpretation is that the payment does not depend on the exact percentage of positive ratings, above a certain threshold; this behavior is consistent with a step premium function. On the other hand, the results of [43] and [16] suggest increasing returns at least in some part of the domain. Plausible interpretations are that the premium function may either exhibit increasing returns throughout its domain, or have the form of a sigmoid function (i.e., increasing returns initially and decreasing returns for large values of the percentage, cf. Figure 2.2). The latter is a smooth approximation of a step function, which was studied in Section 2.3.3.

The fact that there are such high percentages of positive ratings on eBay may also shed light on the form of the premium function. In particular, potential explanations of this phenomenon are that (1) the reputation system works well in general, (2) users tend to not report negative ratings out of fear of retaliation [23], and (3) sellers with bad records exit and possibly re-enter under a new identity. We next give evidence of (3) and argue that it implies that a low percentage of positive ratings significantly decreases the expected payment to the seller.

A seller with a non-negligible percentage of negatives may re-enter the market under a new identity, especially given the high levels of competition on eBay (anecdotal evidence is provided in [43]). Moreover, data from eBay show that an increase in the percentage of negatives in a seller's record translates into an increase in exit probability [16]. Such behavior may imply that a seller without an almost perfect percentage of positive ratings is better off rejoining the system than keeping his current score. This suggests that a low percentage of positives has a significant effect on the expected payment to the seller, and is consistent with increasing returns to reputation or a step premium function.

2.3.5 A Simplified Condition for Truthfulness

In this section we present an assumption that simplifies the task of checking whether the seller is always truthful; notably, this assumption is satisfied whenever the premium function is logarithmically concave. We use this assumption in the following subsection when we study dominance relationships among premium functions, as well as throughout our analysis in Section 2.4.

Assumption 1 *For any parameters δ and q_H, if it is not optimal for the seller to deviate from the truthful policy at $\vec{r} = \vec{1}$, then it is optimal for him to be truthful at all rating vectors.*

Assumption 1 says that if (2.5) holds at $\vec{r} = \vec{1}$, then (2.5) holds at all $\vec{r}$, and it is optimal for the seller to always be truthful. Thus, if Assumption 1 holds, then in order to check whether truthfulness is optimal for the seller, it suffices to check whether (2.5) holds at $\vec{r} = \vec{1}$. Using the facts that $s(\vec{1}) = 1$, $s_i(\vec{1}) = 1$ and $s_i(0, \vec{1}) = 1 - w_i$, we can prove the following lemma.

Lemma 3 *If Assumption 1 holds, then it is optimal for the seller to be truthful at all $\vec{r}$ if and only if*

$$(v_H - v_L)b_{\vec{w}}(1) \leq q\sum_{i=0}^{\infty}\delta^{i+1}(b_{\vec{w}}(1) - b_{\vec{w}}(1-w_i)). \tag{2.6}$$

The following lemma shows that a large class of premium functions satisfies Assumption 1. In particular, the assumption holds for premium functions whose logarithm is concave in the seller's score.

Lemma 4 *A logarithmically concave premium function satisfies Assumption 1.*

We observe that if the premium function $b_{\vec{w}}$ is convex, then its logarithm may or may not be concave depending on how fast the slope of the premium function increases compared to the payment itself. We note, however, that log-concavity is not a particularly restrictive assumption. For example, the functions x^n, e^{nx} are logarithmically concave for any $n > 0$.

We note that various empirical studies suggest that the expected premium is logarithmically concave in the mean value of ratings of the seller [e.g., 16, 48, 31]. These studies regress the payment to the seller or its logarithm against some function of the average rating of the seller or some other function of the number of positive and negative ratings (in the case of eBay). Even though these studies consider all the ratings that the seller has received in his lifetime as a seller or in the last twelve months, we can get some insight in the dependence of the payment on the percentage of positive ratings in the last T transactions by fixing the total number of transactions to T. [16] uses a data set from eBay and regresses the logarithm of the price against the percentage of positive ratings of the seller. Since the logarithm of the price is a linear function of the average rating, the expected payment is logarithmically concave in the seller's score (by definition). [31] uses a data set from the Amazon Marketplace and regresses the logarithm of the premium against the mean rating of the seller in the last twelve months. We note that the Amazon Marketplace asks buyers to rate sellers out of five stars. Again, this corresponds to a logarithmically concave function. [48] uses a data set from eBay and regresses the price against the logarithm of the number of positive ratings and the logarithm of the number of negative ratings. Setting T equal to the sum of the number of positive and negative ratings, we observe that according to this regression the price is a logarithmically concave function of the seller's score.

2.3.6 Dominance

In this section we define a dominance relation such that if the premium function $b_{\vec{w}}$ dominates the premium function $\tilde{b}_{\vec{w}}$, then optimality of truthfulness under $\tilde{b}_{\vec{w}}$ implies optimality of truthfulness under $b_{\vec{w}}$. The dominance relationship we consider informally captures the idea that a "steeper" premium function should lead the seller to be "more truthful."

To have a fair and interesting comparison between two premium functions $b_{\vec{w}}$ and $\tilde{b}_{\vec{w}}$, they should take the same values at minimum and maximum scores; i.e., $b_{\vec{w}}(0) = \tilde{b}_{\vec{w}}(0)$ and $b_{\vec{w}}(1) = \tilde{b}_{\vec{w}}(1)$. Under this condition and Assumption 1, the dominance relation for truthfulness is $b_{\vec{w}}(s) \leq \tilde{b}_{\vec{w}}(s)$ for all $s \in [0,1]$, as the following proposition shows.

Proposition 4 *Let $b_{\vec{w}}$, $\tilde{b}_{\vec{w}}$ be premium functions such that*

(i) $b_{\vec{w}}(s) \leq \tilde{b}_{\vec{w}}(s)$ *for all* $s \in [0,1]$,

(ii) $b_{\vec{w}}(0) = \tilde{b}_{\vec{w}}(0)$, $b_{\vec{w}}(1) = \tilde{b}_{\vec{w}}(1)$,

(iii) $b_{\vec{w}}$ *satisfies Assumption 1.*

For any $q < 1$, $\delta < 1$, and $\vec{w}$, if it is optimal for the seller to be always truthful under $\tilde{b}_{\vec{w}}$, then it is optimal for him to be always truthful under $b_{\vec{w}}$.

For example, if always advertising truthfully is optimal for a seller with parameters δ, q when the premium function is $\tilde{b}_{\vec{w}}(s) = s^k$ where $k > 1$, then always advertising truthfully is also optimal under $b_{\vec{w}}(s) = s^m$ for any $m \geq k$. In particular, this observation suggests that increasing the elasticity of the premium function improves incentives for the seller to be truthful. This is consistent with our observation in Section 2.3.2 that it is easier to incentivize truthfulness with convex premium functions than with concave premium functions.

2.4 Weighted Aggregation Mechanism: Design

The results of the preceding section show that for a given vector of weights $\vec{w}$ and a given premium function $b_{\vec{w}}$, under certain conditions there exists a range of $q < 1$, $\delta < 1$, for which the seller is always truthful. Can we use this insight to guide the design of the aggregation mechanism, i.e., to maximize the range of parameters for which the seller is always truthful? In this section we consider the problem of *designing* a good aggregation mechanism; in particular, we consider a setting where the system designer is considering choosing the weight vector $\vec{w}$ in a Weighted Aggregation Mechanism.

Throughout this chapter, we have considered a "good" aggregation mechanism to be one that ensures sellers are truthful. Accordingly, we assume that the mechanism designer's goal is to *maximize the range of seller parameters q and δ for which truthfulness can be guaranteed.* Throughout this section we assume that the premium function satisfies Assumption 1. We begin in Section 2.4.1 by formulating the design problem under a range of assumptions regarding the information available to the designer regarding q and δ. We subsequently consider specific examples, corresponding to the Window Aggregation Mechanism (Section 2.4.2) and Weighted Aggregation Mechanism (Section 2.4.4), both of which were previously introduced in Section 2.3. Ultimately, our analysis of this design problem lends qualitative insight into the design of aggregation mechanisms; in particular, we find that aggregation mechanisms that average over a longer past history of ratings are *more likely* to incentivize patient sellers to be truthful, but *less likely* to incentivize high quality sellers to be truthful.

2.4.1 The Design Problem

In this section we formulate a general design problem for the Weighted Aggregation Mechanism. We assume the mechanism designer chooses the weights $\vec{w}$ from a set W; special cases are considered in the subsequent subsections.

By Lemma 3, if Assumption 1 holds, then it is optimal for the seller to be always truthful if and only if

$$b_{\vec{w}}(1) \le \frac{q}{v_H - v_L} \sum_{i=0}^{\infty} \delta^{i+1} (b_{\vec{w}}(1) - b_{\vec{w}}(1 - w_i)).$$

We conclude that *the seller is always truthful if and only if q, δ, and $\vec{w}$ jointly satisfy the* *following constraint:*

$$q \cdot F(\vec{w}, \delta) \geq 1 \tag{2.7}$$

where

$$F(\vec{w}, \delta) = \frac{1}{v_H - v_L} \frac{\sum_{i=0}^{\infty} \delta^i (b_{\vec{w}}(1) - b_{\vec{w}}(1 - w_i))}{b_{\vec{w}}(1)}. \tag{2.8}$$

Note that $F(\vec{w}, \delta)$ is increasing in δ for any fixed $\vec{w}$.

As discussed above, our approach is to maximize the *range* of parameter values for which the seller will be truthful, given available information regarding δ and/or q. In what follows, our analysis depends on analyzing the set $\vec{w}^*(\delta)$ defined as follows for each δ:

$$\vec{w}^*(\delta) = \arg\max_{\vec{w} \in W} F(\vec{w}, \delta). \tag{2.9}$$

For tractability, we assume for the duration of this section that for all δ, the set $\vec{w}^*(\delta)$ is nonempty.

We now consider the optimal choice of weights depending on the available information regarding δ and q.

1. *Both* δ *and* q *are known by the mechanism designer.* In this case, the goal is to find weights $\vec{w} \in W$ such that (2.7) holds. Whether or not this will be possible depends on whether q and δ are large enough, given the premium functions $b_{\vec{w}}$. In particular, we can ensure the seller is always truthful if and only if $\max_{\vec{w} \in W} F(\vec{w}, \delta) \geq 1/q$; in this case any choice of weights in $\vec{w}^*(\delta)$ is optimal.

2. *The mechanism designer knows* δ*, but not* q*.* A reasonable choice of $\vec{w}$ is one which maximizes the range of values of q for which the seller will be always truthful. From (2.7), this implies we should maximize $F(\vec{w}, \delta)$ subject to $\vec{w} \in W$, i.e., any $\vec{w} \in \vec{w}^*(\delta)$ is an optimal choice.

3. *The mechanism designer knows* q*, but not* δ*.* A reasonable choice of $\vec{w}$ is one which maximizes the range of values of δ for which the seller will be always truthful; thus,

given q, we solve:

$$\text{minimize } \delta \tag{2.10}$$

$$\text{subject to } q \cdot F(\vec{w}, \delta) \geq 1; \tag{2.11}$$

$$\vec{w} \in W. \tag{2.12}$$

Let $\delta^*(q)$ denote the optimal value of the preceding problem; this is the smallest value of δ such that a seller with quality q and discount factor δ can be guaranteed to be truthful under *some* weight vector. It then follows that any $\vec{w} \in \vec{w}^*(\delta^*(q))$ is an optimal choice of weights. Observe that since the constraint is increasing in q, it follows that $\delta^*(q)$ is decreasing in q; we will use this fact in our subsequent analysis to characterize the dependence of the set $\vec{w}^*(\delta^*(q))$ on q.

Of course, there is a fourth possibility: the mechanism designer may know *neither* δ nor q. In this scenario, the mechanism designer will typically need to consider whether it is preferable to make sellers truthful over a greater range of δ, or a greater range of q. We delineate such a tradeoff in subsequent subsections.

A critical theme emerges from the preceding discussion: *regardless* of the information available to the mechanism designer, solving (2.9) is an important step in finding the best weight vector. We apply this insight in the following subsections to study both the Window Aggregation Mechanism and the more general Weighted Aggregation Mechanism.

2.4.2 Window Aggregation Mechanism

The Window Aggregation Mechanism (introduced in Example 3) is characterized by the window size, which we denote by T. As previously discussed, it is a special case of a Weighted Aggregation Mechanism, where $w_0 = w_1 = \cdots = w_{T-1} = 1/T$, and all other weights are zero; the feasible set W consists of all weight vectors of this form for window sizes $T \geq 1$. Thus the window size is the only design choice.

In this section, we apply the methodology of the Section 2.4.1 to determine an optimal window. Our main result is that an optimal window size is *increasing* in the discount factor δ, and *decreasing* in the quality q. Thus a longer window is *more likely* to make patient sellers always truthful, but *less likely* to make high quality sellers always truthful.

In a slight abuse of notation, we write $b_T(\cdot)$ for the premium function to indicate that the buyers' behavior may depend on window size T. Throughout the section we implicitly assume that the family of premium functions b_T is known by the mechanism designer. We further assume that $b_T(\cdot)$ is strictly convex and satisfies Assumption 1 for each T; as noted in Lemma 4, this Assumption is satisfied by logarithmically concave premium functions.

In the case of a Window Aggregation Mechanism, by a slight abuse of notation, we redefine F from (2.8) as:

$$F(T,\delta) = \frac{1}{v_H - v_L} \frac{b_T(1) - b_T(1-1/T)}{b_T(1)} \sum_{i=1}^{T} \delta^i.$$

Calculation of the optimal weight vector in (2.9) then reduces to calculating the optimal window size via the following optimization:

$$T^*(\delta) = \arg\max_{T \geq 1} F(T,\delta). \tag{2.13}$$

As before, we assume for the duration of this section that for all δ, the set $T^*(\delta)$ is nonempty.

The following proposition characterizes the behavior of $T^*(\delta)$.

Proposition 5 *If Assumption 1 holds, then $T^*(\delta)$ is increasing in δ in the following sense: for $\delta \geq \delta'$,*

(i) $\max\{T : T \in T^*(\delta)\} \geq \max\{T : T \in T^*(\delta')\}$*; and*

(ii) $\min\{T : T \in T^*(\delta)\} \geq \min\{T : T \in T^*(\delta')\}$.

Surprisingly, note that this result holds *regardless* of the dependence of b_T on T.

Using this proposition, we now consider the optimal choice of window depending on the available information regarding δ and q, as in the preceding subsection.

1. *Both δ and q are known by the mechanism designer.* As in Section 2.4.1, any window size in $T^*(\delta)$ is optimal, and whether the seller can be made truthful depends on

whether q and δ are sufficiently large to ensure that (2.7) holds. In this case any window size in $T^*(\delta)$ is optimal.

2. *The mechanism designer knows* δ*, but not* q*.* As in Section 2.4.1, if we wish to maximize the range of values of q for which the seller will be truthful, then any $T \in T^*(\delta)$ is an optimal choice. Note that in this case, from Proposition 5, *the set of optimal windows is increasing in* δ. This is an intuitive result, since sellers with larger δ are more patient, and thus an aggregation mechanism with longer memory can successfully couple current behavior with distant future payoffs.

3. *The mechanism designer knows* q*, but not* δ. As in Section 2.4.1, let $\delta^*(q)$ denote an optimal solution to (2.10)-(2.12), where again the set W consists of all possible weight vectors corresponding to Window Aggregation Mechanisms. It then follows any $T \in T^*(\delta^*(q))$ is an optimal choice of window size. Recall that $\delta^*(q)$ is decreasing in q. From Proposition 5, we conclude *the set of optimal windows* $T^*(\delta^*(q))$ *is decreasing in* q. This is an intuitive result, since as q increases, it is possible to make less patient sellers truthful, and for such sellers a smaller window size is more appropriate.

We can now see the tradeoff discussed at the beginning of the section: informally, increasing the window size is more likely to make patient sellers (those with high δ) truthful. On the other hand, it is less likely to make high quality sellers (those with high q) truthful. When q is high and the window is large, the seller is likely to have a high score regardless of what actions he takes when he receives a low value item, because most items are high quality. This makes a smaller window more desirable, because it magnifies the impact of the seller's actions in those periods where he has a low value item for sale.

By choosing a window size T, the system designer determines *how much* and for *how long* the seller's future scores decrease if he does not describe his current item truthfully. In particular, if the seller exaggerates his description in the current period, then his score will decrease by $1/T$ in each of the next T periods (relative to being truthful). We observe the following tradeoff between the intensity and the duration of this score reduction: the intensity is decreasing in T, while the duration is increasing in T. The optimal value of T will depend on the available information on the seller's attributes and the premium function.

As δ increases the duration effect becomes more important, so the optimal window size increases. On the other hand, as q increases the intensity effect becomes more important, so the optimal window size decreases.

Finally, we observe that this tradeoff is also faced by a mechanism designer who knows *neither* q nor δ: a choice must be made regarding the incentives provided to patient sellers and those provided to high quality sellers.

2.4.3 Window Aggregation Mechanism with Imperfect Monitoring

Throughout this chapter we are focusing on perfect monitoring, i.e., we are assuming that the seller receives a rating which accurately reflects his action after every transaction. However, various studies have shown that monitoring may be imperfect in practice [e.g., 23, 18, 14]. In this section we relax the assumption on perfect monitoring in two ways. First, we consider the case of missing feedback, where after some transactions the buyer does not rate the seller. Second, we consider the case where ratings may not always reflect the seller's action. For both cases we show the same tradeoff as in the perfect monitoring case: a larger window size tends to better incentivize patient sellers, but does not incentivize high quality sellers as well.

Missing Feedback

Let p_{av} be the probability that the seller receives no rating when his action is $a \in \{t, u\}$ for being truthful and untruthful respectively, and the true value of the item is $v \in \{v_H, v_L\}$. If the seller receives a rating, we assume that it accurately reflects his action: he receives a good rating (of value 1) for describing his item truthfully, and a bad rating (of value 0) for exaggerating the value of a low value item in his description. We assume that p_{tH} is not significantly larger than p_{uH}, so that it is optimal for the seller to describe a high value item truthfully.

Let $V(\vec{r})$ be the maximum infinite horizon discounted value when the current vector of ratings is $\vec{r}$; then:

$$
\begin{aligned}
V(\vec{r}) =& q_H \left(v_H \cdot b_T \left(\sum_{i=1}^{T} r_i / T \right) + \delta((1-p_{tH}) V(1,\vec{r}) + p_{tH} V(\vec{r})) \right) \\
&+ (1-q_H) \max \left\{ v_H \cdot b_T \left(\sum_{i=1}^{T} r_i / T \right) + \delta((1-p_{uL}) V(0,\vec{r}) + p_{uL} V(\vec{r})), \right. \\
&\left. v_L \cdot b_T \left(\sum_{i=1}^{T} r_i / T \right) + \delta((1-p_{tL}) V(1,\vec{r}) + p_{tL} V(\vec{r})) \right\} \qquad (2.14)
\end{aligned}
$$

In particular, with probability q_H the seller has a high value item for sale, which he describes truthfully. The immediate payment he receives is $v_H \cdot b_T(\sum_{i=1}^{T} r_i/T)$; with probability $1-p_{tH}$ he receives a good rating and with probability p_{tH}^n he receives no rating. With probability $1-q_H$ the seller has a low value item for sale. If he describes it as a high value item, his payoff is $v_H \cdot b_T(\sum_{i=1}^{T} r_i/T) + \delta((1-p_{uL})V(0,\vec{r}) + p_{uL}V(\vec{r}))$, since he receives $v_H \cdot b_T(\sum_{i=1}^{T} r_i/T)$ now, but his ratings "decrease" to $(0,\vec{r})$ with probability $1-p_{uL}$ and remain the same with probability p_{uL}. If he describes the item truthfully, he receives a lower payment now, but his ratings "increase" to $(1,\vec{r})$ with probability $1-p_{tL}$. The seller will choose the description that maximizes his payoff.

Let $p \equiv q_H \cdot p_{tH} + (1-q_H) \cdot p_{tL}$ be the ex ante probability (before the value of the item is known) that the seller receives no rating if he is truthful.

Let

$$
F_m(T,\delta) = \frac{1-p_{uL}}{v_H - v_L} \frac{b_T(1) - b_T(1-1/T)}{b_T(1)} \sum_{i=0}^{\infty} \delta^{i+1} \sum_{j=0}^{\min(T-1,i)} \binom{i}{j} (1-p)^j p^{i-j}
$$

We use the subscript m to denote that we are considering the possibility of missing feedback. The function F_m is increasing in δ.

Lemma 5 *If b_T is logarithmically concave and $p_{uL} \geq p_{tL}$, then it is optimal for the seller to be always truthful if and only if*

$$
q \cdot F_m(T,\delta) \geq 1. \qquad (2.15)
$$

Lemma 5 reduces the problem of finding whether it is optimal for the seller to be always truthful to checking whether the following inequality holds.

$$(v_H - v_L)b_T(1) \leq (1-p_{uL})q(b_T(1) - b_T(1-1/T))\sum_{i=0}^{\infty}\delta^{i+1}\sum_{j=0}^{\min(T-1,i)}\binom{i}{j}(1-p)^j p^{i-j}$$

This condition requires that the seller does not deviate from being truthful when his current score is equal to 1. The seller does not deviate from being truthful if the discounted future gains for being truthful are greater (in expectation) than the current gains for being untruthful. The current payment to the seller increases by $(v_H - v_L)b_T(1)$ if the seller deviates from being truthful. On the other hand, the future payments to the seller decrease by $b_T(1) - b_T(1-1/T)$ in every period that is affected by the current rating. This is the case until the seller receives T new ratings. The seller receives exactly j new ratings in i periods with probability $(1-p)^j p^{i-j}$. Therefore, the probability that the seller has not received T new ratings in i periods is $\sum_{j=0}^{\min(T-1,i)}\binom{i}{j}(1-p)^j p^{i-j}$.

Let $T_m^*(\delta) = \arg\max_T\{F_m(T,\delta)\}$.

Proposition 6 $T_m^*(\delta)$ *is increasing in* δ *in the following sense: for* $\delta \geq \delta'$*,*

(i) $\max\{T : T \in T_m^*(\delta)\} \geq \max\{T : T \in T_m^*(\delta')\}$*; and*

(ii) $\min\{T : T \in T_m^*(\delta)\} \geq \min\{T : T \in T_m^*(\delta')\}$.

As in Section 2.4.2, we conclude in the case of missing feedback (under the assumptions of Lemma 5) when p is fixed: (1) if δ is known and the goal is to maximize the range of q for which the seller is always truthful, then the set of optimal windows is increasing in δ; (2) if q is known and the goal is to maximize the range of δ for which the seller is always truthful, then the set of optimal windows is decreasing in q. We note that we are assuming that p remains fixed as δ and q change. Since $p = q_H \cdot p_{tH} + (1-q_H)\cdot p_{tL}$ and $q = q_H \cdot v_H + (1-q_H)\cdot v_L$, we can assume that q_H is fixed and q changes through v_H and v_L. Alternatively, we can assume that q_H is changing, but also p_{tH} and p_{tL} change accordingly so that p is fixed.

We conclude by discussing two cases that can be viewed as special cases of (2.14). First, consider the setting where the item that the seller has for sale is not always sold. If the probability that the item is not sold depends on the description that the seller posts, but not on his score, then we can use (2.14) and still interpret $b_T(s)$ as the expected premium to the seller.

Second, we note that our current model aggregates ratings from rounds in which the seller has a low value item in exactly the same way as ratings from rounds in which the seller has a high value item. This modeling decision was motivated by the way ratings are aggregated in online markets. However, the seller faces no moral hazard for describing a high value item, and it may be reasonable to assume that the seller is not given a rating for advertising a high value item truthfully. Then we can assume that $p_{tH} = 1$ in (2.14) and still conclude that the optimal window is increasing in δ and decreasing in q.

Inaccurate Ratings

The rating that the seller receives may not reflect his action, because the buyer may make a mistake. In this section we identify conditions under which in the presence of inaccurate ratings the optimal window is increasing in δ and decreasing in q.

We model inaccurate ratings by assuming that with probability p_{av} the seller receives the wrong rating when his action is $a \in \{t, u\}$ for being truthful and untruthful respectively, and the true value of the item is $v \in \{v_H, v_L\}$. We assume that the probability p_{tH} is sufficiently small so that the seller always describes a high value item truthfully. Let $V(\vec{r})$ be the maximum infinite horizon discounted payoff when the current vector of ratings is $\vec{r}$. Then, the optimization problem of the seller is given by the following dynamic program.

$$
\begin{aligned}
V(\vec{r}) = & q_H \left(v_H \cdot b_T \left(\sum_{i=1}^{T} r_i / T \right) + \delta \left((1 - p_{tH}) V(1, \vec{r}) + p_{tH} V(0, \vec{r}) \right) \right) \\
& + (1 - q_H) \max \left\{ v_H \cdot b_T \left(\sum_{i=1}^{T} r_i / T \right) + \delta \left((1 - p_{uL}) V(0, \vec{r}) + p_{uL} V(1, \vec{r}) \right), \right. \\
& \left. \qquad v_L \cdot b_T \left(\sum_{i=1}^{T} r_i / T \right) + \delta \left((1 - p_{tL}) V(1, \vec{r}) + p_{tL} V(0, \vec{r}) \right) \right\}
\end{aligned}
$$

In particular, if the seller describes a low value item as a high value item, he receives a bad rating with probability $1 - p_{uL}$ and a good rating with probability p_{uL}. On the other hand, if the seller has a low value item for sale and describes it truthfully, then he receives an immediate payment of $v_L \cdot b_T\left(\sum_{i=1}^{T} r_i/T\right)$ in expectation, and gets a good rating with probability $1-p_{tL}$; with probability p_{tL} he gets a bad rating despite the fact that he described the item truthfully.

Let $p \equiv q_H \cdot p_{tH} + (1-q_H)\cdot p_{tL}$ be the ex ante probability (before the value of the item is known) that the seller receives a negative rating if he is truthful. Let

$$F_w(T,\delta) = \frac{1-p_{uL}-p_{tL}}{v_H-v_L}\sum_{i=0}^{T-1}\delta^{i+1}\sum_{k=0}^{i}\binom{i}{k}p^k(1-p)^{i-k}\frac{b_T(1-k/T)-b_T(1-(k+1)/T)}{b_T(1)}$$

We use the subscript w to denote that we are considering that possibility of wrong feedback, i.e., that the rating may not accurately reflect the seller's action. We observe that $F_w(T,\delta)$ is increasing in δ.

Lemma 6 *Suppose that one of the following conditions holds.*

(i) b_T is concave; or

(ii) b_T is logarithmically linear, i.e., $b_T(s) = e^{\alpha\cdot s+\beta}$ with $\alpha > 0$; or

(iii) b_T is strictly logarithmically concave and p is sufficiently small.

Then, it is optimal for the seller to be truthful at all $\vec{r}$ if and only if:

$$q \cdot F_w(T,\delta) \geq 1. \tag{2.16}$$

Lemma 6 reduces the problem of finding whether it is optimal for the seller to be always truthful to checking whether the following inequality holds.

$$(v_H-v_L)b_T(1) \leq (1-p_{uL}-p_{tL})q\sum_{i=0}^{T-1}\delta^{i+1}\sum_{k=0}^{i}\binom{i}{k}p^k(1-p)^{i-k}\frac{b_T(1-k/T)-b_T(1-(k+1)/T)}{b_T(1)}$$

This inequality checks whether the seller would deviate from being truthful if his current score is equal to 1 and he has a low value item for sale. The seller does not deviate from being truthful if the discounted future gains for being truthful are greater (in expectation) than the current gains for being untruthful. If the seller deviates from being truthful by describing a low value item as a high value item, then his current payment increases by

$(v_H - v_L)b_T(1)$. If the seller gets a good rating now and is truthful in future periods, then in i periods from now his score is $1-k/T$ with probability $\binom{i}{k}p^k(1-p)^{i-k}$. On the other hand, if the seller gets a bad rating now and is truthful in future periods, then in i periods from now his score is $1-(k+1)/T$ with probability $\binom{i}{k}p^k(1-p)^{i-k}$. The right hand side considers the difference in the expected payments in the next T periods, and discounts appropriately.

We conjecture that the condition that p must be sufficiently small in condition (iii) of Lemma 6 can be weakened. Numerical experiments with specific premium functions suggest that in the result of Lemma 6 holds for any value of p under condition (iii). However, technical verification of this fact in general remains an open problem.

In order to conclude that the optimal window is increasing in δ and decreasing in q, we will derive conditions under which the set

$$T_w^*(\delta) = \arg\max_T \{F_w(T,\delta)\}$$

is increasing in δ. This is done in the following Proposition.

Proposition 7 *Let*

$$h_{T,T'}(k) = \frac{b_{T'}(1-k/T')}{b_{T'}(1)} - \frac{b_T(1-k/T)}{b_T(1)}.$$

If for every $T' > T$ there exists a $k_0 \in \{0,\ldots,T\}$ such that $h_{T,T'}(k) \le h_{T,T'}(k+1)$ for $k < k_0$ and $h_{T,T'}(k) \ge h_{T,T'}(k+1)$ for $k > k_0$, then $T_w^*(\delta)$ is increasing in δ in the following sense: for $\delta \ge \delta'$,

(i) $\max\{T : T \in T_w^*(\delta)\} \ge \max\{T : T \in T_w^*(\delta')\}$*; and*

(ii) $\min\{T : T \in T_w^*(\delta)\} \ge \min\{T : T \in T_w^*(\delta')\}$.

We conclude (as in Section 2.4.2) that if the premium function satisfies one of the first two conditions of Lemma 6 and the condition of Proposition 7, then the optimal window is increasing in δ and decreasing in q (when p is fixed). Moreover, if condition (iii) of Lemma 6 and the condition of Proposition 7 hold, then for any $\bar{\delta} < 1$ and for sufficiently small p the optimal window is increasing in δ in the interval $[0,\bar{\delta}]$.[5] Similarly, if q is known and the

[5]For condition (iii) of Lemma 6, the upper bound on p is some increasing function of T, say $u(T)$.

goal is to maximize the range of δ for which the seller is always truthful, for every range of q and for sufficiently small p the optimal window is decreasing in q.

We have seen in Section 2.4 that requiring the premium function to be logarithmically concave is not a particularly strong condition. Moreover, the condition of Proposition 7 is satisfied by many premium functions. As an example, consider premium functions of the form $b_T(s) = \alpha(T) \cdot b(s) + \gamma(T)$, where $\alpha(T)$ is nondecreasing in T and $\gamma(T)$ is nonincreasing in T. This form captures the following intuition: *as the window size increases, buyers trust the information that is aggregated in the seller's score more*. For instance, if the seller has the maximum possible score ($s = 1$), we expect that the premium will increase as T increases; this is captured by the assumption that $\alpha(T)$ is increasing. On the other hand, if the seller has the minimum possible score ($s = 0$), we expect the premium to decrease as T increases; this is captured by the assumption that $\gamma(T)$ is nonincreasing. Simple calculations show that the condition of Proposition 7 is satisfied for various functions of this form; e.g., if $b(s) = s^n$ or $b(s) = e^{ns}$, and $\alpha(\cdot)$, $\gamma(\cdot)$ are arbitrary functions of T. We note that many empirical studies use regression forms that correspond to premia of the form $\alpha(T) \cdot b(s) + \gamma(T)$, where $b(s) = e^{ns}$ [e.g., 16, 48, 31].

The following Corollary restricts attention to premia that do not explicitly depend on the window size.

Corollary 1 *Suppose $b_T(\cdot) \equiv b(\cdot)$. If $b'(s)$ is logarithmically concave, then $T^*_w(\delta)$ is increasing.*

Examples of functions with a logarithmically concave derivative are $b(s) = s^n$, $b(s) = e^s$ and the logistic function $b(s) = 1/(1 + e^{a-bs})$ for $b > 0$. We note that the conclusion of Corollary 1 holds more generally if $b'(1-y) - yb''(1-y) < 0$ for some $y \in [0,1]$ implies that $b'(1-z) - zb''(1-z) < 0$ for $z > y$.

We conclude by summarizing the results of this section. In the case of inaccurate ratings, if any of the assumptions of Lemma 6 and the assumption of Proposition 7 hold, then assuming that p is fixed: (1) if δ is known and the goal is to maximize the range of q for which the seller is always truthful, then the set of optimal windows is increasing in δ; (2) if

Consider $\bar{\delta} < 1$. If b_T satisfies the condition of Proposition 7, then the optimal window is increasing for $\delta \in [0, \bar{\delta}]$ if $p \leq u(\max\{T^*_w(\bar{\delta})\})$.

q is known and the goal is to maximize the range of δ for which the seller is always truthful, then the set of optimal windows is decreasing in q.

2.4.4 Weighted Aggregation Mechanism

In the previous section we studied the problem of finding the best mechanism in the class of Window Aggregation Mechanisms. In this section we consider the optimization within the general class of Weighted Aggregation Mechanisms; i.e., we let W denote the set of all weight vectors such that $\sum_i w_i = 1$, and $w_i \geq 0$ for all i, with $w_0 \geq w_1 \geq w_2 \cdots$.

In this section we make the following assumption.

Assumption 2 *The premium function is $b(\cdot)$, and does not depend on the weight vector $\vec{w}$. Further, it is strictly convex and logarithmically concave.*

Note that we assume that the premium depends on the weights of the aggregation mechanism only through the score. Since Assumption 2 implies Assumption 1, we conclude (as before) by Lemma 3 that it is optimal for the seller to be always truthful if and only if condition (2.7) holds.

Applying the formulation of Section 2.4.1, it can be easily shown that finding the optimal weight vector $\vec{w}^*(\delta)$ in (2.9) is equivalent to solving the following optimization problem.[6]

$$\text{minimize} \sum_{i=0}^{\infty} \delta^i \cdot b(1 - w_i) \tag{2.17}$$

$$\text{subject to} \sum_{i=0}^{\infty} w_i = 1; \tag{2.18}$$

$$w_i \geq 0, \forall i. \tag{2.19}$$

By Assumption 2, the premium function is strictly convex and thus (2.17)-(2.19) has a unique solution.

For the general Weighted Aggregation Mechanism it may seem plausible that all past

[6] It is straightforward to show that at the optimal solution the weights will be nonincreasing, and thus we do not need to explicitly include the constraints $w_i \geq w_j$ for $i < j$.

transactions of the seller should be assigned strictly positive weights in order to best incentivize truthfulness. The following lemma shows that this is *not* the case under Assumption 2: it is optimal to only include a *finite* number of ratings in the seller's score. In particular, we note that this result implies that the Exponential Aggregation Mechanism introduced in Example 4 never arises as an optimal weight vector.

Lemma 7 *If the premium function satisfies Assumption 2 and* $\delta \in (0,1)$*, then only finitely many weights are positive in* $\vec{w}^*(\delta)$.

In light of Lemma 7, let $N^*(\delta)$ be the number of strictly positive weights at the optimal solution; i.e., $w_i^*(\delta) > 0$ for $i = 0,\ldots,N^*(\delta)-1$, and $w_i^*(\delta) = 0$ for $i \geq N^*(\delta)$. The following proposition characterizes $\vec{w}^*(\delta)$ and $N^*(\delta)$.

Proposition 8 *If Assumption 2 holds, then*

(i) $w_i^*(\delta)$ *is strictly decreasing in i for* $i < N^*(\delta)$*;*

(ii) For sufficiently small i, $w_i^*(\delta)$ *is decreasing in* δ*; and*

(iii) $N^*(\delta)$ *is nondecreasing in* δ.

According to (i), the optimal weights are strictly decreasing in i for $i < N^*(\delta)$, and thus each strictly positive weight is different at the optimal solution. This implies that the Window Aggregation Mechanism, which has weights $w_i = 1/T$ for $i = 0,\ldots,T-1$, does not arise as an optimal weight vector. The preceding observation is particularly interesting when one considers that nearly all online marketplaces that use mechanisms which only weight recent ratings, tend to use a Window Aggregation Mechanism to do so; for example, eBay's current mechanism weights all ratings equally over the past 12 months of feedback. Our insight suggests that a more robust mechanism might be one that weights more recent ratings even more heavily than older ratings.

Following Section 2.4.1, we now consider the optimal choice of weights depending on the available information regarding δ and q.

1. *Both* δ *and q are known by the mechanism designer.* As in Section 2.4.1, the vector of weights $\vec{w}^*(\delta)$ is optimal is this case, and whether or not it will be possible to make

the seller always truthful depends on whether q and δ are large enough to ensure (2.7) is satisfied.

2. *The mechanism designer knows* δ*, but not* q*.* As in Section 2.4.1, the weights $\vec{w}^*(\delta)$ maximize the range of values of q for which the seller will be truthful. From Proposition 8 (ii), as δ increases, the most recent transactions are weighted less. Informally, this happens because the value to the seller of the distant future increases relative to the near future. This effect implies that the aggregation mechanism should weight distant past ratings more heavily relative to recent ratings when δ increases. Moreover, Proposition 8 (iii) states that *the optimal number of strictly positive weights is nondecreasing in* δ. As in the case of the Window Aggregation Mechanism, the intuition for this is that sellers with higher discount factors care more about the future, and so truthfulness is better incentivized when information from more transactions is included.

3. *The mechanism designer knows* q*, but not* δ*.* As in Section 2.4.1, let $\delta^*(q)$ be the solution to (2.10)-(2.12). Then the weights $\vec{w}^*(\delta^*(q))$ maximize the range of values of δ for which the seller will be truthful. Recall that $\delta^*(q)$ is decreasing in q; thus, from Proposition 8 (ii), as q increases, the most recent transactions are weighted more. Moreover, from Proposition 8 (iii), *the optimal number of strictly positive weights is decreasing in* q. As before, this is an intuitive result, since as q increases, it is possible to make less patient sellers truthful, and such sellers value the future less.

We observe that the dependence of the number of strictly positive weights on δ and q is similar to the corresponding results on the optimal Window Aggregation Mechanism, i.e., that the optimal window size is nondecreasing in δ, and nonincreasing in q. Again, this tradeoff directly affects a mechanism designer who knows neither δ nor q.

2.5 Non-Binary Values and Ratings

Throughout this chapter we assumed that the value of the item is either high or low, and that ratings are either good or bad (binary rating system). In this section we discuss how the modeling and the results shown can be generalized when we relax this assumption. Thus, our results are robust even if the seller has potentially many "levels" of dishonesty possible.

We assume that the possible values of the item are $\Omega = \{0, 1/Q, 2/Q, \ldots, 1\}$ for some integer $Q \geq 1$, instead of only $\{\text{low}, \text{high}\}$. Let p_v be the probability that the seller has an item of value v for sale, so $\sum_{v \in \Omega} p_v = 1$.[7] Note that we allow the possibility $p_v = 0$ for all but one value in Ω; in this case, the seller has the same item in every period, but may be tempted to exaggerate its value in his advertisement.

Let $v_a \cdot b_{\vec{w}}(s)$ be the expected payment to the seller when he advertises an item as having value v_a and his score is s. Thus the highest quality item is worth $b_{\vec{w}}(s)$, and the expected payment to the seller is scaled in proportion to the advertised quality. We assume that $b_{\vec{w}}(s)$ is an increasing function of s.

Further, we assume that when the advertised value is v_a and the true value v, then the seller receives rating $r_i = 1 - (v_a - v)^+$, where $x^+ \equiv \max(x, 0)$ denotes the positive part of x. That is, the buyer "penalizes" the seller by the difference between the advertised and the true value, whenever this difference is positive.[8] Since the payment $v_a \cdot b_{\vec{w}}(s)$ is increasing in v_a and s, the seller will only ever consider exaggerating the value of his advertisement, and we can safely assume that $r_i = 1 - (v_a - v)$. We consider the Weighted Aggregation Mechanism: the seller's score that the buyers see is $s(\vec{r}) = \sum_j w_j r_j$, where $w_i \geq 0$, and $w_0 \geq w_1 \geq \ldots$, with $\sum_i w_i = 1$.

Let $s_i(\vec{r})$ be the seller's score in i periods if his current ratings are $\vec{r}$ and he is truthful in all other periods. Our main insight is the following lemma.

Lemma 8 *If $b_{\vec{w}}$ is convex, then it is optimal for the seller to be always truthful if and only if*

$$b_{\vec{w}}(s(\vec{r})) \leq \left(\sum_{v' \in \Omega} p_{v'} \cdot v' \right) \sum_{i=0}^{\infty} \delta^{i+1} \cdot \left(b_{\vec{w}}(s_i(1, \vec{r})) - b_{\vec{w}}(s_i(0, \vec{r})) \right) \tag{2.20}$$

for all rating vectors $\vec{r}$.

In words, Lemma 8 says that if the premium function is convex, then the seller is more

[7] It is straightforward to extend the result to continuous probability distributions [8].

[8] The same results hold if $r_i = 1 - |v_a - v|$ instead. In this case the buyer "penalizes" the seller by the absolute difference between the advertised and true value. However, it seems more realistic and fair to only "penalize" the seller when he exaggerates the value of the item in his advertisement.

tempted to post an advertisement that is significantly higher than the true value. The intuition for this is that a small lie now is associated with a small gain now, but a relatively large reduction in future payments (because of convexity). If the payment is convex and the seller considers exaggerating the value of the item in his description, he is better off significantly exaggerating it. Then, the reduction in future payments is not that large relative to the current gain. When the premium function is concave, the reverse occurs: small lies are more beneficial for the seller. This is because under concavity a small lie only imposes a small reduction in future payments (relative to current gains).

Inspired by the preceding result, define:

$$q \equiv \sum_{v' \in \Omega} p_{v'} \cdot v',$$

to be the seller's *quality*. As a result, *condition (2.20) is equivalent to (2.5)*. From this observation, we immediately obtain analogs of most results of the paper.

If we assume that $b_{\vec{w}}$ is *strictly convex and logarithmically concave*, then we get analogs of the design results of Section 2.4. In particular, the optimal weights are strictly decreasing whenever positive, and the number of strictly positive weights is finite. Thus, neither the Window Aggregation Mechanism nor the Exponential Aggregation Mechanism is optimal. Moreover, the tradeoff discussed for the binary case still holds: informally, increasing the number of strictly positive weights is more likely to make patient sellers (those with high δ) truthful, while it is less likely to make high quality sellers (those with high q) truthful. Finally, we note that the results on the design of the Window Aggregation Mechanism with perfect monitoring can also be extended to the setting with non-binary values and ratings [8].

2.6 Proofs for Chapter 2

Proof of Proposition 1: We will show that there exists a sufficiently large N such that if the seller is truthful up to that time, then he is better off advertising a low value item as a high value item at time N. We assume that the seller is truthful up to time N and so his score is $(s_P, s_T) = (N, N)$. Substituting in (2.2), the seller will advertise a low value item truthfully at reputation (N,N) only if

$$(v_H - v_L)b(N,N) \le \delta(V(N+1,N+1) - V(N,N+1)). \tag{2.21}$$

We wish to upper bound $V(N+1,N+1) - V(N,N+1)$. Let $\alpha_M = b(M,M) - b(M-1,M)$. By (2.3), $\alpha_M \to 0$ as $M \to \infty$. Choose $k^*(N) \in \arg\max_{k\ge 0} \alpha_{N+k+1}$. Then,

$$V(N+1,N+1) - V(N,N+1) \le \sum_{k=0}^{\infty} \delta^k \alpha_{N+k+1} \le \frac{1}{1-\delta}\alpha_{N+k^*(N)+1} \to 0 \text{ as } N \to \infty.$$

Thus, for every $\varepsilon > 0$ there exists N_1 such that $V(N+1,N+1) - V(N,N+1) < \varepsilon$ for all $N \ge N_1$. Since $b(N,N)$ is assumed to be bounded away from zero as $N \to \infty$, there exist $\varepsilon > 0$ and N_2 such that $(v_H - vL)b(N,N) > \varepsilon$ for all $N \ge N_2$. Thus, there exists N for which (2.21) is not satisfied, which implies that there exists an s_T at which it is optimal for the seller to falsely advertise a low value item.

■

Proof of Lemma 2:

We first prove (i). Fix a vector of ratings $\vec{r}$. Note that if $b_{\vec{w}}(s(\vec{r})) = 0$, then the condition trivially holds. We thus assume that $b_{\vec{w}}(s(\vec{r})) > 0$ in what follows.

We first observe that if $w_i > 0$, then $s_i(1,\vec{r}) > w_i$. In particular, we could only have $s_i(1,\vec{r}) = w_i > 0$ if $i = 0$ and $w_0 = 1$, which contradicts the assumption that $w_0 < 1$.

If $w_i > 0$, then

$$b_{\vec{w}}(s_i(1,\vec{r})) - b_{\vec{w}}(s_i(0,\vec{r})) = b_{\vec{w}}(s_i(1,\vec{r})) - b_{\vec{w}}(s_i(1,\vec{r}) - w_i) > w_i \cdot \frac{b_{\vec{w}}(s_i(1,\vec{r}))}{s_i(1,\vec{r})} \ge w_i \cdot b_{\vec{w}}(s(\vec{r})).$$

The equality holds because $s_i(0,\vec{r}) = s_i(1,\vec{r}) - w_i$. The first inequality follows because $b_{\vec{w}}$ is strictly convex and $b_{\vec{w}}(0) = 0$, and the second because $s(\vec{r}) \le s_i(1,\vec{r}) \le 1$. We conclude that $b_{\vec{w}}(s_i(1,\vec{r})) - b_{\vec{w}}(s_i(0,\vec{r})) > w_i \cdot b_{\vec{w}}(s(\vec{r}))$ for all i with $w_i > 0$. Summing over all i and using the fact that $\sum_i w_i = 1$, we show (i).

We now show (ii). By concavity of $b_{\vec{w}}$,

$$\sum_{i=0}^{\infty} (b_{\vec{w}}(1) - b_{\vec{w}}(1 - w_i)) \le \sum_{i=0}^{\infty} w_i b_{\vec{w}}(1) = b_{\vec{w}}(1),$$

since $b_{\vec{w}}(0) \geq 0$, and $\sum_i w_i = 1$. $\blacksquare$

Proof of Corollary 2:

We first prove (i) by showing that there is no profitable deviation from the policy of always advertising truthfully when $b_{\vec{w}}$ is strictly convex and q and δ are sufficiently large. Fix a vector of ratings $\vec{r}$. Note that if $b_{\vec{w}}(s(\vec{r})) = 0$, then the seller has no reason to deviate from the truthful policy. We thus assume that $b_{\vec{w}}(s(\vec{r})) > 0$ in what follows.

We first observe that if $w_i > 0$, then $s_i^1(\vec{r}) > w_i$. In particular, we could only have $s_i^1(\vec{r}) = w_i > 0$ if $i = 0$ and $w_0 = 1$. But this can not happen under the assumption $w_1 > 0$, since $w_i \geq 0$ and $\sum_i w_i = 1$.

If $w_i > 0$, then

$$b_{\vec{w}}(s_i^1(\vec{r})) - b_{\vec{w}}(s_i^0(\vec{r})) = b_{\vec{w}}(s_i^1(\vec{r})) - b_{\vec{w}}(s_i^1(\vec{r}) - w_i) > w_i \cdot \frac{b_{\vec{w}}(s_i^1(\vec{r}))}{s_i^1(\vec{r})} \geq w_i \cdot b_{\vec{w}}(s(\vec{r})).$$

The equality holds because $s_i^0(\vec{r}) = s_i^1(\vec{r}) - w_i$. The first inequality follows because $b_{\vec{w}}$ is strictly convex and $b_{\vec{w}}(0) = 0$, and the second because $s(\vec{r}) \leq s_i^1(\vec{r}) \leq 1$. We conclude that $b_{\vec{w}}(s_i^1(\vec{r})) - b_{\vec{w}}(s_i^0(\vec{r})) > w_i \cdot b_{\vec{w}}(s(\vec{r}))$ for all i with $w_i > 0$. Let

$$\varepsilon = \frac{1}{q\delta} \sum_{i:w_i>0} \delta^i \left(\frac{b_{\vec{w}}(s_i^1(\vec{r})) - b_{\vec{w}}(s_i^0(\vec{r}))}{b_{\vec{w}}(s(\vec{r}))} - w_i \right)$$

The seller will always be truthful if δ and q are such that,

$$b_{\vec{w}}(s(\vec{r})) \leq q\delta \sum_i \delta^i w_i b_{\vec{w}}(s(\vec{r})) + \varepsilon b_{\vec{w}}(s(\vec{r})),$$

(so that (2.5) holds), or equivalently

$$1 - \varepsilon \leq q\delta \sum_i \delta^i w_i. \tag{2.22}$$

We next show that for any $\varepsilon > 0$, we can choose $\delta < 1$ and $q < 1$ such that (2.22) holds. We consider the following cases.

(a) There exists k such that $w_i = 0$ for $i > k$. Then it suffices to choose $\delta < 1$ and $q < 1$ such that $q\delta^{k+1} \geq 1 - \varepsilon$.

(b) $w_i > 0$ for all i. Then there exists a finite k such that $\sum_{i=0}^{k} w_i \geq 1-\varepsilon/2$. By choosing $\delta, q < 1$ such that $\delta^{k+1} \geq (1-\varepsilon)/(1-\varepsilon/2)$, we have

$$q\delta \sum_{i} \delta^i w_i \geq q\delta \sum_{i=0}^{k} \delta^i w_i \geq \frac{1-\varepsilon}{1-\varepsilon/2} \sum_{i=0}^{k} w_i \geq 1-\varepsilon.$$

If (2.22) holds for given $q < 1$ and $\delta < 1$, then it also holds for all q' and δ' with $q' \geq q$ and $\delta' \geq \delta$. This concludes the proof of (i).

We now show (ii) by demonstrating that condition (2.5) is not satisfied at $\vec{r} = \vec{1}$ when $b_{\vec{w}}$ is concave. At $\vec{r} = \vec{1}$, condition (2.5) is $b_{\vec{w}}(1) \leq \delta q \sum_{i=0}^{\infty} \delta^i (b_{\vec{w}}(1) - b_{\vec{w}}(1-w_i))$. By concavity of $b_{\vec{w}}$, we upper bound the right hand side

$$\delta q \sum_{i=0}^{\infty} \delta^i (b_{\vec{w}}(1) - b_{\vec{w}}(1-w_i)) \leq \delta q \sum_{i=0}^{\infty} \delta^i w_i b_{\vec{w}}(1) < b_{\vec{w}}(1),$$

since $q, \delta < 1$, $b_{\vec{w}}(0) \geq 0$, and $\sum_i w_i = 1$. Thus, it is optimal for the seller to deviate from always being truthful. This shows that always advertising truthfully is not optimal for the seller under decreasing returns to reputation. ■

Proof of Proposition 3: We note that if $s(\vec{r}) < s^*$, then the seller has no reason to deviate from the truthful policy. We thus consider what happens for $s(\vec{r}) \geq s^*$.

We first show that if $s^* \leq 1 - w_1$, then the seller is better off deviating from the truthful policy at $s(\vec{r}) = 1$. If he receives a low value item and does not deviate, his infinite horizon expected payoff is $v_H q_H \sum_{i=1}^{\infty} \delta^i b_{\vec{w}}(1)$, while if he does deviate his expected payoff is at least $v_H(1 + q_H \sum_{i=2}^{\infty} \delta^i) b_{\vec{w}}(1)$, which is strictly greater for any $\delta < 1$. This shows (i).

We now show (ii). If $s^* > 1 - w_1$, then $w_0 \geq w_1 > 1 - s^*$. We will show that (2.5) is satisfied for any vector of ratings $\vec{r}$ with $s(\vec{r}) \geq s^*$. We note that $b_{\vec{w}}(s(\vec{r})) = b_{\vec{w}}(s(1,\vec{r})) = b_{\vec{w}}(1)$. Moreover, $s(0,\vec{r}) = s_i(1,\vec{r}) - w_i$, so $s_0(0,\vec{r}), s_1(0,\vec{r}) < s^*$ and $b_{\vec{w}}(s_0(0,\vec{r})) = b_{\vec{w}}(s_1(0,\vec{r})) = 0$. Thus,

$$\delta q_H \sum_{i=0}^{\infty} \delta^i (b_{\vec{w}}(s_i(1,\vec{r})) - b_{\vec{w}}(s_i(0,\vec{r}))) \geq \delta q_H \sum_{i=0}^{1} \delta^i (b_{\vec{w}}(s_i(1,\vec{r})) - b_{\vec{w}}(s_i(0,\vec{r}))) = q_H(\delta + \delta^2) b_{\vec{w}}(1),$$

and thus if $q_H(\delta + \delta^2) \geq 1$, the seller does not deviate.

$\blacksquare$

Proof of Lemma 4:

By Lemma 1, it is optimal for the seller to always be truthful if $q \geq q^*(b_{\vec{w}}, \delta, \vec{r})$ for all $\vec{r}$, where $q^*(b_{\vec{w}}, \delta, \vec{r}) \equiv (v_H - v_L) b_{\vec{w}}(s(\vec{r})) / \sum_{i=0}^{\infty} \delta^{i+1}(b_{\vec{w}}(s_i(1, \vec{r})) - b_{\vec{w}}(s_i(0, \vec{r})))$. If it is not optimal for the seller to deviate from the truthful policy at maximum reputation, then $q \geq q^*(b_{\vec{w}}, \delta, \vec{1})$.

Let $b_{\vec{w}}$ be logarithmically concave. To show the statement of the proposition, it suffices to show that $q^*(b_{\vec{w}}, \delta, \vec{1}) \geq q^*(b_{\vec{w}}, \delta, \vec{r})$ for all $\vec{r}$, i.e., if for all $\vec{r}$, $b_{\vec{w}}(1) \cdot \sum_{i=0}^{\infty} \delta^i (b_{\vec{w}}(s_i(1, \vec{r})) - b_{\vec{w}}(s_i(0, \vec{r}))) \geq b_{\vec{w}}(s(\vec{r})) \cdot \sum_{i=0}^{\infty} \delta^i (b_{\vec{w}}(1) - b_{\vec{w}}(1 - w_i))$.

It suffices to show that

$$b_{\vec{w}}(1) \cdot (b_{\vec{w}}(s_i^1(\vec{r})) - b_{\vec{w}}(s_i^0(\vec{r}))) \geq b_{\vec{w}}(s(\vec{r})) \cdot (b_{\vec{w}}(1) - b_{\vec{w}}(1 - w_i)). \tag{2.23}$$

To conclude the proof we show that (2.23) is satisfied by a logarithmically concave function.

Let $w > 0$ and $x \in [w, 1]$. Since $b_{\vec{w}}$ is logarithmically concave, then $\log(b_{\vec{w}}(x)) - \log(b_{\vec{w}}(x - w))$ is nonincreasing in x, which implies that $b_{\vec{w}}(x - w)/b_{\vec{w}}(x)$ is nondecreasing in x. Thus, $b_{\vec{w}}(1 - w_i)/b_{\vec{w}}(1) \geq b_{\vec{w}}(s_i(0, \vec{r}))/b_{\vec{w}}(s_i(1, \vec{r}))$ and $(b_{\vec{w}}(1) - b_{\vec{w}}(1 - w_i))/b_{\vec{w}}(1) \leq (b_{\vec{w}}(s_i(1, \vec{r})) - b_{\vec{w}}(s_i(0, \vec{r})))/b_{\vec{w}}(s_i(1, \vec{r}))$. Using the fact that $b_{\vec{w}}(s_i(1, \vec{r})) \geq b_{\vec{w}}(s_i(0, \vec{r}))$, this implies that (2.23) holds and concludes the proof. $\blacksquare$

Proof of Proposition 4: If the seller is always truthful under $\tilde{b}_{\vec{w}}$, it must be that $(v_H - v_L)\tilde{b}_{\vec{w}}(1) \leq \delta q \sum_{i=0}^{\infty} \delta^i (\tilde{b}_{\vec{w}}(1) - \tilde{b}_{\vec{w}}(1 - w_i))$, since in particular he is truthful when his score is equal to 1. Since $b_{\vec{w}}(s) \leq \tilde{b}_{\vec{w}}(s)$ for all s, and $b_{\vec{w}}(1) = \tilde{b}_{\vec{w}}(1)$, we have $\tilde{b}_{\vec{w}}(1) - \tilde{b}_{\vec{w}}(1 - w_i) \leq b_{\vec{w}}(1) - b_{\vec{w}}(1 - w_i)$ for all w_i. Thus,

$$(v_H - v_L) b_{\vec{w}}(1) = (v_H - v_L)\tilde{b}_{\vec{w}}(1) \leq \delta q \sum_{i=0}^{\infty} \delta^i (\tilde{b}_{\vec{w}}(1) - \tilde{b}_{\vec{w}}(1 - w_i)) \leq \delta q \sum_{i=0}^{\infty} \delta^i (b_{\vec{w}}(1) - b_{\vec{w}}(1 - w_i)).$$

Since $b_{\vec{w}}$ satisfies Assumption 1 and condition (2.6), the optimal policy of the seller is to always advertise truthfully under $b_{\vec{w}}$ (by Lemma 3). $\blacksquare$

Proof of Proposition 5: We will show that $\log(F(T, \delta))$ satisfies increasing differences.

Let $T' \geq T$.

$$\log(F(T',\delta)) - \log(F(T,\delta)) = \log\left(\frac{1-\delta^{T'}}{1-\delta^{T}}\right) + t(T,T'),$$

where $t(T,T')$ does not depend on δ. Thus, to show that $\log(F(T,\delta))$ has increasing differences in (T,δ) it suffices to show that $(1-\delta^{T'})/(1-\delta^{T})$ is increasing in δ. The first derivative with respect to δ is positive if and only if

$$\frac{T\cdot\delta^{T-1}}{1-\delta^{T}} \geq \frac{T'\cdot\delta^{T'-1}}{1-\delta^{T'}}.$$

Since $T' \geq T$ it suffices to show that $r(x) \equiv (x\cdot\delta^{x-1})/(1-\delta^{x})$ is decreasing. We proceed by differentiating r:

$$r'(x) = \frac{\delta^{x-1}}{(1-\delta^{x})^2}(1-\delta^{x} - x\ln(1/\delta)).$$

To complete the proof, we show that $\delta^{T} + T\ln(1/\delta) > 1$ holds for $T \geq 1$, $\delta \in (0,1)$. First note that $\delta^{T} + T\ln(1/\delta)$ is increasing in T, since

$$\frac{\partial(\delta^{T} + T\ln(1/\delta))}{\partial T} = \ln(1/\delta)\cdot(1-\delta^{T}) > 0.$$

So it suffices to show that $\hat{g}(\delta) \equiv \delta + \ln(1/\delta) > 1$. g is strictly decreasing in $(0,1)$, because

$$\hat{g}'(\delta) = 1 + \frac{-1/\delta^2}{1/\delta} = \frac{\delta-1}{\delta} < 0,$$

and $\hat{g}(1) = 1$. So, $\hat{g}(\delta) > 1$ for $\delta \in (0,1)$.

This proves that $\log(F(T,\delta))$ has increasing differences in (T,δ); the result follows by applying Topkis' Theorem [68]. ■

Proof of Lemma 5: It is optimal for the seller to be always truthful if and only if any one step deviation from the truthful policy (i.e., the policy of always advertising items truthfully) does not yield a higher payoff. The seller may consider to exaggerate an item in his description if in that period he has a low value item for sale. Let $\hat{V}(\vec{r})$ be the expected infinite horizon discounted payoff to the seller if he is always truthful and his current vector of ratings is $\vec{r}$. The seller does not deviate from the truthful policy at $\vec{r}$ if and only if

$$(v_H - v_L) \cdot b_T \left(\sum_{i=1}^{T} r_i / T \right) \le \delta[(1 - p_{tL})\hat{V}(1, \vec{r}) - (1 - p_{uL})\hat{V}(0, \vec{r}) + (p_{tL} - p_{uL})\hat{V}(\vec{r})]. \tag{2.24}$$

With probability $a(i, j) \equiv \binom{i}{j}(1-p)^j p^{i-j}$ exactly j new ratings arrive in i periods. Let $s_j(\vec{r})$ be the score of the seller after j good rating arrive with initial rating vector $\vec{r}$. Then $s_0(\vec{r}) = \sum_{i=1}^{T} r_i/T$, and

$$\begin{aligned}
&(1 - p_{tL})\hat{V}(1, \vec{r}) - (1 - p_{uL})\hat{V}(0, \vec{r}) + (p_{tL} - p_{uL})\hat{V}(\vec{r}) = \\
&q \cdot \sum_{i=0}^{\infty} \delta^i \sum_{j=0}^{\min(T-1,i)} a(i, j)[(1 - p_{tL})b_T(s_j(1, \vec{r})) - (1 - p_{uL})b_T(s_j(0, \vec{r})) + (p_{tL} - p_{uL})b_T(s_j(\vec{r}))] = \\
&q \cdot \sum_{i=0}^{\infty} \delta^i \sum_{j=0}^{\min(T-1,i)} a(i, j)[(1 - p_{tL})(b_T(s_j(1, \vec{r})) - b_T(s_j(\vec{r}))) + (1 - p_{uL})(b_T(s_j(\vec{r})) - b_T(s_j(0, \vec{r})))]
\end{aligned}$$

This shows that $q \cdot F_m(T, \delta) \ge 1$ is equivalent to (2.24) for $\vec{r} = \vec{1}$, and thus it is a necessary condition for the seller to be always truthful. To show sufficiency, it suffices to show that for $j < T$,

$$\begin{aligned}
&\frac{(1 - p_{tL})(b_T(s_j(1, \vec{r})) - b_T(s_j(\vec{r}))) + (1 - p_{uL})(b_T(s_j(\vec{r})) - b_T(s_j(0, \vec{r})))}{b_T(s_0(\vec{r}))} \ge \\
&\frac{(1 - p_{uL})(b_T(1) - b_T(1 - 1/T))}{b_T(1)}
\end{aligned}$$

Note that $s_j(1, \vec{r}) \ge s_j(\vec{r}) \ge s_j(0, \vec{r})$ and $s_j(1, \vec{r}) = s_j(0, \vec{r}) + 1/T$. Therefore, either $s_j(1, \vec{r}) = s_j(\vec{r}) + 1/T$ or $s_j(\vec{r}) = s_j(0, \vec{r}) + 1/T$.

We first assume that $s_j(1,\vec{r}) = s_j(\vec{r}) + 1/T$ and $s_j(\vec{r}) = s_j(0,\vec{r})$. Then

$$\begin{aligned}
&\frac{(1-p_{tL})(b_T(s_j(1,\vec{r})) - b_T(s_j(\vec{r}))) + (1-p_{uL})(b_T(s_j(\vec{r})) - b_T(s_j(0,\vec{r})))}{b_T(s_0(\vec{r}))} = \\
&(1-p_{tL})\frac{b_T(s_j(1,\vec{r})) - b_T(s_j(\vec{r}))}{b_T(s_0(\vec{r}))} \geq \\
&(1-p_{tL})\frac{b_T(1) - b_T(1-1/T)}{b_T(1)} \geq \\
&(1-p_{uL})\frac{b_T(1) - b_T(1-1/T)}{b_T(1)}
\end{aligned}$$

This first inequality is a consequence of the fact that b_T is logarithmically concave; and the second inequality holds because $p_{uL} \geq p_{tL}$.

We next assume that $s_j(1,\vec{r}) = s_j(\vec{r})$ and $s_j(\vec{r}) = s_j(0,\vec{r}) + 1/T$. Then

$$\begin{aligned}
&\frac{(1-p_{tL})(b_T(s_j(1,\vec{r})) - b_T(s_j(\vec{r}))) + (1-p_{uL})(b_T(s_j(\vec{r})) - b_T(s_j(0,\vec{r})))}{b_T(s_0(\vec{r}))} = \\
&(1-p_{uL})\frac{b_T(s_j(\vec{r})) - b_T(s_j(0,\vec{r}))}{b_T(s_0(\vec{r}))} \geq \\
&(1-p_{uL})\frac{b_T(1) - b_T(1-1/T)}{b_T(1)}
\end{aligned}$$

This again holds because b_T is logarithmically concave.

We conclude that if $p_{uL} \geq p_{tL}$, then the seller is always truthful if and only if $q \cdot F_m(T,\delta) \geq 1$. ■

Proof of Proposition 6: Let

$$g(T,\delta) \equiv \frac{b_T(1) - b_T(1-1/T)}{b_T(1)} \sum_{i=0}^{\infty} f(T,i)\delta^i,$$

$$f(T,i) = \sum_{j=0}^{\min(T-1,i)} \binom{i}{j} (1-p)^j p^{i-j}.$$

Also let $\alpha(i,j) \equiv \binom{i}{j}(1-p)^j p^{i-j}$. Clearly, $T_m^*(\delta) = \arg\max_T\{g(T,\delta)\}$.

We will show that g satisfies the single crossing property in (T,δ), i.e., that $g(T',\delta) > g(T,\delta)$ implies $g(T',\delta') > g(T,\delta')$ for $\delta' > \delta$ and $T' > T$. This will imply that $T^*(\delta)$ is increasing in δ [51]. Let

$$c \equiv \frac{(b_{T'}(1) - b_{T'}(1 - 1/T'))/b_{T'}(1)}{(b_T(1) - b_T(1 - 1/T))/b_T(1)}.$$

Equivalently we will show that if

$$\sum_{i=0}^{\infty} (c \cdot f(T', i))\delta^i > \sum_{i=0}^{\infty} f(T, i)\delta^i,$$

then the inequality also holds for $\delta' > \delta$.

We first show that if $c \cdot f(T', i) > f(T, i)$, then $c \cdot f(T', i+1) > f(T, i+1)$. We consider the cases $i < T$ and $i \geq T$ separately.

Suppose $i < T$. Then $cf(T', i) > f(T, i)$ implies that

$$c \sum_{j=0}^{i} a(i, j) > \sum_{j=0}^{i} a(i, j),$$

which can only happen if $c > 1$. Also $\min(i+1, T'-1) = i+1$, while $\min(i+1, T-1)$ is $i+1$ if $i < T-1$; and i if $i = T-1$. In either case, $c \cdot f(T', i+1) > f(T, i+1)$.

Now suppose that $i \geq T$, and let $k \equiv \min(i, T'-1)$. Assume $c \cdot f(T', i) - f(T, i) > 0$. Then

$$\sum_{j=T}^{k} a(i, j) > \frac{1-c}{c} \sum_{j=0}^{T-1} a(i, j).$$

We observe that

$$a(i+1, j) = (1-p)\frac{i+1}{i+1-j} a(i, j) = (1-p)\left(1 + \frac{j}{i+1-j}\right) a(i, j).$$

Then,

$$\begin{aligned}
&\sum_{j=T}^{k} a(i+1, j) - \frac{1-c}{c} \sum_{j=0}^{T-1} a(i+1, j) \\
&= (1-p) \sum_{j=T}^{k} \left(1 + \frac{j}{i+1-j}\right) a(i, j) - (1-p)\frac{1-c}{c} \sum_{j=0}^{T-1} \left(1 + \frac{j}{i+1-j}\right) a(i, j)
\end{aligned}$$

Moreover,

$$\sum_{j=T}^{k}\frac{i+1}{i+1-j}a(i,j)\geq\frac{i+1}{i+1-T}\sum_{j=T}^{k}a(i,j)>\frac{1-c}{c}\frac{i+1}{i+1-T}\sum_{j=0}^{T-1}a(i,j)\geq\frac{1-c}{c}\sum_{j=0}^{T-1}\frac{i+1}{i+1-j}a(i,j).$$

Since

$$\sum_{j=T}^{k}a(i,j)-\frac{1-c}{c}\sum_{j=0}^{T-1}a(i,j)>0,$$

we have that

$$\sum_{j=T}^{k}\frac{i+1}{i+1-j}a(i,j)\geq\frac{i+1}{i+1-T}\sum_{j=T}^{k}a(i,j)>\frac{1-c}{c}\frac{i+1}{i+1-T}\sum_{j=0}^{T-1}a(i,j)\geq\frac{1-c}{c}\sum_{j=0}^{T-1}\frac{i+1}{i+1-j}a(i,j)$$

We conclude that if $c\cdot f(T',i)-f(T,i)>0$, then

$$\sum_{j=T}^{k}a(i+1,j)>\frac{1-c}{c}\sum_{j=0}^{T-1}a(i+1,j).$$

Since $a(i+1,j)\geq 0$, $\min(i+1,T'-1)\geq k$ and $\min(i+1,T-1)=T-1$, this implies that

$$c\cdot f(T',i+1)-f(T,i+1)>0.$$

The final step of the proof is to show that if $\sum_{i=0}^{\infty}(c\cdot f(T',i))\delta^i>\sum_{i=0}^{\infty}f(T,i)\delta^i$, then the inequality also holds for $\delta'>\delta$. Let $T'>T$ and $e_i=c\cdot f(T',i)-f(T,i)$. We have shown that if $e_i>0$ then $e_{i+1}>0$. If $e_i>0$ for all i, then trivially $\sum_{i=0}^{\infty}e_i\delta^i>0$ for all δ.

Now suppose $e_0<0$ and let $k=\max\{i:e_i<0\}$. If $\sum_{i=0}^{\infty}e_i\delta^i>0$, then

$$\sum_{i=k+1}^{\infty}|e_i|\delta^i>\sum_{i=1}^{k}|e_i|\delta^i.$$

Moreover,

$$\sum_{i=k+1}^{\infty}i|e_i|\delta^{i-1}>\sum_{i=1}^{k}i|e_i|\delta^{i-1}.$$

The last inequality is equivalent to

$$\frac{\partial}{\partial\delta}\left(\sum_{i=0}^{\infty} e_i\delta^i\right) = \sum_{i=1}^{\infty} ie_i\delta^{i-1} > 0,$$

which concludes the proof. ■

Proof of Lemma 6: It is optimal for the seller to describe a low value item truthfully at $\vec{r}$ if

$$(v_H - v_L)\cdot b_T\left(\sum_{i=1}^{T} r_i/T\right) \leq \delta(1-p_{uL}-p_{tL})(V(1,\vec{r})-V(0,\vec{r})).$$

It is optimal for the seller to be always truthful if and only if any one step deviation from the truthful policy (i.e., the policy of always advertising items truthfully) does not yield a higher payoff. The seller may consider to exaggerate an item in his description if in that period he has a low value item for sale. Let $\hat{V}(\vec{r})$ be the infinite horizon discounted payoff to the seller if he is always truthful. The seller will not deviate from being truthful when his ratings are $\vec{r}$ if

$$(v_H - v_L)\cdot b_T\left(\sum_{j=1}^{T} r_j/T\right) \leq \delta(1-p_{uL}-p_{tL})(\hat{V}(1,\vec{r})-\hat{V}(0,\vec{r})).$$

We observe that the payments to the seller from $\hat{V}(1,\vec{r})$ and $\hat{V}(0,\vec{r})$ may differ in the next T periods, but not after that. Let $s_i(\vec{r})$ be the seller's score in i periods if his current rating vector is $\vec{r}$ and he gets good ratings in all future periods. Then,

$$\hat{V}(1,\vec{r})-\hat{V}(0,\vec{r}) = q\sum_{k=0}^{i}\binom{i}{k}p^k(1-p)^{i-k}\left(b_T(s_i(1,\vec{r})-k/T)-b_T(s_i(1,\vec{r})-(k+1)/T)\right)$$

In particular, after i periods the seller gets k bad ratings (which are inaccurate) with probability $a(i,k)\equiv\binom{i}{k}p^k(1-p)^{i-k}$. We conclude that the seller will not deviate from being truthful when his ratings are $\vec{r}$ if

$$(v_H-v_L)\cdot b_T(s_0(\vec{r})) \leq \delta(1-p_{uL}-p_{tL})q\sum_{k=0}^{i}a(i,k)\left(b_T(s_i(1,\vec{r})-k/T)-b_T(s_i(1,\vec{r})-(k+1)/T)\right).$$

To prove the Lemma it suffices to show that

$$\sum_{k=0}^{i} a(i,k)\left(\frac{b_T(s_i(1,\vec{r})-k/T)-b_T(s_i(1,\vec{r})-(k+1)/T)}{b_T(s_0(\vec{r}))}-\frac{b_T(1-k/T)-b_T(1-(k+1)/T)}{b_T(1)}\right)\geq 0. \tag{2.25}$$

This implies that if the seller does not deviate from being truthful at $\vec{1}$, then he does not deviate at any $\vec{r}$. The remainder of the proof shows that if (i), (ii) or (iii) is satisfied, then (2.25) holds.

We first consider condition (i) and assume that b_T is concave. Then,

$$b_T(s_i(1,\vec{r})-k/T)-b_T(s_i(1,\vec{r})-(k+1)/T)\geq b_T(1-k/T)-b_T(1-(k+1)/T)$$

by the concavity of b_T, and

$$b_T(1)\geq b_T\left(\sum_{j=1}^{T} r_j/T\right)$$

since b_T is increasing. We conclude that in this case

$$\left(\frac{b_T(s_i(1,\vec{r})-k/T)-b_T(s_i(1,\vec{r})-(k+1)/T)}{b_T\left(\sum_{i=1}^{T} r_i/T\right)}-\frac{b_T(1-k/T)-b_T(1-(k+1)/T)}{b_T(1)}\right)\geq 0,$$

and thus (2.25) holds.

Now assume that $b_T(s)=e^{\alpha\cdot s+\beta}$ and $\alpha>0$ (condition (ii)). Then

$$\begin{aligned}
&\frac{b_T(s_i(1,\vec{r})-k/T)-b_T(s_i(1,\vec{r})-(k+1)/T)}{b_T\left(\sum_{i=1}^{T} r_i/T\right)}-\frac{b_T(1-k/T)-b_T(1-(k+1)/T)}{b_T(1)}=\\
&\frac{e^{\alpha(s_i(1,\vec{r})-k/T)+\beta}-e^{\alpha(s_i(1,\vec{r})-(k+1)/T)+\beta}}{e^{\alpha\left(\sum_{j=1}^{T} r_j/T\right)+\beta}}-\frac{e^{\alpha(1-k/T)+\beta}-e^{\alpha(1-(k+1)/T)+\beta}}{e^{\alpha+\beta}}=\\
&\frac{e^{\alpha(s_i(1,\vec{r})-k/T)}-e^{\alpha(s_i(1,\vec{r})-(k+1)/T)}}{e^{\alpha\left(\sum_{j=1}^{T} r_j/T\right)}}-\frac{e^{\alpha(1-k/T)}-e^{\alpha(1-(k+1)/T)}}{e^{\alpha}}=\\
&\left(\frac{e^{\alpha(s_i(1,\vec{r}))}}{e^{\alpha\left(\sum_{j=1}^{T} r_j/T\right)}}-1\right)(e^{-\alpha k/T}-e^{-\alpha(k+1)/T})\geq 0
\end{aligned}$$

because $s_i(1,\vec{r})\geq\sum_{j=1}^{T} r_j/T$ and $\alpha>0$.

We now show (iii). We will show that if b_T is strictly logarithmically concave, then

(2.25) holds for sufficiently small p. If b_T is strictly log-concave, then $\log(b_T(x)) - \log(b_T(x-w))$ is strictly decreasing in x. Thus $b_T(x-w)/b_T(x)$ is strictly increasing in x and $b_T(1-1/T)/b_T(1) > b_T(x-1/T)/b_T(x)$. Let

$$c(T) \equiv \frac{b_T(1-1/T)}{b_T(1)} - \frac{b_T(1-2/T)}{b_T(1-1/T)}.$$

Then

$$\frac{b_T(1-1/T)}{b_T(1)} - \frac{b_T(x-1/T)}{b_T(x)} \geq c(T)$$

for $x \in \{1/T, 2/T, ..., (T-1)/T\}$, and $c(T) > 0$. Moreover, there exists λ such that

$$\frac{b_T(x-k/T) - b_T(x-(k+1)/T)}{b_T(x)} - \frac{b_T(1-k/T) - b_T(1-(k+1)/T)}{b_T(1)} \geq -\lambda$$

for $x \in \{1/T, 2/T, ..., (T-1)/T\}$ and $k \in \{1, 2, ..., Tx-1\}$. For instance, the aforementioned inequality holds for any premium function b_T if $\lambda = 2$, since

$$\frac{b_T(x-(k+1)/T)}{b_T(x)} + \frac{b_T(1-k/T)}{b_T(1)} \leq \frac{b_T(x)}{b_T(x)} + \frac{b_T(1)}{b_T(1)} = 2.$$

We conclude that if

$$p \leq 1 - \left(\frac{\lambda}{c(T)+\lambda}\right)^{1/T}$$

then condition (2.25) holds. ■

Proof of Proposition 7: Let

$$c(i) = \sum_{k=0}^{i} \binom{i}{k} p^k (1-p)^{i-k} f(k),$$

$$f(k) = \frac{b_T(1-k/T) - b_T(1-1/T-k/T)}{b_T(1)} - \frac{b_{T'}(1-k/T') - b_{T'}(1-1/T'-k/T')}{b_{T'}(1)},$$

$$g(T,\delta) = \sum_{i=0}^{T-1} \delta^i \sum_{k=0}^{i} \binom{i}{k} p^k (1-p)^{i-k} \frac{b_T(1-k/T) - b_T(1-(k+1)/T)}{b_T(1)}.$$

Clearly, $T_w^*(\delta) = \arg\max_T \{g(T,\delta)\}$.

The proof consists of three steps. First, we show that if $h_{T,T'}$ satisfies the assumption of the proposition, then $f(k) < 0$ implies $f(k+1) < 0$. The second step is to show that $c(i) < 0$ implies that $c(i+1) < 0$. Then we show that g satisfies the single crossing property in (T, δ) and conclude that $T_w^*(\delta)$ is increasing.

Step 1: Let

$$h(k) = \frac{b_{T'}(1 - k/T')}{b_{T'}(1)} - \frac{b_T(1 - k/T)}{b_T(1)}$$

We observe that

$$f(k) = \frac{b_{T'}(1-(k+1)/T')}{b_{T'}(1)} - \frac{b_T(1-(k+1)/T)}{b_T(1)} - \frac{b_{T'}(1-k/T')}{b_{T'}(1)} + \frac{b_T(1-k/T)}{b_T(1)} = h(k+1) - h(k).$$

Since there exists a $k_0 \in \{0, ..., T\}$ such that h_k is increasing in k for $k < k_0$ and decreasing in k for $k \geq k_0$, we conclude that $f_k < 0$ if and only if $k \geq k_0$. Thus, if $f(k) < 0$ then $f(k+1) < 0$.

Step 2: Let

$$a(i,k) = \binom{i}{k} p^k (1-p)^{i-k}.$$

The key property we exploit is that:

$$\frac{a(i+1,k)}{a(i,k)} = (1-p)\frac{i+1}{i+1-k}$$

is strictly increasing in k. We have shown that $f(k) \geq 0$ for $k < k_0$ and $f(k) < 0$ for $k \geq k_0$. Suppose $c(i) < 0$. Of course, in this case we must have $i \geq k_0$. Then:

$$\sum_{k=0}^{k_0-1} a(i,k) f(k) < -\sum_{k=k_0}^{i} a(i,k) f(k).$$

But now note that for all $k < k_0$, $i+1-k > i+1-k_0$; and for all k such that $k_0 \leq k < i$,

$i+1-k \leq i+1-k_0$. So we get:

$$\begin{aligned}\sum_{k=0}^{k_0-1} a(i+1,k)f(k) &= (i+1)(1-p)\sum_{k=0}^{k_0-1} a(i,k)f(k)/(i+1-k)\\ &< (i+1)(1-p)\sum_{k=0}^{k_0-1} a(i,k)f(k)/(i+1-k_0)\\ &< -(i+1)(1-p)\sum_{k=k_0}^{i} a(i,k)f(k)/(i+1-k_0)\\ &< -(i+1)(1-p)\sum_{k=k_0}^{i} a(i,k)f(k)/(i+1-k)\\ &< -\sum_{k=k_0}^{i+1} a(i+1,k)f(k)\end{aligned}$$

where the last inequality follows since $f(i+1) < 0$. We conclude that $c(i+1) < 0$, as required.

Step 3: Let $T' > T$ and $\delta' > \delta$. The function g satisfies the single crossing property in (T,δ) if $g(T',\delta) > g(T,\delta)$ implies that $g(T',\delta') > g(T,\delta')$. We observe that

$$g(T',x) - g(T,x) =$$
$$\sum_{i=T}^{T'-1}\delta^i\sum_{k=0}^{i}\binom{i}{k}p^k(1-p)^{i-k}\frac{b_T(1-k/T')-b_T(1-(k+1)/T')}{b_{T'}(1)} - \sum_{i=0}^{T-1}\delta^i\sum_{k=0}^{i}\binom{i}{k}p^k(1-p)^{i-k}f(k) =$$
$$\sum_{i=T}^{T'-1}\delta^i\sum_{k=0}^{i}\binom{i}{k}p^k(1-p)^{i-k}\frac{b_T(1-k/T)-b_T(1-(k+1)/T)}{b_T(1)} - \sum_{i=0}^{T-1}\delta^i c(i)$$

According to Step 2, there exists some i_0 such that we can rewrite the previous difference as

$$g(T',x) - g(T,x) = -\sum_{i=0}^{i_0-1}x^i d_i + \sum_{i=i_0}^{T'-1}x^i d_i,$$

where $d_i \geq 0$ for all i. Assume that $g(T',\delta) - g(T,\delta) > 0$. Then,

$$\sum_{i=0}^{i_0-1} i\delta^{i-1}d_i = \sum_{i=0}^{i_0-1}\frac{i}{\delta}\delta^i d_i \leq \frac{i_0-1}{\delta}\sum_{i=0}^{i_0-1}\delta^i d_i \leq \frac{i_0-1}{\delta}\sum_{i=i_0}^{T'-1}\delta^i d_i = \sum_{i=T}^{T'-1}(i_0-1)\delta^{i-1}d_i \leq \sum_{i=i_0}^{T'-1} i\delta^{i-1}d_i.$$

This implies that if $g(T',\delta)-g(T,\delta)>0$, then $g'(T',\delta)-g'(T,\delta)\geq 0$. We conclude that if $g(T',\delta)-g(T,\delta)>0$, then $g(T',\delta')-g(T,\delta')>0$ for $\delta'>\delta$. This shows that the objective satisfies the single crossing property. Thus, we apply Theorem 4 from [51] to conclude the proof. ■

Proof of Corollary 1: We will first show a stronger result: if $b'(1-y)-yb''(1-y)<0$ for some $y\in[0,1]$ implies that $b'(1-z)-zb''(1-z)<0$ for $z>y$, then $T_w^*(\delta)$ is increasing in δ. Then, we will show that this is the case if $\log(b'(s))$ is concave.

Let

$$g(x)\equiv b(1)\cdot h'_{T',T}(x)=\frac{1}{T}b'(1-x/T)-\frac{1}{T'}b'(1-x/T').$$

Clearly, $g(0)>0$. It suffices to show that if $g(x)<0$ then $g(x')<0$ for $x'>x$ in order to apply Proposition 7.

$$\begin{aligned}
g(x) &= \frac{1}{T}b'(1-x/T)-\frac{1}{T'}b'(1-x/T') \\
&= \int_{1/T'}^{1/T}\frac{\partial}{\partial y}[yb'(1-yx)]dy \\
&= \int_{1/T'}^{1/T}[b'(1-yx)-yxb''(1-yx)]dy \\
&= \frac{1}{x}\int_{x/T'}^{x/T}[b'(1-z)-zb''(1-z)]dz
\end{aligned}$$

If $b'(1-y)-yb''(1-y)>0$ for all $y\in[0,1]$ then $g(x)>0$ for all $x\in[0,T]$. Otherwise there exists $z_0\in(0,1]$ such that $b'(1-z_0)-z_0b''(1-z_0)=0$ and $b'(1-z)-zb''(1-z)>0$ for $z<z_0$; $b'(1-z)-zb''(1-z)<0$ for $z>z_0$.

If $g(x)<0$, then

$$\int_{x/T'}^{z_0}[b'(1-z)-zb''(1-z)]dz<\int_{z_0}^{x/T}|b'(1-z)-zb''(1-z)|dz.$$

Let $x'>x$. Then $x'/T'>x/T'$ and $x'/T>x/T$. If $x'/T'>z_0$, then $g(x')<0$. If $x'/T'<z_0$, then

$$\int_{x/T'}^{z_0}[b'(1-z)-zb''(1-z)]dz>\int_{x'/T'}^{z_0}[b'(1-z)-zb''(1-z)]dz$$

and

$$\int_{z_0}^{x/T} |b'(1-z) - zb''(1-z)|dz < \int_{z_0}^{x'/T} |b'(1-z) - zb''(1-z)|dz.$$

We conclude that

$$\int_{x'/T'}^{z_0} [b'(1-z) - zb''(1-z)]dz < \int_{z_0}^{x'/T} |b'(1-z) - zb''(1-z)|dz$$

which implies that $g(x') < 0$.

We have shown the result for the case that $b'(1-y) - yb''(1-y) < 0$ implies that $b'(1-z) - zb''(1-z) < 0$ for $z > y$. A sufficient condition for this is that $(1-x)b''(x)/b'(x)$ is decreasing. The latter holds if b' is logarithmically concave. ■

Proof of Lemma 7:

We first show that if b is logarithmically concave, then $b'(1) < \infty$. Suppose not, and let $f(x) = \log(b(x))$. Then f is concave and $f'(x) \to \infty$ as $x \to 1$. However, this is not possible since if f is concave, then f' is decreasing on $[0,1]$.

We next show that if $b'(1) < \infty$, then at the optimal solution only a finite number of weights are positive. Consider the vector $\vec{u}^*(\delta)$, which is the solution of (2.17)-(2.19). The optimality conditions are:

$$b'(1-u_i^*) = \frac{\lambda - \mu_i}{\delta^i};\ u_i^* \cdot \mu_i = 0,\ \text{for all } i,$$

where λ and μ_i are the Lagrange multipliers of constraint (2.18) and (2.19) respectively. Suppose $u_i^* > 0$ for all i. Then $b'(1-u_i^*) = \lambda/\delta^i$ for some $\lambda > 0$. But $b'(1-u_i^*) < b'(1) < \infty$ for all i, while $\lambda/\delta^i \to \infty$ as $i \to \infty$, which is a contradiction. Thus it must be that $u_i^* > 0$ for only a finite number of i's, i.e., $N^*(\delta) < \infty$.

■

Proof of Proposition 8: We first consider $\vec{u}^*$, i.e., the solution of (2.17)-(2.19). Let λ be the Lagrange multiplier of (2.18). The optimality conditions imply that for $i = 0, \ldots, N^*(\delta) - 1$, $u_i^*(\delta) = 1 - b'^{-1}(\lambda/\delta^i)$. Since b is strictly convex, both b' and b'^{-1} are strictly increasing. Thus the optimal weights are strictly decreasing for $i = 0, \ldots, N^*(\delta) - 1$. For $\vec{v}^*$, we observe that the optimal weights are strictly decreasing for $i = 0, \ldots, N^*(\delta^*(q)) - 1$ (according to the proof of Lemma 7). This proves (i).

Substituting the formula for $u_i^*(\delta)$ in (2.18),

$$\sum_{i=0}^{N^*(\delta)-1} b'^{-1}(\lambda/\delta^i) = N^*(\delta) - 1. \tag{2.26}$$

In order to find the optimal weights, it suffices to find $\lambda, N^*(\delta)$ such that equation (2.26) is satisfied and

$$\frac{\lambda}{\delta^{N^*(\delta)-1}} < b'(1) \leq \frac{\lambda}{\delta^{N^*(\delta)}}, \tag{2.27}$$

so that $u_{N^*-1}^*(\delta) > 0$ and $u_{N^*}^*(\delta) = 0$.

Suppose $\delta_1 < \delta_2$. Let N_1 and N_2 be the corresponding optimal numbers of strictly positive weights, and λ_1 and λ_2 be the corresponding Lagrange multipliers. By (2.26),

$$\sum_{i=0}^{N_1-1} b'^{-1}(\lambda_1/\delta_1^i) = N_1 - 1;$$

$$\sum_{i=0}^{N_2-1} b'^{-1}(\lambda_2/\delta_2^i) = N_2 - 1.$$

We will show that $N_1 \leq N_2$. Suppose not. Then, $N_1 > N_2$, and subtracting the latter from the former equation,

$$\sum_{i=N_2}^{N_1-1} b'^{-1}(\lambda_1/\delta_1^i) + \sum_{i=0}^{N_2-1} \left(b'^{-1}(\lambda_1/\delta_1^i) - b'^{-1}(\lambda_2/\delta_2^i)\right) = N_1 - N_2.$$

Since $b'^{-1}(\lambda_1/\delta_1^i) < 1$ for $i \in \{N_2, ..., N_1-1\}$, it follows that $\sum_{i=N_2}^{N_1-1} b'^{-1}(\lambda_1/\delta_1^i) < N_1 - N_2$. Thus the previous equality can only hold if $b'^{-1}(\lambda_1/\delta_1^i) > b'^{-1}(\lambda_2/\delta_2^i)$ for some $i \in \{0, 1, ..., N_2 - 1\}$. Since b is strictly convex, this implies that $\lambda_1/\delta_1^i > \lambda_2/\delta_2^i$ for some $i \in \{0, 1, ..., N_2 - 1\}$. We observe that if this is the case for i, it is also the case for $i+1$ (since $1/\delta_1 > 1/\delta_2$). We conclude that it must hold for $i = N_2 - 1$, i.e., $\lambda_1/\delta_1^{N_2-1} > \lambda_2/\delta_2^{N_2-1}$. Thus $b'^{-1}(\lambda_1/\delta_1^{N_2}) > b'^{-1}(\lambda_2/\delta_2^{N_2})$, which is a contradiction, since $N_1 > N_2$ implies $b'^{-1}(\lambda_1/\delta_1^{N_2}) < 1$ and $b'^{-1}(\lambda_2/\delta_2^{N_2}) \geq 1$ (by (2.27)). Thus, if $\delta_1 < \delta_2$, it must be $N_1 \leq N_2$. This proves that $N^*(\delta)$ is increasing in δ.

To show that $M^*(q)$ is decreasing in q, we observe that $\delta^*(q)$ (defined in the proof of

Lemma 7) is decreasing in q, and that $M^*(q) \equiv N^*(\delta^*(q))$. This concludes the proof of (ii). ■

Proof of Lemma 8: The Bellman equation for the discounted infinite horizon expected payoff of the seller is then given by:

$$V(\vec{r}) = \sum_{v \in \Omega} p_v \max_{v_a \geq v, v_a \in \Omega} \left\{ v_a b_{\vec{w}}(s(\vec{r})) + \delta \cdot V\left(1 - (v_a - v), \vec{r}\right) \right\}.$$

Similar to the analysis in Section 2.3.1, it is optimal for the seller to always be truthful if and only if

$$(v_a - v) \cdot b_{\vec{w}}(s(\vec{r})) \leq \left(\sum_{v' \in \Omega} p_{v'} \cdot v' \right) \sum_{i=0}^{\infty} \delta^{i+1} \cdot \left(b_{\vec{w}}(s_i(1, \vec{r})) - b\left(s_i(1 - (v_a - v), \vec{r})\right)\right) \quad (2.28)$$

for all rating vectors $\vec{r}$ and all $v_a, v \in \Omega$ with $v_a > v$.

We observe that (2.28) depends only on the difference $v_a - v$ and not on the specific values v_a or v. Moreover, if $b_{\vec{w}}$ is convex, since $(b_{\vec{w}}(s_i(1,\vec{r})) - b_{\vec{w}}(s_i(1-d,\vec{r})))/d = (b_{\vec{w}}(s_i(1,\vec{r})) - b_{\vec{w}}(s_i(1,\vec{r}) - w_i d))/d$ is decreasing in d, we conclude that if (2.28) is satisfied for v_a, v, then it is also satisfied for v'_a, v' with $v'_a - v' \leq v_a - v$. This establishes the lemma. ■

Chapter 3

Market Mechanisms for Peer-to-Peer Systems

In peer-to-peer systems, users share files or resources with each other. By sharing, a peer incurs a cost (because uploading a file consumes network resources), but no direct benefit. Thus, if there is no mechanism that stimulates sharing, a peer has a strong incentive to free ride, i.e., use the resources of other peers without contributing his own. Such behavior is observed in existing peer-to-peer systems; for instance, early data showed that nearly 70 percent of peers of Gnutella were sharing no files, and nearly 50 percent of all responses were returned by the top 1 percent of sharing hosts [1]. A more recent study shows that 85 percent of Gnutella peers share no files [38]. Even worse, according to [1], there were peers in Gnutella who were free riding on the system *despite* sharing files: the files that they were sharing were unpopular, and hence not widely uploaded.

Incentive mechanisms that penalize free riders or reward peers that share have been proposed. In [15] peers enjoy different levels of service according to how much they share their resources, while in [28] free riders are excluded from the system with some probability. In [24], a distributed rating scheme for tackling the free-rider problem is suggested. More general reputation mechanisms, such as those proposed in [33], can be used to obtain a system-wide reputation for each peer. Using this information, each peer will give priority to peers with high reputation.

An alternate approach is to design a system where resource sharing is *required* to be able to use the resources of other peers. This is the case in BitTorrent [19], where peers download pieces of the file and at the same time upload the pieces they already have. Analogously, in [3] peers directly trade resources between themselves.

Another option is to use monetary incentives to solve the problem of free riding. In this case, peers must pay to download files from other peers. The payments may either be in monetary terms (e.g., [32]), or in an internal non-monetary currency. In the latter case, the budget of a peer decreases every time he downloads a file, and increases every time he uploads a file. Such models are considered in [34], [39] and [70]. Friedman et al. study system performance as a function of the total amount of internal currency available [29] .

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

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

With this formulation we can also avoid cheap pseudonyms [30], which are a drawback in most approaches for solving free-riding. Peers cannot benefit by leaving the system and joining with a new identity, since peer performance is determined only by the files uploaded. This naturally introduces a "transaction cost" into the system that prevents peers from taking advantage of multiple identities. Of course, one shortcoming here is that peers who join the system with little content of interest to others may be unable to download anything. One solution is to require such peers to upload a file that is not desired by anyone. The price of this file can be set to be less than the price of any other file. In this way, new peers do not get anything for free, and thus existing peers do not have any incentive to rejoin the system with a different name.

In Section 3.1 we describe the model in more detail. In Section 3.2, we show the existence of a competitive equilibrium: a vector of prices at which demand of each file is equal to the corresponding supply. It is well known that such a vector is Pareto efficient. We derive conditions that guarantee uniqueness of the competitive equilibrium (up to scaling). In Section 3.3, we study the tâtonnement price adjustment process [9], and show that under some assumptions the rate of convergence around the equilibrium is linear in the number of peers. This means that in a large system, the prices will rapidly converge.

A key aspect of our approach consists of a proposal to clear the market even out of equilibrium. In Section 3.4, we propose an allocation mechanism to allocate rates when demand is not equal to supply. We study the Nash equilibria when peers anticipate how their actions affect the resulting allocation, and show that in large peer-to-peer systems, fully strategic behavior by the peers will not ultimately cause large deviations from competitive equilibrium behavior.

Section 3.5 incorporates a complex network structure in the model. We compare various pricing schemes and show that pricing per peer and pricing per file are equivalent in a setting with a trivial network structure, i.e., when all the constraints in the network are imposed by the upload capacities of peers. Moreover, we show that pricing per peer is strictly better when the network structure is not trivial. In Section 3.6 we provide an overview of how we envision market mechanisms to be used in practice, and discuss user incentives.

3.1 Model

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

Assumption 3 *The preference relation of a peer on the set of feasible rate vectors is represented by a continuous utility function $v_i : \Re_+^{|T_i|+1} \to \Re$, which is strictly increasing in each download rate x_{ij}, $j \in T_i$; and strictly decreasing in the upload rate y_i.*

($\Re_+$ denotes the interval $[0, \infty)$.)

We introduce strictly positive prices in the system and consider a particular peer i. Each peer is assumed to have a constraint on the available upload rate; let B_i denote this upper bound for peer i. A rate vector is feasible for a peer as long as the upload rate is at most equal to the peer's upload capacity. We assume that peers do not face any constraint on their download rate; this is consistent with most end peer connections today, where upload capacity is far exceeded by download capacity.[1] Given a vector of prices $\vec{p} \gg 0$ (i.e. $p_j > 0$ for $j \in F$), peer i can find the upload rate y_i and vector of download rates $\vec{x}_i$ that maximize his utility by solving the following optimization problem:

Multilateral Peer Optimization:

$$\text{maximize} \quad v_i(\vec{x}_i, y_i) \tag{3.1}$$

$$\text{subject to} \quad \sum_{j \in T_i} x_{ij} \cdot p_j \leq (\max_{j \in S_i} p_j) \cdot y_i; \tag{3.2}$$

$$y_i \leq B_i; \tag{3.3}$$

$$y_i \geq 0; \ \ x_{ij} \geq 0, \text{ for all } j \in T_i. \tag{3.4}$$

We refer to this exchange as *multilateral*, because it allows peers to trade multilaterally. In Chapter 4 we compare a system that allows multilateral exchange with a system that restricts exchange to being bilateral.

By assumption, the utility function of a peer only depends on his upload rate and not on which files he is uploading. Thus peer i will only choose to upload files that have the highest price among all files in S_i. The constraint (3.2) guarantees that the expenses of a peer are at most equal to his revenue from uploading. The constraint (3.3) guarantees

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

that a peer does not upload at a higher rate than his upload capacity B_i. Finally, all rates must be non-negative (constraint (3.4)). For any price vector $\vec{p} \gg 0$, the feasible region of the Multilateral Peer Optimization problem is compact and by Assumption 3 the objective function is continuous; thus an optimal solution exists for any price vector $\vec{p} \gg 0$. The following lemma captures an important feature of this optimal solution; the proof follows immediately using strict monotonicity of the utility function v_i.

Lemma 9 *If Assumption 3 is satisfied, the budget constraint will bind in the Multilateral Peer Optimization for any price vector $\vec{p} \gg 0$.*

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

Assumption 4 *For every peer $i \in U$, the corresponding Multilateral Peer Optimization problem has a unique solution $(\vec{x}_i, y_i)$ for any price vector $\vec{p} \gg 0$.*

For instance, Assumption 4 is satisfied if each utility function is strictly concave, since the feasible region of the optimization problem of each peer is convex. Let $x_{ij}(\vec{p})$ and $y_i(\vec{p})$ be the optimal values of x_{ij} and y_i respectively when the price vector is $\vec{p} \gg 0$.

We now define exchange economy [50] and relate it to our model. In an *exchange economy* there is a finite number of agents and a finite number of commodities. Each agent is endowed with a bundle of commodities, and has a preference relation on the set of commodity vectors. Given a price vector, each agent finds a vector of commodities to exchange that maximizes his utility. In particular, if $\vec{p}$ is the vector of prices and agent i has endowment $\vec{w}_i$, he sells it at the market and obtains wealth $\vec{p} \cdot \vec{w}_i$. Then the agent buys goods for his consumption at the same price (he may buy back some of the goods he sold).

A straightforward reformulation reveals that our model shares much in common with a standard exchange economy. Consider the constraints of the peer optimization problem (3.1)-(3.4). The constraint $y_i \geq 0$ is implied by the other constraints as long as all prices are non-negative. The remaining constraints can equivalently be written as:

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

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

3.2 Competitive Equilibrium

In this section we define competitive equilibrium. In Section 3.2.1, we then proceed to show that there always exists at least one for the model described in Section 3.1. In Section 3.2.2, we give conditions that guarantee uniqueness.

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

Definition 2 *Given a vector of prices $\vec{p} \gg 0$, a vector $(z_j, j \in F)$ is an* aggregate excess demand vector *if there exist y_{ij}, $i \in U$, $j \in F$, such that:*

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

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

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

Definition 3 *The rate allocation* $(\vec{x}_i^*, i \in U)$ *and* $(y_i^*, i \in U)$ *and the price vector* $\vec{p}^* \gg 0$ *constitute a competitive equilibrium if the following conditions are satisfied:*

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

Note that because of Assumption 3, at competitive equilibrium all prices are strictly positive; otherwise peers would want to download all free files at unboundedly large rates. For this reason, we can restrict competitive equilibria to strictly positive price vectors without loss of generality.

Our goal is to show that a competitive equilibrium exists. *We emphasize that competitive equilibria are desirable because they are Pareto efficient*; this is the content of the first fundamental theorem of welfare economics [50]. However, we do not expect equilibria to exist without any restrictions on the sets S_i and T_i of files being uploaded and downloaded, respectively, by peer i. For example, suppose there is a file that some peers want to download, but no peer has available for upload. Then in general, such a file will have positive demand, while supply will always be zero. Thus the excess demand for such a file will be positive unless its price is sufficiently high. Setting a sufficiently high price is equivalent to considering a system without that file.

To avoid such pathological situations, we introduce a natural diversity assumption. We define the *user graph* as the directed graph $G = (V, E)$ with $V = U$, and $E = \{(i, j) : S_i \cup T_j \neq \emptyset\}$. In other words, G is a graph where nodes correspond to peers. There is a directed edge from peer i to peer j if i has a file that j desires.

Assumption 5 *The user graph consists only of strongly connected components.*

If Assumption 5 is not satisfied, then an equilibrium may not exist. We will therefore assume that Assumption 5 holds.

3.2.1 Existence of Competitive Equilibrium

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

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

The following proposition shows properties of the aggregate excess demand correspondence that are used to prove existence of a competitive equilibrium. The proof is an extension of an argument typically used to prove existence of competitive equilibrium in exchange economies. The key difficulty is in addressing the fact that peers may simultaneously upload multiple files; as discussed in Section 3.1, this feature means our basic model is not a standard exchange economy.

Proposition 9 *If Assumptions 3, 4 and 5 hold, then the aggregate excess demand correspondence $\vec{z}(\cdot)$ defined on $(0, \infty)^F$ satisfies the following properties:*

1. *For every $\vec{p} \gg 0$ and $\vec{z} \in \vec{z}(\vec{p})$, $\vec{p} \cdot \vec{z} = 0$.*

2. *$\vec{z}(\cdot)$ is convex-valued.*

3. $\vec{z}(\cdot)$ *is homogeneous of degree 0.*

4. $\vec{z}(\cdot)$ *is upper hemicontinuous.*

5. *There is an* $s > 0$ *such that* $z_j > -s$ *for any* $\vec{z} \in \vec{z}(\vec{p})$, *for every file* $j \in F$ *and every price vector* $\vec{p} \gg 0$.

6. *If* $\vec{p}^m \to \vec{p} \neq \vec{0}$, $\vec{z}^m \in \vec{z}(\vec{p}^m)$ *and* $p_j = 0$ *for some* j, *then* $\max\{z_j^m : j \in F\} \to \infty$.

Now the existence of a competitive equilibrium follows from standard results in microeconomics; see, e.g., [50], Exercise 17.C.2 .

Theorem 1 *If Assumptions 3, 4 and 5 hold, then there exists a competitive equilibrium.*

Corollary 2 *If the utility function of each peer is strictly concave, and Assumptions 3 and 5 are satisfied, then there exists a competitive equilibrium.*

In Section 3.4, we assume that each peer has a separable utility function, and experiences a cost of uploading that is linear in the upload rate. In this case, the utility function is not strictly concave. The following corollary of Theorem 1 shows existence of a competitive equilibrium for that case.

Corollary 3 *If the utility function of peer* i *is* $v_i(\vec{x}_i, y_i) = u_i(\vec{x}_i) - y_i$, *where* $u_i(\vec{x}_i)$ *is continuous, strictly concave, and strictly increasing in each* x_{ij}, *and Assumption 5 is satisfied, then there exists a competitive equilibrium.*

3.2.2 Uniqueness of Competitive Equilibrium

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

Definition 4 *The function* $\vec{z}(\cdot)$ *has the gross substitutes property if whenever* $\vec{p}' \gg 0$ *and* $\vec{p} \gg 0$ *are such that for some* l, $p'_l > p_l$ *and* $p_k = p'_k$ *for* $k \neq l$, *we have* $z_k(\vec{p}') > z_k(\vec{p})$ *for* $k \neq l$.

If the aggregate excess demand is a function that satisfies the gross substitutes property, then there is at most one competitive equilibrium up to scaling of the price vector [50]. In our model, the aggregate excess demand is a function if and only if each peer is uploading exactly one file, i.e., $|S_i| = 1$ for all $i \in U$. When some peers i have $|S_i| > 1$, the aggregate excess demand is a correspondence, so the preceding result does not apply. In order to adapt that result, we use the following definition.

Definition 5 *The Multilateral Peer Optimization of peer i satisfies the gross substitutes property if whenever $\vec{p}' \gg 0$ and $\vec{p} \gg 0$ are such that for some l, $p'_l > p_l$ and $p'_k = p_k$ for $k \neq l$, the following conditions hold:*

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

We interpret this definition as follows. When the price of a file that is relevant to peer i increases, peer i demands more of all other files he is downloading, and supplies less of the file he is uploading. As one example, it is straightforward to verify that peer i's optimization problem satisfies gross substitutes if $T_i = \{j\}$ and $v_i(x_{ij}, y) = x_{ij}^{\alpha}/\alpha - y$, where $0 < \alpha < 1$.

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

Proposition 10 *If the optimization problem of each peer satisfies the gross substitutes property, and $\forall j, k \in F$ there exists $i \in U$ such that $j, k \in T_i$, then there is at most one competitive equilibrium up to scaling of the price vector.*

3.3 Tâtonnement Process

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

When every peer is uploading a single file, the aggregate excess demand is a function. A reasonable way to adjust the prices in order to reach a competitive equilibrium is to increase the prices of the files whose excess demand is positive, and decrease the prices of the files whose excess demand is negative. This motivates the *tâtonnement process* [50], where the price adjustment rate is equal (or in general proportional) to excess demand:

$$\frac{dp_j}{dt} = z_j(\vec{p}). \tag{3.5}$$

The following theorem is a restatement of Proposition 17.H.1 [50] for our model.

Theorem 2 *If $|S_i| = 1$ for all $i \in U$, the gross substitutes property holds for the aggregate excess demand function, and Assumptions 3, 4 and 5 are satisfied, then the relative prices of any solution trajectory of (3.5) converge to the unique equilibrium (up to scaling of the price vector).*

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

Theorem 3 *Suppose that $|S_i| = 1$ for all $i \in U$, the gross substitutes property holds for the aggregate excess demand function, and Assumptions 3, 4 and 5 are satisfied. Suppose also that $U = U_1 \cup \ldots \cup U_K$ with $U_k \cap U_l = \emptyset$ whenever $k \neq l$, and:*

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

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

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

3.4 Proportional Allocation

Although the tâtonnement process provides a price adjustment mechanism that (under reasonable assumptions) ensures that a competitive equilibrium is reached, it has a serious shortcoming from a system design standpoint: the tâtonnement process does not specify how agents should engage in trade *before* equilibrium is reached. Thus, in addition to adjusting the price vector according to the tâtonnement process (3.5), our system design should specify a mechanism for allocating rates out of equilibrium.

Mechanisms for exchange out of equilibrium have been proposed for an exchange economy in the economics literature [35], but do not directly apply to our model. These mechanisms work by performing part of the exchange, then updating the endowment of each peer. However, in a peer-to-peer system, the endowment of a peer at any given time is determined by the file with the maximum price, and the amount he owns is his upload capacity. Therefore, the amount remains the same even after a peer uploads the file once.

We consider an alternate system design, which is similar to the approach given in [63] (which builds on [65]): we will ask peers to report the *total upload rate* they are willing to allow, and the *proportions* of their budget they wish to spend on downloading various files. For analytical simplicity, we consider the case where each peer uploads a single file, i.e., $|S_i| = 1$ for all i. Let $f(i)$ denote the file peer i is uploading. Suppose that each peer i

optimizes with respect to the current prices $\vec{p} \gg 0$, and reports his optimal upload rate to the system, $y_i(\vec{p})$. If peer i is interested in downloading multiple files, i.e., $|T_i| > 1$, then he also reports what *proportion* $\pi_{ij}(\vec{p})$ of his budget he wants to spend on each file in $j \in T_i$. In terms of $\vec{x}_i(\vec{p})$ and $y_i(\vec{p})$, if $y_i(\vec{p}) > 0$, we have:

$$\pi_{ij}(\vec{p}) = \frac{p_j x_{ij}(\vec{p})}{p_{f(i)} y_i(\vec{p})}. \tag{3.6}$$

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

In order to formally motivate the proportional allocation mechanism, we first consider the outcome at a competitive equilibrium. From Definition 3 we know that the rates $(x^*_{ij}, i \in U, j \in T_i)$, $(y^*_i, i \in U)$ and the price vector $\vec{p}^*$ constitute a competitive equilibrium if the following conditions are satisfied.

1. Each peer optimizes, i.e., $x^*_{ij} = x_{ij}(\vec{p}^*)$, $y^*_i = y_i(\vec{p}^*)$ for all $i \in U$, for all $j \in T_i$.

2. The market clears, i.e., $\sum_{i: j \in T_i} x^*_{ij} = \sum_{i: f(i) = j} y^*_i$ for all $j \in F$.

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

In particular, given prices $\vec{p}$, suppose peer i reports his desired upload rate $y_i(\vec{p})$ and what proportion of his budget he wants to spend on each $j \in T_i$, π_{ij}; we do not assume anything about $\vec{\pi}$, other than $\pi_{ij} \geq 0$ and $\sum_{j \in T_i} \pi_{ij} = 1$. Because $\vec{p}$ may not be a competitive equilibrium price vector, in general it is not possible to choose download rates at the current prices that ensure each peer i spends exactly the desired proportion π_{ij} on file j. Instead, we will use $\vec{\pi}$ and $\vec{y}(\vec{p})$ to compute a *download* rate allocation $\hat{\vec{x}}_i$ to each peer i, together with a *new price vector* $\hat{\vec{p}}$ such that each peer i earns a budget of $\hat{p}_{f(i)} y_i(\vec{p})$, and spends exactly a proportion π_{ij} on file j; i.e., (3.6) is satisfied for all i and $j \in T_i$ with $y_i(\vec{p}) > 0$. This is a relaxation of Condition 1 above: of course the resulting allocation may not be optimal for each peer given the prices $\hat{p}$; however, the following *budget constraint* will hold:

$$\sum_{j \in T_i} \hat{p}_j \hat{x}_{ij} = \hat{p}_{f(i)} y_i(\vec{p}). \tag{3.7}$$

This ensures every agent has maximally spent their available budget under the new prices $\hat{\vec{p}}$; this is a requirement of optimality, cf. Lemma 9.

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

Proposition 11 *Suppose $|S_i| = 1$ for all $i \in U$. Suppose each peer i reports an upload rate y_i, and a vector $\vec{\pi}_i$ describing the proportion π_{ij} of his eventual budget to be spent on file j. Then there exists a pair $\hat{\vec{p}}$ and $\hat{\vec{x}} = (\hat{\vec{x}}_i, i \in U)$ such that:*

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

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

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

In practice, such a mechanism suggests a natural means to adapting prices as well as allocations. In particular, the following corollary ensures that if the upload rates and requested proportions arose from a competitive equilibrium, then the allocation mechanism given in Proposition 11 will yield the competitive equilibrium allocation.

Corollary 4 *Suppose $\vec{p}^* \gg 0$ is a competitive equilibrium, and $\vec{y}^* = \vec{y}(p^*)$ and $\vec{\pi}^* = \vec{\pi}(\vec{p}^*)$. Let $\hat{\vec{p}}$ and $\hat{\vec{x}}$ the corresponding prices and download rates, respectively, of Proposition 11. Then $\hat{\vec{x}} = \vec{x}(\vec{p}^*)$.*

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

3.4.1 Two Files, Two Peer Types

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

Assumption 6 *The utility of peer i of type k when he is downloading at rate $x_{ki} \geq 0$ and uploading at rate $y_{ki} \geq 0$ is $u_{ki}(x_{ki}) - y_{ki}$, where $u_{ki}(x_{ki})$ is continuously differentiable, strictly concave, and strictly increasing.*[2]

In the next section, we characterize competitive equilibria for this system. In Section 3.4.1, we study Nash equilibria of a game where peers have utilities that satisfy Assumption 6, and anticipate that prices and allocations are chosen according to the proportional allocation mechanism. We establish a competitive limit theorem: in the large system limit,

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

it is as if each peer optimizes as a price taker. In Section 3.4.2, we specialize further to a case where all peers of the same type share the same utility function. This allows us to establish uniqueness of the Nash equilibrium as well, and gives a more precise characterization of the Nash equilibrium rates. Finally, in Section 3.4.3, we study the efficiency of the rate allocation obtained at a Nash equilibrium.

Competitive Equilibrium

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

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

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

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

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

$$p_1 u'_{1i}(0) \leq 1, \text{ if } y_{1i} = 0; \tag{3.8}$$

$$p_1 u'_{1i}(p_1 y_{1i}) = 1, \text{ if } 0 < y_{1i} < B_{1i}; \tag{3.9}$$

$$p_1 u'_{1i}(p_1 B_{1i}) \geq 1, \text{ if } y_{1i} = B_{1i}. \tag{3.10}$$

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

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

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

We know from Corollary 3 that a competitive equilibrium always exists. The following proposition characterizes the competitive equilibria.

Proposition 12 *If* $\sup_i u'_{1i}(0) \cdot \sup_i u'_{2i}(0) \leq 1$, *then at any competitive equilibrium* $y_{1i} = y_{2i} = 0$ *for all i. On the other hand, if* $\sup_i u'_{1i}(0) \cdot \sup_i u'_{2i}(0) > 1$, *then at any competitive equilibrium there exist* i, j *such that* $y_{1i} > 0$ *and* $y_{2j} > 0$.

Nash Equilibrium

We use the proportional allocation mechanism to clear the market out of equilibrium. The results in Proposition 11 are simplified in this case, because there are only two files ($|F| = 2$) and each peer is downloading a single file. Thus peers only report upload rates; it is clear that they will spend their entire budget on the single file they wish to download. Let y_{ki} be the upload rate that peer i of type k reports and $Y_k = \sum_i y_{ki}$. If $Y_1 > 0$ and $Y_2 > 0$, it is straightforward to check that the proportional allocation mechanism will use the following price to clear the market:

$$\hat{p}_1 = \frac{Y_2}{Y_1}. \tag{3.11}$$

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

$$\Pi_{1i}(y_{1i}) = \begin{cases} u_{1i}(y_{1i}Y_2/Y_1) - y_{1i}, & \text{if } Y_1 > 0; \\ u_{1i}(0), & \text{if } Y_1 = 0. \end{cases} \tag{3.12}$$

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

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

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

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

$$u'_{1i}(0) \leq \frac{Y_1}{Y_2}, \text{ if } y_{1i} = 0; \tag{3.13}$$

$$u'_{1i}\left(y_{1i}\frac{Y_2}{Y_1}\right)\left(1 - \frac{y_{1i}}{Y_1}\right) = \frac{Y_1}{Y_2}, \text{ if } 0 < y_{1i} < B_{1i}; \tag{3.14}$$

$$u'_{1i}\left(B_{1i}\frac{Y_2}{Y_1}\right)\left(1 - \frac{B_{1i}}{Y_1}\right) \geq \frac{Y_1}{Y_2}, \text{ if } y_{1i} = B_{1i}. \tag{3.15}$$

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

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

Theorem 4 *If Assumption 6 is satisfied, $u'_{1i}(0) > N_1/(N_1-1)$ for all i, and $u'_{2i}(0) > N_2/(N_2-1)$ for all i, then there exists a Nash equilibrium $(\vec{y}_1, \vec{y}_2)$ at which not all rates are equal to zero.*

We now develop a *competitive limit*, where the number of peers of each type becomes large. Suppose that $N_1, N_2 \to \infty$, and consider a sequence of Nash equilibria $\vec{y}^N$ indexed by $N = N_1 + N_2$; by taking subsequences if necessary, we can assume the Nash equilibria converge, say to $\vec{y}$. Let $Y_k^N = \sum_i y_{ki}^N$. Suppose that $y_{ki}^N / Y_k^N \to 0$ for all peers i of type k, but that $Y_2^N / Y_1^N \to p_1 \in (0, \infty)$; we normalize $p_2 = 1$. Under these assumptions, since the optimality conditions (3.13)-(3.15) are continuous, they become identical to the optimality conditions (3.8)-(3.10) for a competitive equilibrium. Thus informally, we expect that the Nash equilibrium rates should approach competitive equilibrium rates.

Formally, recall that we define $x_{ki}(p_1)$ and $y_{ki}(p_1)$ as the optimal solutions for a price taking peer (i.e., a peer solving (3.1)-(3.4)), given a price p_1. We then have the following theorem.

Theorem 5 *Let $N = N_1 + N_2$ be the total number of peers. Suppose that as $N \to \infty$, both $N_1 \to \infty$ and $N_2 \to \infty$. Suppose that Assumption 6 holds for the utility function of each peer, $\sup_i B_{ki} < \infty$, $\sup_i u'_{ki}(0) < \infty$, and $\inf u'_{ki}(0) > 1$ for $k = 1, 2$. Let $\vec{y}^N$ denote a nonzero Nash equilibrium when $N = N_1 + N_2$ peers are in the system, and let $p_1^N = Y_2^N / Y_1^N$, where $Y_k^N = \sum_i y_{ki}^N$. Then:*

1. $0 < \inf_N p_1^N$ *and* $\sup_N p_1^N < \infty$.
2. *For all i and k, $y_{ik}^N / Y_k^N \to 0$ as $N \to \infty$, while $Y_k^N \to \infty$ as $N \to \infty$.*
3. *Any limit point $(p_1, \vec{y})$ of the sequence $(p_1^N, \vec{y}^N)$ satisfies the competitive equilibrium optimality conditions* (3.8)-(3.10).

The preceding theorem shows that in the large system limit, it is as if each peer optimizes as a price taker. Observe that from the proof, in the limit we have infinite upload rates for both types of files; thus we cannot directly interpret the limit point as a competitive equilibrium. However, we can make the following precise statement: asymptotically, peers choose upload rates that are nearly equal to their optimal upload rate *if* they were acting as price takers. One way to interpret such a theorem is that in large peer-to-peer systems, fully strategic behavior by the peers will not ultimately cause large deviations from competitive equilibrium behavior.

3.4.2 Homogeneous Utilities

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

Throughout this section, to avoid boundary conditions, we will make the following additional simplifying assumption about the utility functions; the analysis can be extended to study the case where the assumption does not hold, but without a significant change in insight.

Assumption 7 *The function $u(\cdot)$ satisfies $u'(x) \to \infty$ as $x \to 0$, and $u'(x) \to 0$ as $x \to \infty$.*

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

Competitive Equilibrium We show that under Assumptions 6 and 7, the price $p_1 = 1$ is always a competitive equilibrium of this economy. First suppose that $u'(B) < 1$. Since $u'(\cdot)$ is continuous, there is a $y \in (0,B)$ such that $u'(y) = 1$. Then, when $p_1 = 1$, a peer of type 1 will choose to upload and download $u'^{-1}(1)$, and the same for all peers of type 2. Since the total upload and download rates of a file are equal, this is a competitive equilibrium. On the other hand, if $u'(B) \geq 1$, then when $p_1 = 1$, all peers will choose to upload B. In this case the upload rate constraint binds, and we again have a competitive equilibrium.

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

Lemma 10 *Suppose Assumptions 6 and 7 are satisfied, and $u(\cdot)$ is twice differentiable. Then, the following are equivalent:*

1. *For all $B > 0$, the Multilateral Peer Optimization of each peer satisfies the gross substitutes property.*

2. *$pu'^{-1}(p)$ is nonincreasing on $(0,\infty)$.*

3. *$xu'(x)$ is nondecreasing.*

In this case the competitive equilibrium is unique.

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

Lemma 11 *If $u(\cdot)$ is a strictly concave function, then peers of the same type will have the same upload rate at any Nash equilibrium.*

For the remainder of this section, we suppose that Assumptions 6 and 7 hold. If $u'(x) \to \infty$ as $x \to 0$, the optimality condition (3.13) will never apply. Let y_1, y_2 be the rates at which peers of type 1 and type 2 upload, respectively, at a Nash equilibrium; and recall that N denotes the number of peers of each type. If $y_1 > 0$ and $y_2 > 0$, the optimality conditions (3.14) and (3.15) can be equivalently written:

$$u'(y_2) = \frac{N}{N-1}\frac{y_1}{y_2}, \text{ if } 0 < y_1 < B; \tag{3.16}$$

$$u'(y_2) \geq \frac{N}{N-1}\frac{B}{y_2}, \text{ if } y_1 = B; \tag{3.17}$$

Similarly, for peers of type 2 the following conditions hold.

$$u'(y_1) = \frac{N}{N-1}\frac{y_2}{y_1}, \text{ if } 0 < y_2 < B; \tag{3.18}$$

$$u'(y_1) \geq \frac{N}{N-1}\frac{B}{y_1}, \text{ if } y_2 = B; \tag{3.19}$$

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

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

a contradiction.

If $u'(0) > N/(N-1)$, then there exists a Nash equilibrium with $y_1 > 0$ and $y_2 > 0$: for example, $y_1 = y_2 = \min(u'^{-1}(N/(N-1)), B)$. When the upload rates are positive we define the Nash price as $p^{NE} = y_2/y_1$. Thus $p^{NE} = 1$ is a possible Nash price and we know that $p^* = 1$ is a competitive equilibrium price. In particular, if $u'(0) > N/(N-1)$, there exists a Nash equilibrium with the same price as the unique competitive equilibrium. Theorem 5 does not apply here, because of Assumption 7. However, by Lemma 11, $y_{ki}/Y_k = 1/N \to 0$ as $N \to \infty$, and thus any limit point $(p_1, \vec{y})$ of the sequence $(p_1^N, \vec{y}^N)$ satisfies the competitive equilibrium optimality conditions. Moreover, it can be shown that the rates of any sequence of Nash equilibria converge to the rates of a competitive equilibrium.

The following proposition is our key result for the model with homogeneous peers: we show that if gross substitutes holds, there exists a unique Nash equilibrium where the upload rates are strictly positive. The proof uses the characterization of gross substitutes shown in Lemma 10.

Proposition 13 *Suppose that Assumptions 6 and 7 hold and $u(\cdot)$ is twice differentiable. Then, if the Multilateral Peer Optimization of each peer satisfies the gross substitutes property, there is a unique Nash equilibrium with strictly positive rates. At the equilibrium $y_1 = y_2 = u'^{-1}(N/(N-1))$.*

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

3.4.3 Efficiency

We consider a Nash equilibrium of the game that results from the proportional allocation mechanism at which not all rates are zero; i.e., $Y_1 > 0$ and $Y_2 > 0$. Notice that when type 1 peers choose their optimal upload rate, they take Y_2 as *given*. Thus, we can interpret the rates y_{1i} reported by type 1 peers as a Nash equilibrium to the following auction game. Suppose that the available upload rate of file 2 is fixed and equal to Y_2. Type 1 peers submit bids to acquire a share of the available file transfer rate for file 2; each peer has to pay his bid, and is allocated a download rate proportional to his bid. In [41], it is shown for this game that if Assumption 6 is satisfied, and for all i $u_{1i}(0) \geq 0$, then:

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

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

On the other hand, Nash equilibria need not be Pareto efficient. Suppose that Assumption 6 holds, peers are homogeneous, and $1 < u'(0) \leq N/(N-1)$, where N is the number of peers in each type. Then at any competitive equilibrium, all upload and download rates will be strictly positive, while at any Nash equilibrium all rates will be zero. Since each peer has the option of uploading and downloading zero in the competitive equilibrium, this shows that each peer is strictly worse off at the Nash equilibrium.

3.5 Comparing Pricing Schemes

In preceding sections we formulated a peer-to-peer system as an exchange economy and studied multilateral exchange. Our previous analysis has two limitations. First, it considers a setting where transfers are only constrained by peer upload capacity; however, often the network structure is more complex. Second, it doesn't provide a practical way to set prices. In this section we address these limitations by discussing how we can incorporate a more complex network structure in the model, as well as decentralized dynamics.

With a complex network structure, we envision that network links are also priced in order to capture network constraints. Then, it may not suffice to have one price per file, and at first it may seem necessary to maintain one price per file per peer. However, this is not the case. In Section 3.5.1, we compare a range of pricing schemes for multilateral exchange and conclude that simply maintaining a single price per peer suffices to achieve the benefits of price-based multilateral exchange. This result means that a more complex approach of variable pricing (at each peer) for different files or different chunks of files is, in fact, neither necessary nor beneficial. Intuitively, even with a single price per peer, variable prices for files still arise *across the network* since different peers supply different files. In Section 3.5.2, we consider how explicit prices, pricing per peer, and currency benefit system dynamics.

3.5.1 Equilibria

This section compares three pricing schemes for multilateral exchange: (1) one price per file (denoted PF); (2) one price per peer (denoted PP); and (3) one price per file per peer (denoted PFP). We compare the schemes in terms of equilibrium existence. First, we show that all three are equivalent when transfers are only constrained by peer upload capacity. However, we then demonstrate that PF may be strictly worse than PP if the network topology is non-trivial. Finally, we show that PP and PFP yield equivalent equilibria, even when the network topology is non-trivial. Since explicitly pricing every file of every peer is much more complicated than only maintaining a single price per peer, our analysis suggests PP is the most desirable scheme.

In Section 3.2 we introduced the PF pricing scheme and studied some of its properties in preceding sections. Here, we introduce the other two pricing schemes: PP and PFP. To formulate the optimization problems of peers it is convenient to introduce the following notation: let r_{ijf} be the rate at which peer i sends file f to peer j. We also reformulate the

PF pricing scheme with this notation. We now give the peer optimization problem.

$$
\begin{aligned}
&\text{maximize } v_i(\vec{x}_i, y_i) \\
&\text{subject to } x_{if} = \sum_j r_{jif}, \forall f \in T_i \\
&\qquad\qquad y_i = \sum_{j,f} r_{ijf} \\
&\qquad\qquad r_{kjf} = 0, \text{ if } f \notin F_k \\
&\qquad\qquad y_i \leq B_i \\
&\qquad\qquad \text{Budget Constraint} \\
&\qquad\qquad \vec{r} \geq 0.
\end{aligned}
$$

The first four constraints (giving download rates, upload rate, ensuring peers only upload files they possess, and meeting the bandwidth constraint) are identical for all pricing schemes. The budget constraint ensures that the capital peer i spends for downloading does not exceed the capital the peer accrues by uploading. Given prices $(p_i, i \in N)$, $(p_f, f \in F)$ and $(p_{if}, i \in N, f \in F)$ for the PP, PF and PFP pricing schemes respectively, the corresponding budget constraints are shown in the following table.

Pricing Scheme	**Budget Constraint**
PF	$\sum_{j,f} p_f r_{jif} \leq \sum_{j,f} p_f r_{ijf}$
PP	$\sum_{j,f} p_j r_{jif} \leq p_i \sum_{j,f} r_{ijf}$
PFP	$\sum_{j,f} p_{jf} r_{jif} \leq \sum_{j,f} p_{if} r_{ijf}$

Analogously to Definition 3, a *competitive equilibrium* is a combination of a rate allocation vector and a price vector such that all peers have solved their corresponding optimization problems. In this case, the prices have exactly aligned supply and demand: for any i, j, f, the transfer rate r_{ijf} is simultaneously an optimal choice for both the uploader i and downloader j. We next give the equilibrium definition for the PP pricing scheme. (Similarly, an equilibrium can be defined for the two other pricing schemes.)

Definition 6 *The rate allocation $\vec{r}^* = (r^*_{ijf}, i, j \in N, f \in F)$ and the peer prices $(p^*_i, i \in N)$ with $p^*_i > 0$ for all $i \in N$ constitute a* PP competitive equilibrium if for each peer j, $\vec{r}^*$ solves the Peer Optimization problem given prices $(p_i^*, i \in N)$.

We first observe that if transfers are only constrained by the upload capacity of peers, then the PF and PP schemes are equivalent in terms of equilibria, i.e., an equilibrium for one scheme exists if and only if there exists an equilibrium with the same rate allocation for the other pricing scheme. In particular, given PF equilibrium prices $(p_f^*, f \in F)$, peer i will only be uploading files in $\arg\max_{f \in F_i}\{p_f^*\}$. Setting $p_i^* = \max_{f \in F_i}\{p_f^*\}$, we get an equilibrium for the PP scheme, since the optimization problem of each peer does not change. Similarly, if $(p_i^*, i \in F)$ are PP equilibrium prices, then $p_f^* = \min_{i: f \in F_i}\{p_i^*\}, f \in F$ are PF equilibrium prices which yield the same rate allocation.

Given a non-trivial network topology, however, links other than peer access links may be congested. These links need to be priced as well to ensure efficient network usage. Again abusing notation, we denote the price of link ℓ by p_ℓ. For now we assume that we can price every link in the network; we relax this assumption in Section 3.6. When peer i is downloading from peer j, i pays j, but also all links that i's traffic traverses.

Network links are priced in order to make peers internalize their effect on the network, and not for profit. Thus, it is desirable to rebate any payments related to network costs to peers. We will assume that whatever is paid to traverse links in the network is rebated equally to all peers; however, our results also hold for other rebating schemes.

When the network topology is non-trivial we can modify the budget constraints to include the payment to the network on the left hand side and the rebate from the network on the right hand side. A *competitive equilibrium* now is a combination of a rate allocation vector and a price vector (which includes prices for links) such that all peers have solved their corresponding optimization problems and the total traffic that traverses each link does not exceed its capacity.

The following example shows that when the network is non-trivial, equilibria may fail to exist under the PF scheme, even though they exist for the PP and PFP schemes.

Example 6 *There are four peers and two files, with file allocation and demand as shown in Figure 3.1. The network has two clusters, consisting of peers* $\{1,2\}$ *and* $\{3,4\}$*, with a bidirectional link* ℓ *of capacity* 1 *connecting them. Peers* $\{1,3\}$ *have file* f *and want file* g*; peers* $\{2,4\}$ *have file* g *and want file* f*. The peers' upload capacities are* $B_1 = B_4 = 8$ *and*

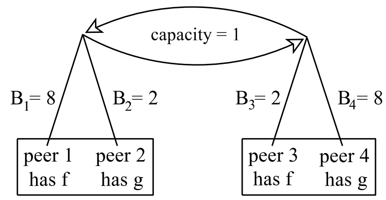

Figure 3.1: Example with a PP equilibrium, but no PF equilibrium (see Example 6). System with peers $\{1,2,3,4\}$ and files $\{f,g\}$. Peers are located in two clusters; transfers are constrained by bandwidth constraints of peers and the inter-cluster link.

$B_2 = B_3 = 2$.

This system has no equilibrium under the PF scheme. If $p_\ell = 0$, peers demand $x_1 = (p_f/p_g)B_1$, $x_2 = (p_g/p_f)B_2$, $x_3 = (p_f/p_g)B_3$ and $x_4 = (p_g/p_f)B_4$ when optimizing. The market clears only if $x_1 + x_3 = B_2 + B_4$, which implies $p_f/p_g = 1$. But then $x_1 = 8$, which is not feasible, since the maximum total rate at which peer 1 can download is 3. If $p_\ell > 0$ and there is a single price per file, then peers only download locally and, from the market clearing condition, we get $p_f/p_g = 4$ in one cluster and $p_f/p_g = 1/4$ in the other, which is a contradiction.

On the other hand, under the PP scheme, there is an equilibrium with prices $(p_1, p_2, p_3, p_4, p_\ell) = (1,3,3,1,2)$ and rates $(x_1,x_2,x_3,x_4) = (3,7,7,3)$. Peers 1 and 4 download at rate 2 locally and at rate 1 remotely from each other. The revenue collected from the link ℓ is rebated equally to all peers, which allows peer 2 to download more than $(p_1/p_2)B_2$.

The preceding example shows that for general network topologies, the existence of a competitive equilibrium for the PP scheme does not imply the existence of a competitive equilibrium for the PF scheme. The reason is that a file may be uploaded at different prices at different parts of the network.

PFP is the most general pricing scheme. The following result shows that PP is equivalent to PFP in terms of equilibria. The proof is similar to the proof of equivalence between PP and PF in the trivial network setting. It shows that given equilibrium prices for one pricing scheme we can construct prices for the other pricing scheme that yield the same equilibrium allocation.

Proposition 14 *For any network topology, there exists a competitive equilibrium for the PP scheme if and only if there exists a competitive equilibrium for the PFP scheme.*

We conclude that one price per peer is sufficient to identify heterogeneity in the system. Intuitively, the upload capacity of a peer's access link is the local resource that becomes congested, and in market design it is typically the case that one price is required for each congestible resource; hence one price per peer suffices for competitive equilibrium. The argument holds not only for different files, but also for different chunks of the same file. In light of Example 6, a price for each peer is the minimal amount of information needed.

The PP scheme provides many practical benefits as well. It greatly reduces the number of prices that need to be maintained, compared to PFP pricing. Further, it simplifies price discovery and leads to a natural service discipline for uploading files, as discussed in Section 3.5.2.

3.5.2 Dynamics

The preceding section considers equilibria, i.e., a *static* setting. However, since it is hard to know a system's equilibrium prices or allocation in advance, we need to consider several issues related to *dynamics*: how downloaders and uploaders are matched (peer discovery), how out-of-equilibrium prices are updated (price discovery), and which requests uploaders satisfy (service discipline). In this section we discuss the role of *explicit* per-peer pricing in aiding dynamics; we also briefly discuss the advantages of explicit currency.

A significant advantage of *explicit* per-peer prices is that they enable fast peer discovery. This short discovery time significantly improves on that needed by systems with *implicit* prices, e.g., such as BitTorrent's rate-based exchange ratios, which have long discovery times. These implicitly-priced systems need to perform a brute-force search across their peers; this has been found in practice to sometimes take tens of minutes [13], and, at least for high-bandwidth peers, requires an asymptotically-linear search to find similar reciprocation rates.

For price discovery, a simple mechanism is to update the prices of peers and links according to the corresponding *excess demand*.[3] In particular, a price should increase if demand exceeds supply, and decrease if demand trails. But how to define supply and demand for a peer? If the PP scheme is employed, a peer's observed demand is the sum of

[3]This is a decentralized version of the tâtonnement process discussed in Section 3.3.

all received rate requests, and his supply is equal to the upload capacity of his access link (within a fixed time period). If the PFP scheme is used, excess demand is more complex, since it needs to be calculated for each file separately, and a peer's supply for a file depends on the prices (it is optimal for each peer to only upload his most expensive file). Thus the PP scheme leads to simpler price dynamics. The benefits of the PP pricing scheme are even more apparent for unpopular files for which requests are relatively rare.

PP pricing leads to a natural service discipline for uploading files, well-aligned with a peer's incentives: serve requests sequentially and without preemption. The service discipline for PFP pricing is less clear, however: serving requests only for the highest-priced file may not fully utilize a peer's available resources, while serving requests sequentially is not profit maximizing.

We note that the PF pricing scheme has a further disadvantage in terms of price dynamics. If all peers in the system change the same price for a file, is not possible to update this file price in a decentralized fashion. A centralized method, such as that the tâtonnement process (Section 3.3), is required.

Allowing peers to store and exchange currency over time also has significant benefits for a system in a dynamic setting. First, peers can engage in trade *before* equilibrium prices are reached. Second, peers can reach a rate allocation by *trading in a decentralized fashion*, without requiring a central authority to clear the market by matching uploaders and downloaders.

3.6 Implications for System Design

In this section we discuss some system design issues at a high level, and argue that users are incentivized to contribute valuable content and use network resources efficiently.

In the previous sections we argued that explicit prices and currency facilitate multilateral exchange and showed that one price per peer suffices to achieve the benefits of multilateral exchange. In [5] we discuss a system design (PACE) that enables multilateral exchange through currency and the per-peer pricing scheme. The high level picture is the following. Given prices of peers and network links, a peer requests download rates from other peers in the network. For downloading, a peer pays the uploading peer and the links its traffic traverses. Payments to links are rebated equally to all peers. A peer serves requests sequentially without preemption, and updates its price according to the mismatch between requests received and available capacity.

Network links are priced in order to make peers internalize their effect on the network. However, associating a price with every link in the network is probably infeasible in practice, given a lack of network topology and routing information, inaccurate bandwidth capacity estimations, and computational complexity. Alternatively, we can separately price most of the bottlenecks that lead to supply constraints. For instance, to a first approximation, we can view the entire network as composed of local clusters, connected together by the wide-area core. In many cases we can assume that most bottlenecks – especially those due to transient congestion – are at access links, where we can accurately estimate capacity and adapt prices accordingly. Rare bottlenecks in the core are captured via slowly-changing network prices. Different network prices may be associated with different pairs of clusters.

One typical complaint against explicit pricing and currency is that the process of setting prices and bidding for goods becomes a usability hurdle. Indeed, this hurdle is seen by the designers of systems such as MojoNation as their major reason for failing to be widely adopted. However, pricing mechanisms, rather than being directly exposed to end-users, can serve as algorithmic devices to ensure efficient exchange: We can expose a very simple interface to users and have users' software optimally compute their buy and sell behavior. We expect that even strategic users cannot gain any significant benefit from operating in a manner other than that specified by our algorithms.

Finally, incorporating currency into a system introduces its own complications, since peer exchanges and credit balances need to be secured. [5] demonstrates how these issues can be handled.

User Incentives

In this section we study the incentives provided to users in our multilateral exchange model. *First*, we note that our system encourages efficient use of resources in a *large* peer-to-peer system. If the system is large, users cannot accurately anticipate how their actions affect prices (e.g., Section 3.4), and users have difficulty in predicting the evolution of demand, supply, and prices.

Second, users are given incentives to contribute. In particular, users are implicitly incentivized both to contribute a high percentage of their upload capacities and to share high-value content, since high-value files will typically increase a user's price. Moreover, with currency stored over time, users are incentivized to contribute even if they are not currently downloading. We also note that a user is paid for delivered rates, so a user does not profit by advertising a higher upload capacity than the one he has.

Third, network prices align user incentives with efficient network usage. If network links are priced correctly, we expect network prices to reflect congestion. Then users internalize their effect on the network, since they have to pay for all network links they use. For instance, among peers with the same price, a user prefers to download from the peer with the smallest total network cost, which we expect to correspond to the least congested route. On the other hand, sellers do not benefit from network-cost-related payments, and so the system does not create a perverse incentive for sellers to prefer remote transfers.

Finally, one potential concern in a market-based system is the phenomenon of *market power*: Users may try to manipulate prices higher than the laws of supply and demand dictate. This effect is mitigated significantly in a market with many sellers, where market manipulation cannot significantly increase profit [50]. To illustrate this, suppose k peers have a file desired by other users in the system. Let $d(p)$ be the demand for the file at price p, and B_j be the bandwidth of peer j. The equilibrium price p^* satisfies $d(p^*) = \sum_{j=1}^{k} B_j$. Peer i would benefit from increasing his price to $p > p^*$ only if

$$\frac{B_i}{d(p)} \geq -\frac{d(p) - d(p^*)}{p - p^*} \cdot \frac{p}{d(p)}.$$

The right-hand side represents the demand elasticity, i.e., the percentage change in quantity demanded, divided by the percentage change in price. Thus, user i will not be able to exert market power as long as enough sellers compete to upload the file relative to the elasticity of demand.

Market power may still be an issue if only a few users have a file, yet any system can suffer if such users choose to dictate terms to the remainder. In our setting, such users are often the "seeders" of files and *want* to see their content disseminated. Many peers in existing peer-to-peer systems exhibit such altruistic behavior.

More generally, sellers *create* other uploaders in the very act of uploading; thus, market power is at best a transient phenomenon, since other sellers quickly emerge as competitors. Further, the number of competitors grows exponentially in time: if a file chunk is always transferred from one user to another in one time period, then after t time periods, $O(2^t)$-times more users will have it.

3.7 Proofs for Chapter 3

Proof of Proposition 9: Suppose that the budget constraint does not bind. Then there is an optimal solution $(\vec{x}_i, y_i)$ with $\sum_{j \in T_i} x_{ij} p_j < (\max_{j \in S_i} p_j) \cdot y_i$. And since $x_{ij} \geq 0$ for all j, we will have $y_i > 0$. However, we can choose a small ε such that the solution $(\vec{x}'_i, y'_i) = (\vec{x}_i, y_i - \varepsilon)$ is feasible and $v_i(\vec{x}_i, y_i) < v_i(\vec{x}'_i, y'_i)$ because of Assumption 3. This contradicts the assumption that $(\vec{x}_i, y_i)$ is optimal. ■

Proof of Proposition 9: If Assumption 5 holds, then either the user graph is strongly connected, or the system can be decomposed to subsystems for which the user graphs are strongly connected. Therefore, without loss of generality, in this proof we assume that the user graph is strongly connected.

By Lemma 9, for any user the budget constraint will bind at the optimal solution. In particular, given any choice of $(y_{ij}, i \in U, j \in F)$ that satisfies the conditions of Definition 2, we have for each i:

$$\sum_{j \in T_i} p_j x_{ij}(\vec{p}) - \sum_{j \in S_i} p_j y_{ij}(\vec{p}) = 0.$$

By summing over all users, we obtain Property 1.

Fix a price vector $\vec{p} \gg 0$. The set of vectors $(y_{ij}, i \in U, j \in S_i)$ that satisfy Conditions 2, 3 and 4 of Definition 2 is convex. Thus the aggregate excess demand $\vec{z}(\cdot)$ is a convex valued correspondence (Property 2).

Consider a price vector $\vec{p} \gg 0$, and fix a constant $t > 0$. It is clear that the feasible region (3.2)-(3.4) remains unchanged if we replace the price vector $\vec{p}$ by $t\vec{p}$; we conclude that $\vec{x}_i(\vec{p}) = \vec{x}_i(t\vec{p})$, and $y_i(\vec{p}) = y_i(t\vec{p})$; i.e., $\vec{x}_i$ and y_i are homogeneous of degree zero. Thus by Definition 2, the aggregate excess demand is also homogeneous of degree zero (Property 3).

We now show that the aggregate excess demand correspondence has a closed graph. We start by showing that $\vec{x}_i(\cdot)$ and $y_i(\cdot)$ are continuous functions. By Assumption 3 $v(\cdot)$ is a continuous function. From the Theorem of the Maximum [10] it follows that $x_{ij}(\vec{p})$ and $y_i(\vec{p})$ are continuous functions.

Consider the sequences $\vec{p}^m \to \vec{p} \gg 0$ and $\vec{w}^m \to \vec{w}$ such that $\vec{w}^m \in \vec{z}(\vec{p}^m)$. Since $\vec{w}^m \in \vec{z}(\vec{p}^m)$, there exist y_{ij}^m, $i \in U$, $j \in F$, that satisfy Conditions 1, 2, 3 and 4 of Definition 2 for the price vector $\vec{p}^m$ and the aggregate excess demand vector $\vec{w}^m$. We will show that $\vec{w}$ satisfies these conditions when the price vector is $\vec{p}$, and thus $\vec{w} \in \vec{z}(\vec{p})$.

Fix $\varepsilon > 0$. Since $y_i(\cdot)$ is continuous, there exists M such that $y_i(\vec{p}^m) < y_i(\vec{p}) + \varepsilon$, for all $m \geq M$, or equivalently $\sum_{j \in T_i} y_{ij}^m < y_i(\vec{p}) + \varepsilon$, for all $m \geq M$. Moreover, $y_{ij}^m \geq 0$, so for $m \geq M$, y_{ij}^m lies in the compact set $[0, y_i(\vec{p}) + \varepsilon]$. Thus for all $i \in U$ and $j \in F$, the sequence y_{ij}^m has at least one limit point $\bar{y}_{ij}$.

We will show that $\bar{y}_{ij}$ satisfies Conditions 1-4 of Definition 2 with price vector $\vec{p}$ and excess demand vector $\vec{w}$. Since $\vec{w}^m \to \vec{w}$ and $x_{ij}(\cdot)$ are continuous, we have Condition 1 of Definition 2:

$$w_j = \sum_{i \in U : j \in T_i} x_{ij}(\vec{p}) - \sum_{i \in U} \bar{y}_{ij}.$$

We know that $\sum_{j \in F} y_{ij}^m = y_i(\vec{p}^m)$, so by continuity of y_i we have $\sum_{j \in F} \bar{y}_{ij} = y_i(\vec{p})$ (Condition 2). Since $y_{ij}^m \geq 0$ for all m, we have $\bar{y}_{ij} \geq 0$ (Condition 3). Finally, suppose that $j \notin \arg\max_{k \in S_i} p_k$. Then there exists M' such that $j \notin \arg\max_{k \in S_i} p_k^m$ for all $m \geq M'$. Thus $y_{ij}^m = 0$ for all $m \geq M'$, which implies that $\bar{y}_{ij} = 0$ (Condition 4). Thus we conclude $\bar{y}_{ij}$ satisfies all the conditions of Definition 2 with price vector $\vec{p}$ and excess demand vector $\vec{w}$, so $\vec{w} \in \vec{z}(\vec{p})$. This establishes Property 3 of the proposition.

For any price vector $\vec{p} \gg 0$, the feasible region of every User Optimization problem is compact, so we can find an upper bound for the excess demand of any good. Thus for any compact set $B \subset (0, \infty)^F$, $\vec{z}(B)$ is bounded. This completes the proof that $\vec{z}(\cdot)$ is upper hemicontinuous (Property 4).

The upload rate of any user i is upper bounded by his upload rate constraint B_i, so the total supply is upper bounded and the excess demand is bounded from below (Property 5).

If $\vec{p}^m \to \vec{p} \neq 0$ and $p_j = 0$, then $p_k > 0$ for some k. Because of Assumption 5, there is a sequence of users $u_1, u_2, ..., u_l \in U$ and a sequence of files $f_1, ..., f_{l+1}$ such that $f_1 = j$,

$f_{l+1} = k$ and user u_i has file f_i and wants to download file f_{i+1}, so that his utility is strictly increasing in the rate at which he downloads file f_{i+1} (Assumption 3). Thus, there is a user i who has a file $j \in S_i$ to upload whose price approaches a strictly positive limit, and who wants a file $f \in T_i$ whose price approaches zero. The budget of user i approaches a strictly positive limit as $\vec{p}^m \to \vec{p}$ and the amount of f he can afford goes to infinity. On the other hand, the total possible supply is bounded above by the sum of the upload rate constraints B_n of the users n that have $f \in S_n$. Thus $\max\{z_j^m : j \in F\} \to \infty$, establishing Property 6. ■

Proof of Corollary 3: Because of the assumptions on $u_i(\vec{x}_i)$, Assumption 3 is satisfied. By Lemma 9, the budget constraint will bind; thus given $\vec{p} \gg 0$, y_i is a linear function of the download rates of user i. By substituting in the objective function (3.1), we obtain a function of $\vec{x}_i$ that is strictly concave. We conclude the optimization problem of each user has a strictly concave objective and a convex feasible region, and thus a unique solution—i.e., Assumption 4 is satisfied. All the assumption of Theorem 1 are satisfied, so a competitive equilibrium exists. ■

Proof of Proposition 10: It suffices to show that whenever $\vec{p} \gg 0$ and $\vec{p}' \gg 0$ are two price vectors that are not collinear, any corresponding aggregate excess demand vectors can not be equal, i.e. $\vec{z}(\vec{p}) \cap \vec{z}(\vec{p}') = \emptyset$. Since z is homogeneous of degree zero, we can assume that $p'_k \geq p_k$ for all k, and $p_l = p'_l$ for some l. Let $S = \{j : p_j = p_l\}$.

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

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

Proof of Theorem 3: We refer to $\{1, ..., K\}$ as the set of types. Let D_k be the subset of types that download file k and U_k be the subset of types that upload file k. Then we can write the tâtonnement process as $\dot{\vec{p}} = N\vec{f}(\vec{p})$, where $f_k(\vec{p}) = \sum_{i \in D_k} r_i x_{ik}(\vec{p}) - \sum_{i \in U_k} r_i y_i(\vec{p})$. By linearizing around the equilibrium $\vec{p}^*$, we see that the error $\vec{w}(t) = \vec{p}(t) - \vec{p}^*$ satisfies,

$$\dot{\vec{w}}(t) = ND\vec{f}(\vec{p}^*) \cdot \vec{w}(t).$$

This is a system of first order differential equations which has solutions of the form $\vec{w} = \vec{w}_0 e^{\lambda \cdot t}$, where $(ND\vec{f}(\vec{p}^*) - \lambda I)\vec{w}_0 = 0$. The rate of convergence is given by the minimum of $|Re\lambda|$ over all eigenvalues λ of $ND\vec{f}(\vec{p}^*)$. The real part of each eigenvalue of $Df(p^*)$ is negative (because we are assuming convergence) and does not depend on N, so the rate of convergence is linear in N. ■

Proof of Proposition 11: If we multiply through the third condition by $\hat{p}_j$, and substitute from the second condition, we obtain:

$$\sum_{i: j \in T_i} \pi_{ij} \hat{p}_{f(i)} y_i = \hat{p}_j \sum_{i: f(i)=j} y_i. \tag{3.20}$$

Consider a continuous time Markov chain on the state space F, where the transition rate from state (file) j to state (file) k is $Q_{jk} = \sum_{i: f(i)=j} \pi_{ik} y_i$. (Note that $\pi_{ik} = 0$ if $k \notin T_i$.) Let $Q_{jj} = -\sum_{k \neq j} Q_{jk}$. Then (3.20) can be rewritten as:

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

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

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

Proof of Corollary 4: It suffices to note that $\vec{p}^*$ and $\vec{x}(\vec{p}^*)$ satisfy Conditions 1-3 of Proposition 11. Since $\hat{\vec{x}}$ is uniquely determined, it must be the case that $\vec{x}(\vec{p}^*) = \hat{\vec{x}}$. ■

Proof of Proposition 12: We first show that if $\sup_i u'_{1i}(0) \cdot \sup_i u'_{2i}(0) \leq 1$, there does not exist a competitive equilibrium where some upload rate is strictly positive. If such an equilibrium exists, then at least one user from each type must be uploading at a strictly positive rate. Suppose such an equilibrium exists and let $(p_1, 1)$ be the corresponding price vector. Then, there exist users i and j such that:

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

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

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

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

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

Proof of Theorem 4: We use a perturbation approach. Assume there is a "virtual" type 1 user that always uploads ε of file 1, and a "virtual" type 2 user that always uploads ε of file 2. Given strategies $\vec{y}_k$ and $\vec{y}_2$ of type k users, note that $Y_k = \varepsilon + \sum_i y_{ki}$. Thus, for any $\varepsilon > 0$, the utility function of each user i of type k is continuous in the strategies of all users, and concave in y_{ki}. Moreover, the strategy space of each user is compact and convex. Thus, according to Theorem 1 of [62] there exists a Nash equilibrium.

We first show that when $\varepsilon > 0$, at any Nash equilibrium not all upload rates can be zero. Suppose that at some Nash equilibrium $y_{1i} = 0$ for all i. Then, $Y_1/Y_2 = \varepsilon/(\varepsilon + \sum_i y_{2i}) \leq 1$ and (3.13) gives a contradiction. The symmetric argument for type 2 users shows that $y_{2l} > 0$ for some l. Thus, at any Nash equilibrium there exist i, l such that $y_{1i} > 0$ and $y_{2l} > 0$.

Let $\{\varepsilon^n\}$ be a strictly positive sequence such that $\varepsilon^n \to 0$. For each n, let y_{1i}^n, y_{2i}^n be Nash equilibrium rates given ε^n, and let $Y_1^n = \varepsilon^n + \sum_i y_{1i}^n$ and $Y_2^n = \varepsilon^n + \sum_i y_{2i}^n$. Then for all i and k, y_{ki}^n/Y_k^n lies in the compact interval $[0,1]$ and thus has a limit point. Let n_m be a subsequence such that as $m \to \infty$, for all i we have $y_{1i}^{n_m}/Y_1^{n_m} \to \alpha_i$, and $y_{2i}^{n_m}/Y_2^{n_m} \to \beta_i$. Note that there exist a user i of type 1 and user j of type 2 such that $\alpha_i \leq 1/N_1$, and $\beta_j \leq 1/N_2$.

Taking subsequences again if necessary, we also assume that for each user i of type k, $y_{ki}^{n_m}$ converges as $m \to \infty$ (as this sequence take values in the compact strategy space of user i). Suppose that $Y_1^{n_m} \to 0$ and $Y_2^{n_m} \to 0$. Then:

$$
\begin{aligned}
u'_{1i}\left(\frac{y_{1i}^{n_m}}{Y_1^{n_m}} Y_2^{n_m}\right) \cdot \left(1 - \frac{y_{1i}^{n_m}}{Y_1^{n_m}}\right) &\to u'_{1i}(0) \cdot (1 - \alpha_i) > 1; \\
u'_{2j}\left(\frac{y_{2j}^{n_m}}{Y_2^{n_m}} Y_1^{n_m}\right) \cdot \left(1 - \frac{y_{2j}^{n_m}}{Y_2^{n_m}}\right) &\to u'_{2j}(0) \cdot (1 - \beta_i) > 1.
\end{aligned}
$$

These conditions together with the optimality conditions (3.13)-(3.15) imply that there exists k such that $Y_1^k/Y_2^k > 1$ and $Y_2^k/Y_1^k > 1$, which is a contradiction.

Now suppose that as $m \to \infty$, $Y_1^{n_m} \to 0$ but $Y_2^{n_m} \to c > 0$. Then from the optimality conditions (3.13)-(3.15), there exists M such that $y_{1i}^{n_m} > 0$ for all i and all $m \geq M$. Furthermore, there must exist a user i such that $y_{1i}^{n_m}/Y_1^{n_m} \to \alpha_i \leq 1/N_1$. Thus we have:

$$
\frac{Y_1^{n_m}}{Y_2^{n_m}} = u'_{1i}\left(\frac{y_{1i}^{n_m}}{Y_1^{n_m}} Y_2^{n_m}\right) \cdot \left(1 - \frac{y_{1i}^{n_m}}{Y_1^{n_m}}\right) \geq u'_{1i}\left(\frac{1}{N_1} Y_2^{n_m}\right)\left(\frac{N_1 - 1}{N_1}\right).
$$

The right hand side is strictly positive as $m \to \infty$, while the left hand side approaches zero. We conclude that we must have $Y_1^{n_m} \to c_1 > 0$ and $Y_2^{n_m} \to c_2 > 0$.

Suppose that as $m \to \infty$, $y_{ki}^{n_m} \to y_{ki}$ for each user i of type k. We will show that the resulting rates constitute a Nash equilibrium of the original game. If not, then there exists some user with a profitable deviation. Without loss of generality, let this be user i of type 1. Because u_{1i} is continuous, and $Y_k = \sum_j y_{kj} > 0$ for $k = 1, 2$, it is straightforward to check that for sufficiently large m user i will have a profitable deviation as well. This contradicts the assumption that $\vec{y}_1^m$ and $\vec{y}_2^m$ are a Nash equilibrium given ε^{n_m}; as a result, no such profitable deviation can exist. We conclude that $\vec{y}_1$ and $\vec{y}_2$ constitute a nonzero Nash equilibrium, as required.

■

Proof of Theorem 5: It suffices to show that $Y_k^N \to \infty$ as $N \to \infty$, for $k = 1, 2$. In this case, the second property of the theorem holds simply because each y_{ik} is bounded above by the upload rate constraint. Further, the optimality conditions (3.13)-(3.15) will imply both the first and third properties of the theorem: the first property follows because p_1^N cannot go to zero nor become unbounded if some users have positive rate; and the third property follows because the Nash optimality conditions are continuous as long as $Y_1 > 0$ and $Y_2 > 0$.

Suppose that Y_1^N remains bounded as $N \to \infty$; in this case, taking subsequences if necessary, we can assume that $Y_1^N \to c_1 < \infty$ as $N \to \infty$. Suppose also that $\sup_N Y_2^N = \infty$; then for at least one type 2 user i the analogous optimality conditions (3.14) or (3.15) hold. Taking subsequences if necessary, we have as $N \to \infty$, $Y_2^N / Y_1^N \to \infty$ and $Y_1^N / Y_2^N \to 0$, which contradicts either (3.14) or (3.15) for type 2 user i and the assumption $\sup_i u'_{1i}(0) < \infty$. Thus Y_2^N remains bounded as $N \to \infty$, so taking subsequences again if necessary, we assume that $Y_2^N \to c_2 < \infty$ as $N \to \infty$.

Without loss of generality, we can assume that $c_1/c_2 \leq 1$; otherwise we apply the subsequent argument to type 2 users. Again taking subsequences if necessary, we assume y_{ki}^N converges to y_{ki} for all i and k; this is straightforward, as the strategy space of each user is compact. Now since Y_1^N remains bounded for large N, there exists at least one user i of type 1 who has $y_{1i} = 0$. For such a user, taking limits in (3.13)-(3.14), we conclude we must have:

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

This contradicts our assumption that $\inf_i u'_{1i}(0) > 1$. Thus we conclude that in fact $Y_k^N \to \infty$

as $N \to \infty$, for $k = 1, 2$, as required; this establishes the theorem. ■

Proof of Lemma 10: Let $D(p) = u'^{-1}(p)$. As above, we normalize $p_2 = 1$. We first show the equivalence of Properties 1 and 2. Consider a type 1 user i; the argument for type 2 users is symmetric. For a given price $p_1 > 0$, his budget constraint is $x_{1i} \leq p_1 \cdot y_{1i}$ and will bind at any optimal solution (Lemma 9). Thus his objective function is $u(x_{1i}) - x_{1i}/p_1$. The nonnegativity constraint in (3.4) cannot bind, given Assumption 7. The optimal solution is given by $x_{1i}(p_1) = \min\{D(1/p_1), p_1 B\}$, so that $y_{1i}(p_1) = \min\{(1/p_1)D(1/p_1), B\}$.

Since u is strictly concave, $D(\cdot)$ is strictly decreasing. Thus Condition 2 in Definition 5 is satisfied; that is, $x_{1i}(p_1)$ strictly increases if p_1 strictly increases. Furthermore, if Property 2 in the statement of the lemma holds, then $y_{1i}(p_1)$ is nondecreasing in p_1, so Condition 1 of Definition 5 is also satisfied. Conversely, fix $p' > p > 0$, and choose $B > pD(p)$. Then if gross substitutes holds, we have $y_{1i}(1/p) \geq y_{1i}(1/p')$, so $pD(p) \geq p'D(p')$. Thus Property 1 and Property 2 are equivalent above.

Equivalence of the last two Properties follows by standard relationships between the derivatives of u' and u'^{-1}.

■

Proof of Lemma 11: Suppose there is a Nash equilibrium at which $y_{1i} < y_{1k}$ for some $i \neq k$. This means that $0 \leq y_{1i} < B$ and $0 < y_{1k} \leq B$. Then, if $Y_1 - y_{1k} > 0$,

$$\begin{aligned} u'\left(\frac{y_{1i}}{Y_1} Y_2\right) &\leq \frac{1}{Y_1 - y_{1i}} \frac{Y_1^2}{Y_2} \\ &< \frac{1}{Y_1 - y_{1k}} \frac{Y_1^2}{Y_2} \\ &\leq u'\left(\frac{y_{1k}}{Y_1} Y_2\right), \end{aligned}$$

where the first inequality follows from (3.13) and (3.14), and the last inequality follows from (3.14) and (3.15). Since $y_{1k}Y_1/Y_2 > y_{1i}Y_1/Y_2$, this contradicts the assumption that $u(\cdot)$ is strictly concave.

Now suppose that $Y_1 - y_{1k} = 0$. Then, $y_{1j} = 0$ for all $j \neq k$, while y_{1k} is strictly positive. If $Y_2 = 0$ the best response for any type 1 user is to upload zero, so if $y_{1k} > 0$ we must have $Y_2 > 0$. But no such equilibrium exists: user k will always want to decrease his upload rate. This shows that there cannot be an equilibrium at which $0 = y_{1i} < y_{1k}$. A symmetric argument holds for users of type 2. ■

Proof of Proposition 13: By Theorem 4, we know a Nash equilibrium exists. We show there exists at most one Nash equilibrium. Let (y_1, y_2) be Nash equilibrium upload rates, and first suppose that the upload rate constraint does not bind. Let $y_2/y_1 = a$. By substituting in (3.16) and (3.18), we obtain:

$$u'(ay_1) = \frac{N}{(N-1)a}; \;\; u'(y_1) = \frac{N}{N-1}a.$$

We only consider values of $a \in (N/(N-1)(1/u'(0)), Nu'(0)/(N-1))$, since only such values may yield strictly positive rates. The second equation gives $y_1 = u'^{-1}(Na/(N-1))$ and by substituting in the first equation, we conclude:

$$u'\left(au'^{-1}\left(\frac{Na}{N-1}\right)\right) = \frac{N}{(N-1)a}.$$

Clearly, $a = 1$ is a solution, which corresponds to $y_1 = y_2 = u'^{-1}(N/(N-1))$. Since $N/((N-1)a)$ is strictly decreasing in a, if $u'(au'^{-1}(Na/(N-1)))$ is nondecreasing in a, then $a = 1$ will be the unique solution. By Assumption 6, $u'(\cdot)$ is a strictly decreasing function, and from Lemma 10, $xu'^{-1}(x)$ is nonincreasing on $(0,\infty)$. Thus, there exists at most one Nash equilibrium with strictly positive rates at which the rate constraints do not bind.

From Lemma 10, we know that if the Optimization Problem of a user satisfies the gross substitutes property, then $xu'(x)$ is nondecreasing. We now show that if $xu'(x)$ is nondecreasing and the rate constraint binds for one type of user, then the rate constraint will also bind for the other type of user. Suppose that $y_1 = B$ and $y_2 < B$. Then, using (3.17) and (3.18), we have:

$$\frac{N}{N-1}B \leq y_2 u'(y_2) \leq Bu'(B) = \frac{N}{N-1}y_2 < \frac{N}{N-1}B,$$

which is a contradiction.

It remains to show that if $y_1 = y_2 = B$ is a Nash equilibrium, then $u'^{-1}(N/(N-1)) \geq B$.

Indeed, if $y_1 = y_2 = B$ is a Nash equilibrium, then from (3.17) and (3.19), $u'(B) \geq N/(N-1)$. But then, from Assumption 6, $B \leq u'^{-1}(N/(N-1))$. ■

Proof of Proposition 14: First suppose that there exists a competitive equilibrium for the PP scheme with equilibrium prices $(p_i^*, i \in N)$ for peers and the (p_ℓ^*) for links in the network. Then, the peer price vector $(p_{if}^*, i \in N, f \in F_i)$ with $p_{if}^* = p_i^*\ \forall f \in F_i$ and the link price vector (p_ℓ^*) constitute a competitive equilibrium price vector for the PFP scheme. This holds because it gives rise to the same demand and supply as $(p_i^*, i \in N)$ and (p_ℓ^*).

Now suppose that there exists a competitive equilibrium for the PFP scheme with prices $(p_{if}^*, i \in N, f \in F_i)$ and (p_ℓ^*). For each peer i, set $p_i^* = \max_{f \in F_i}$. Peer i only "supplies" his most expensive files at equilibrium. Hence, the prices (p_i^*) and (p_ℓ^*) yield the same demand and supply as (p_{if}^*) and (p_ℓ^*) for each peer. ■

Chapter 4

Bilateral and Multilateral Exchange

4.1 Introduction

This chapter provides a formal comparison of peer-to-peer system designs with bilateral barter, such as BitTorrent, and multilateral exchange of content enabled by a price mechanism to match supply and demand. Multilateral exchange in the context of a peer-to-peer system was studied in Chapter 3.

BitTorrent is a popular peer-to-peer protocol, which accounts for a large percentage of all Internet traffic. According to the BitTorrent protocol, each peer splits its available upload rate among peers from which it gets the highest download rates. As a result, the total download rate of a peer is a nondecreasing function of his total upload rate, and peers are incentivized to contribute. Therefore, the BitTorrent protocol incentivizes users in a bilateral basis: an increase in the upload rate to one peer may increase the download rate from that particular peer.

The difficulties of bilateral exchange (or barter) in an economy have been long known. Several inconveniences arise in the absence of money, the most important being the improbability of coincidence between persons wanting and possessing [40]. In modern economies, the aforementioned difficulty is eliminated by the use of *money*. Money can enable multilateral exchange by serving as a medium of exchange and a common measure of value. Even though modern societies take the use of money for granted, this is not the case in peer-to-peer systems, partly because of the associated system design complexity.

Peer-to-peer systems could potentially also use *market-based multilateral exchange* to match user demand for content to available supply at other users in the system. This can be done by using virtual currency and assigning a budget to each user that decreases when downloading and increases when uploading. Monetary incentives in a virtual currency have been previously proposed to incentivize uploading in peer-to-peer systems [34, 70, 66, 6]; however such designs are usually more complex than bilateral protocols and are not widespread. Thus, there is a significant tradeoff: *bilateral exchange without money is simple; on the other hand multilateral exchange allows more users to trade*. In this paper, we provide a formal comparison of two peer-to-peer system designs: bilateral barter systems, such as BitTorrent; and a market-based exchange of content enabled by a price mechanism to match supply and demand.

We start in Section 4.2 with a fundamental abstraction of content exchange in systems like BitTorrent: *exchange ratios*. The exchange ratio from one user to another gives the download rate received per unit upload rate. Exchange ratios are a useful formal tool because they allow us to define and study the equilibria of bilateral exchange. In a bilateral equilibrium each user optimizes with respect to exchange ratios. In Section 4.3 we also define competitive equilibria (corresponding to multilateral exchange), where users optimize with respect to prices.

In Section 4.4 we compare bilateral and multilateral peer-to-peer systems through the allocations that arise at equilibria. A competitive equilibrium allocation is always Pareto efficient, while bilateral equilibria may be inefficient. Our main result in this section is that a bilateral equilibrium allocation is Pareto efficient if and only if it is a competitive equilibrium allocation. This result provides formal justification of the efficiency benefits of competitive equilibria. The proof exploits an interesting connection between equilibria and Markov chains: an important step of the proof is to show that Pareto efficiency of a bilateral equilibrium rate allocation implies reversibility of an appropriately defined Markov chain, and that this chain has an invariant distribution that corresponds to a price vector of a competitive equilibrium.

In Section 4.5 we perform a quantitative comparison of bilateral and multilateral exchange. As discussed in [40], "there may be many people wanting, and many possessing those things wanted; but to allow of an act of barter, there must be a double coincidence, which will rarely happen." We quantify how rarely this double coincidence of wants occurs under different assumptions on the popularity of files in the system. We first perform an asymptotic analysis assuming that file popularity follows a power law and study two extreme regimes. We find that asymptotically all users are able to trade bilaterally when the file popularity is very concentrated. On the other hand, multilateral exchange performs significantly better than bilateral when the file popularity is not concentrated. We complement our theoretical analysis by studying file popularity from a BitTorrent dataset. Although bilateral equilibria may in general be inefficient, the gap between bilateral and multilateral exchange can be narrowed significantly if each user shares a sufficient number of files (in practice as small as ten).

Our work is related to the study of equilibria in economies where not all trades are allowed. Kakade et. al. introduce a graph-theoretic generalization of classical Arrow-Debreu economics, in which an undirected graph specifies which consumers or economies are permitted to engage in direct trade [42]. However, the inefficiencies of bilateral exchange do not arise in their model. Finally, the monetary economics literature has studied how money reduces the double coincidence problem. The implementation of a competitive equilibrium is a central theme in this literature. The superiority of monetary exchange has been studied [67], and dynamics of bilateral trading processes have been considered [55, 27]. The transactions role of money is surveyed in [54].

4.2 Exchange Ratios in Bilateral Protocols

Many peer-to-peer protocols enable exchange on a *bilateral* basis between users: a user i uploads to a user j if and only if user j uploads to user i in return. Of course, such an exchange is only possible if each user has something the other wants. The foremost examples of such a protocol are BitTorrent and its variants. While such protocols are traditionally studied solely through the rates that users obtain, in this section we provide an interpretation of these protocols through *exchange ratios*. As exchange ratios can be interpreted in terms of prices, these ratios allow us to compare bilateral barter-based peer-to-peer systems with multilateral price-based peer-to-peer systems.

Let r_{ij} denote the rate sent from user i to user j at a given point in time in a bilateral peer-to-peer protocol. We define the *exchange ratio* between user i and user j as the ratio $\gamma_{ij} = r_{ji}/r_{ij}$; this is the download rate received by i from j, per unit of rate uploaded to j. By definition, $\gamma_{ij} = 1/\gamma_{ji}$. Clearly, a rational user i would prefer to download from users with which he has higher exchange ratios.

The exchange ratio has a natural interpretation in terms of prices: an equivalent story emerges if we assume that users charge each other for content in a common monetary unit, but that all transactions are *settlement-free*, i.e., no money ever changes hands. In this case, if user i charged user j a price p_{ij} per unit rate, the exchange of content between users i and j must satisfy:

$$p_{ij}r_{ij} = p_{ji}r_{ji}$$

We refer to p_{ij} as the *bilateral price* from i to j. Note that the preceding condition thus shows the exchange ratio is equivalent to the ratio of bilateral prices: $\gamma_{ij} = p_{ij}/p_{ji}$ (as long as the prices and rates are nonzero).

What is the exchange ratio for BitTorrent? A user splits its upload capacity equally among those users in its active set from which it gets the highest download rates. Let α be the size of the active set. Suppose all rates r_{kj} that user j receives from users $k \neq i$ are fixed and let R_j^α be the α-th highest rate that j receives. Let B_j be the upload capacity of user j. Then, r_{ji} depends on r_{ij}. In particular,

$$r_{ji} = \begin{cases} B_j/\alpha & \text{if } r_{ij} > R_j^\alpha \\ 0 & \text{otherwise} \end{cases}$$

Thus for BitTorrent, the exchange ratio is $\gamma_{ij} = B_j/(\alpha \cdot r_{ij})$ if user i is in the active set, and zero otherwise. Note that the exchange ratios $\gamma_{i_1,j}$ and $\gamma_{i_2,j}$ may be different for two users i_1, i_2 in j's active set.

The exchange ratio γ_{ij} *decreases* with r_{ij} as long as user i is in user j's active set (in which case r_{ji} is constant). Hence, a strategic user i would prefer to choose r_{ij} as small as possible while remaining in j's active set. This behavior is exactly the approach taken by the BitTyrant [57] variation on BitTorrent. In fact, if all users follow this policy, then $r_{ij} = R_j^\alpha$ for all users i in j's active set. Note that in this case, $\gamma_{ij} = B_j/(\alpha \cdot R_i^\alpha)$. Thus, user j has the same exchange ratio to all users i with which he bilaterally exchanges content.

The preceding discussion highlights the fact that the rates in a bilateral peer-to-peer

Bilateral User Optimization:		**Multilateral User Optimization:**	
maximize	$v_i(\vec{x}_i, y_i)$	maximize	$v_i(\vec{x}_i, y_i)$
subject to	$x_{if} = \sum_j r_{jif}, \forall f$	subject to	$x_{if} = \sum_j r_{jif} \forall f$
	$y_i = \sum_{j,f} r_{ijf}$		$y_i = \sum_{j,f} r_{ijf}$
	$\vec{r} \in X$		$\vec{r} \in X$
	$\sum_f r_{jif} \leq \gamma_{ij} \sum_f r_{ijf} \forall j$		$\sum_{j,f} p_j r_{jif} \leq p_i \sum_{j,f} r_{ijf}$

Figure 4.1: Optimization problems for price-based exchange. The two optimization problems differ only in the last constraint (budget constraint).

system can be interpreted via exchange ratios. Thus far we have assumed that *transfer rates* are given, and exchange ratios are computed from these rates. In the next section, we turn this relationship around: we explicitly consider an abstraction of bilateral peer-to-peer systems where users react to given exchange ratios, and compare the resulting outcomes to price-based multilateral exchange.

4.3 Bilateral and Competitive Equilibria

Motivated by the discussion in the preceding section, this section uses exchange ratios to define bilateral equilibria. In Section 3.2 we defined competitive equilibrium where peers optimize with respect to given prices. Similarly, at a bilateral equilibrium peers optimize with respect to exchange ratios (instead of prices).

As in Chapter 3, in the formal model we consider, a set of users U shares a set of files F. User i has a subset of the files $S_i \subseteq F$, and is interested in downloading files in $T_i \subseteq F - S_i$. Throughout, we use r_{ijf} to denote the rate at which user i uploads file f to user j. We then let $x_{if} = \sum_j r_{jif}$ be the rate at which user i downloads file f. We denote the vector of download rates for user i by $\vec{x}_i = (x_{if}, f \in T_i)$. Let $y_i = \sum_{j,f} r_{ijf}$ be the total upload rate of user i. We measure the desirability of a download vector and an upload rate to user i by a *utility function* Assumption 3.

Each user is assumed to have a constraint on the available upload rate; let B_i denote this upper bound for user i. We assume that users do not face any constraint on their download rate; this is consistent with most end user connections today, where upload capacity is far exceeded by download capacity.

Let

$$\mathcal{X} = \left\{ \vec{r} : \vec{r} \geq 0;\ r_{kjf} = 0 \text{ if } f \notin S_k;\ \sum_{j,f} r_{ijf} \leq B_i\ \forall i \in U \right\}$$

be the set of feasible rate vectors. In particular, this ensures that (1) all rates are non-negative, (2) users only upload files they possess, and (3) each user does not violate his bandwidth constraint.

We start by considering users' behavior in bilateral schemes, given a vector of exchange ratios $(\gamma_{ij}, i, j \in U)$. User i solves the bilateral optimization problem given in Figure 4.1.[1] By contrast, in a *multilateral price-based exchange*, the system maintains one price per user, and users optimize with respect to these prices.[2] We denote the price of user i by p_i. Figure 4.1 also gives the user optimization problem in multilateral price-based exchange. Note that the first three constraints (giving download and upload rates and ensuring that the rate allocation is feasible) are identical to the bilateral user optimization. Only the last constraint is different. While the bilateral exchange implicitly requires user i to download only from those users to whom he uploads, no such constraint is imposed on multilateral exchanges: user i accrues capital for uploading, and he can spend this capital however he wishes for downloading.

For bilateral (resp., multilateral) exchange, an *equilibrium* is a combination of a rate allocation vector and an exchange ratio vector (resp., price vector) such that all users have solved their corresponding optimization problems. In this case, the exchange ratios (resp., prices) have exactly aligned supply and demand: for any i, j, f, the transfer rate r_{ijf} is simultaneously an optimal choice for both the uploader i and downloader j. In the next two subsections we provide formal definitions of equilibria for both models.

4.3.1 Bilateral Equilibrium

Definition 7 *The rate allocation $\vec{r}^* \in \mathcal{X}$ and the exchange ratios $\vec{\gamma}^* = (\gamma^*_{ij}, i, j \in U)$ with $\gamma^*_{ij} \cdot \gamma^*_{ji} = 1$ for all i, j, constitute a Bilateral Equilibrium (BE) if for each user i, $\vec{r}^*$ solves*

[1] Note that we allow users to bilaterally exchange content over multiple files, even though this is not typically supported by swarming systems like BitTorrent; in BitTorrent a single file is split into subpieces called chunks, and users exchange chunks.

[2] This is equivalent to having one price per file in our setting [5].

the Bilateral User Optimization problem given exchange ratios $\vec{\gamma}^*$.

Definition 7 requires that (1) all users have optimized with respect to the exchange ratios, and (2) the market clears. Even though the market clearing condition is not explicitly stated, it is implicitly required, since the same vector $\vec{r}^*$ is an optimal solution of the bilateral optimization problems of all users.

We do not expect a BE to exist in general. For example, this is trivially the case if no pair of users has reciprocally desired files; i.e., if for every pair i, j either $S_i \cap T_j = \emptyset$ or $S_j \cap T_i = \emptyset$. To show existence we assume that every user can find every file he desires through bilateral trade. This is formalized in Assumption 8.

Assumption 8 *For every user i and every file $f \in T_i$ there exists a user j such that $f \in S_j$ and $T_j \cap S_i \neq \emptyset$.*

Proposition 15 *If Assumptions 3 and 8 hold, then a BE exists.*

4.3.2 Competitive Equilibrium

For completeness we also give the definition of CE. In particular, we rewrite Definition 3 using the notation of this section and one price per peer.

Definition 8 *The rate allocation $\vec{r}^*$ and the user prices $\vec{p} = (p_i^*, i \in U)$ with $p_i^* > 0$ for all $i \in U$ constitute a Competitive Equilibrium (CE) if for each user i, $\vec{r}^*$ solves the Multilateral User Optimization problem given prices $\vec{p}$.*

Similarly to Definition 7, Definition 8 requires that (1) all users have optimized with respect to prices, and (2) the market clears. Even though the market clearing condition is not explicitly stated, it is implicitly required, since the same vector $\vec{r}$ is used in the optimization problems of all users.

Proposition 16 *If Assumptions 3 and 5 hold, then there exists a CE.*

4.4 Efficiency of Equilibria

This section rigorously analyzes the efficiency properties of bilateral and multilateral exchange. We assume users explicitly react to exchange ratios or prices, and we compare the schemes through their resulting equilibria.

A CE allocation is always Pareto efficient, i.e., there is no way to increase the utility of some user without decreasing the utility of some other user; this is the content of the first fundamental theorem of welfare economics [50]. For completeness, we include the result here.

Theorem 6 *If the rate allocation $\vec{r}^*$ and the user prices $(p_i^*, i \in U)$ with $p_i^* > 0$ for all $i \in U$ constitute a CE, then the allocation $\vec{r}^*$ is Pareto efficient.*

A BE, on the other hand, may not be Pareto efficient, as the following example shows.

Example 7 *Consider a system with n users and n files, for $n > 2$. Each user i has file f_i and wants files f_{i+1} and f_{i-1}. The utility of user i is $v_i(x_{i,f_{i-1}}, x_{i,f_{i+1}}, y_i) = x_{i,f_{i-1}} + 4x_{i,f_{i+1}} + \ln(2 - y_i)$, i.e., user i wants the files of both user $i+1$ and user $i-1$, but derives a higher utility from the file of user $i+1$.*

We first consider a symmetric BE with exchange ratios $\gamma_{i,i+1}^* = 2$ and $\gamma_{i,i-1}^* = 1/2$. The equilibrium rates are $r_{i-1,i}^* = 1$ and $r_{i+1,i}^* = 1/2$, and the download rates are $x_{i,f_{i-1}}^* = 1$ and $x_{i,f_{i+1}}^* = 1/2$. The utility of each user i is $3 - \ln(2) \approx 2.3$. On the other hand, prices $p_i^* = 1$ for all i, and rates $r_{i+1,i}^* = 1.75$, $r_{i-1,i}^* = 0$ constitute a CE. The utility of each user is $7 - \ln(4) \approx 5.61$, i.e., significantly larger than the utility of a user at the BE. This demonstrates that the BE allocation is not Pareto efficient.

The previous examples show that BE may not be Pareto efficient. We next provide an example of a BE rate allocation that is Pareto efficient.

Example 8 *Consider a system with n users and n files, for $n > 2$. Each user i has file f_i and wants files f_{i+1} and f_{i-1}. The utility of user i is $v_i(x_{i,f_{i-1}}, x_{i,f_{i+1}}, y_i) = x_{i,f_{i-1}} + x_{i,f_{i+1}} + \ln(2 - y_i)$.*

We consider a symmetric BE with exchange ratios $\gamma_{i,i+1}^* = 1$ and $\gamma_{i,i-1}^* = 1$. The equilibrium rates are $r_{i-1,i}^* = 1/2$ and $r_{i+1,i}^* = 1/2$. The BE rate allocation is Pareto efficient. In *particular, it corresponds to a CE: prices $p_i^* = 1$ for all i, and rates $r_{i+1,i}^* = 1/2$, $r_{i-1,i}^* = 1/2$ constitute a CE.*

BE may be inefficient, while CE always have Pareto efficient allocations (Theorem 6). In Example 8 the BE rate allocation is Pareto efficient and corresponds to a CE. Our main result is that *a BE allocation is a CE allocation if and only if it is Pareto efficient*. In particular, if a BE allocation is Pareto efficient, then there exist "supporting prices", i.e., prices such that the BE rate allocation is optimal for the multilateral optimization problem of each user. Informally, Pareto efficiency represents the "gap" between BE and CE.

Proposition 17 *Assume that for every user i and any fixed $\vec{x}_i$, $v_i(\vec{x}_i, y_i) \to -\infty$ as $y_i \to B_i$. Let $(\vec{r}^*, \vec{\gamma}^*)$ be a BE. The rate allocation $\vec{r}^*$ is Pareto efficient if and only if there exists a price vector $\vec{p}$ such that $\vec{r}^*$ and $\vec{p}$ constitute a CE.*

Proposition 17 assumes that $v_i(\vec{x}_i, y_i) \to -\infty$ as $y_i \to B_i$ for every user i and every fixed $\vec{x}_i$. This assumption is needed so that upload capacity constraints do not bind at the BE. This is a reasonable assumption for a peer-to-peer setting, since we do not expect users to use all their upload capacity. We note that if the upload capacity constraint binds for some users, then there may exist Pareto efficient BE that do not correspond to CE, simply because users have already "maxed out" their available upload capacity.

We provide an overview of the proof of Proposition 17, which demonstrates an interesting connection between equilibria and Markov chains; the details of the proof are provided in the appendix. From a BE rate allocation $\vec{r}^*$ we construct a transition rate matrix $\vec{Q}$ such that $Q_{ij} = \sum_f r_{ijf}^*$ if $i \neq j$; and $Q_{ii} = -\sum_{j,f} r_{ijf}^*$. We first observe that $\vec{\pi}\vec{Q} = 0$ implies that the multilateral budget constraint is satisfied with price vector $\vec{\pi}$; therefore for any invariant distribution $\vec{\pi}$, $\vec{r}^*$ is feasible for the multilateral optimization problem of every user when prices are equal to $\vec{\pi}$. We then show that there exists an invariant distribution of $\vec{Q}$ (say $\vec{p}$) such that $\vec{r}^*$ is an optimal solution of the multilateral optimization problem of each user when the prices are equal to $\vec{p}$. We conclude that $\vec{r}^*$ and $\vec{p}$ constitute a CE.

A key step of the proof is to show that Pareto efficiency of $\vec{r}^*$ implies *reversibility* of $\vec{Q}$. Let $\vec{\pi}$ be an invariant distribution of $\vec{Q}$. $\vec{Q}$ is reversible if and only if $\gamma_{ij}^* = \pi_i / \pi_j$ for all pairs of users i and j that trade at the BE. This means that if $\vec{Q}$ is reversible, then $\vec{r}^*$ solves the multilateral optimization problem of each user given prices $\vec{\pi}$ *if the user is restricted* *to trade with peers it trades at the BE*. The matrix corresponding to the BE allocation of Example 7 is not reversible, which implies that the BE allocation is not Pareto efficient. On the other hand, the matrix corresponding to the BE allocation of Example 8 is reversible, and the BE allocation is Pareto efficient and corresponds to a CE allocation.

4.5 Quantitative Comparison

Bilateral exchange may be particularly restrictive because a pair of users can exchange only if each has a file that the other wants. On the other hand, allowing multilateral exchange significantly increases the number of possible exchanges, and potentially increases the number of users that can trade, but is also associated with increased complexity. In this section we compare bilateral and multilateral exchange through the corresponding percentages of users that can trade. Though distinct from Pareto efficiency, this metric provides quantitative insight into the comparison of the two types of exchange. We characterize regimes where bilateral exchange performs very well with respect to this metric, and for which — as a result — it may not be worth the effort to use multilateral exchange.

We first perform an asymptotic analysis assuming that file popularity follows a power law. We find that asymptotically all users are able to trade bilaterally when the file popularity is very concentrated. We complement our theoretical analysis by studying file popularity from a BitTorrent dataset. We find that a very large percentage of users is able to trade bilaterally if each user is sharing a sufficiently large number of files; e.g., over 96% of users can trade if each user has 10 files (for the number of users in the dataset).

We start by formally defining the quantities we compare. For a given peer-to-peer system, we define the *system profile* to consist of the specification of which files each user possesses and desires, i.e., $\mathcal{P} = \{T_i, S_i, i \in U\}$. For simplicity, we consider settings where each user is interested in downloading one file, i.e., $|T_i| = 1$ for all $i \in U$. This assumption significantly simplifies the analysis, since we do not need to consider how a user's utility function depends on different files. We thus abstract from specific utility functions and focus on how much bilateral exchange restricts trade.

We say that user i *can trade bilaterally under* $\mathcal{P}$ if there exists some user j such that $S_i \cap T_j \neq \emptyset$ and $S_j \cap T_i \neq \emptyset$, i.e., if i and j have reciprocally desired files. Given a system profile $\mathcal{P}$, let $\rho_{BE}(\mathcal{P})$ be the percentage of users that cannot trade bilaterally. We note that $\rho_{BE}(\mathcal{P})$ is equal to the percentage of users that need to be removed from the system so that a BE exists for $\mathcal{P}$ (if Assumption 3 holds). The condition $\rho_{BE}(\mathcal{P}) = 0$ is equivalent to Assumption 5 when each user possesses one file, as we assume in this section.

Similarly, we say that user *i can trade multilaterally under* $\mathcal{P}$ if there exist users $k_1, k_2, \ldots, k_n$ such that $S_{k_j} \cap T_{k_{j+1}} \neq \emptyset$ for $j = 1, \ldots, n$; $S_i \cap T_{k_1} \neq \emptyset$ and $S_{k_n} \cap T_i \neq \emptyset$. In words, user i is able to trade multilaterally if and only if there exists a cycle of users starting (and ending) at i such that each user possesses a file that is desired by the next user in the cycle. Clearly, if user i can trade bilaterally under $\mathcal{P}$, then he can also trade multilaterally under $\mathcal{P}$. Let $\rho_{ME}(\mathcal{P})$ be the percentage of users that cannot trade multilaterally. We note that $\rho_{ME}(\mathcal{P})$ is equal to the percentage of users that need to be removed from the system so that a CE exists for $\mathcal{P}$ (if Assumption 3 holds). The condition $\rho_{ME}(\mathcal{P}) = 0$ is weaker than Assumption 5; however, it is sufficient for CE existence when each user possesses one file, as we assume in this section.

We assume that the system profile $\mathcal{P}$ is chosen according to some distribution that depends on the popularity of different files. We are interested in comparing the expected values of $\rho_{BE}(\mathcal{P})$ and $\rho_{ME}(\mathcal{P})$. In Section 4.5.1 we consider a large system and perform an asymptotic analysis. In Section 4.5.2 we use a BitTorrent dataset to derive file popularity distributions, and then compare the expected values of $\rho_{BE}(\mathcal{P})$ and $\rho_{ME}(\mathcal{P})$ through simulations.

4.5.1 Asymptotic Analysis

In this section we theoretically study bilateral and multilateral exchange in large systems. We compare the two types of exchange through the expected percentages of users that cannot trade. We focus on large systems, and consider the asymptotic regime as the number of files and users in the system becomes large.

We assume the files that users possess and desire are drawn from a Zipf file popularity distribution independently and identically for each user. Our motivation to study this distribution comes from the fact that Zipf's law has been observed in many settings, and has been suggested as a good model for file popularity (e.g., [17]). Zipf's law states that the popularity of the r-th largest occurrence is inversely proportional to its rank. We adjust this definition to our setting.

Definition 9 *File popularity has a* Zipf distribution *with parameter s if the r-th most popular file has probability proportional to r^{-s}.*

For example, if S_i and T_i are singletons, a user desires the i-th most popular file and possesses the j-th most popular file with probability $(i \cdot j)^{-s} / (\sum_{k' \neq k} (k' \cdot k)^{-s})$.

Note that $s = 0$ corresponds to the uniform distribution. On the other hand, as s increases the distribution becomes more concentrated.

We are interested in the expected percentage of users that cannot trade. This is a function of the number of users N, the number of files K and the Zipf exponent s. Let $\overline{\rho}_{BE}(K,N,s)$ and $\overline{\rho}_{ME}(K,N,s)$ be the expected percentages of users that cannot trade bilaterally and multilaterally. In particular, $\overline{\rho}_{BE}(K,N,s)$ (resp., $\overline{\rho}_{ME}(K,N,s)$) is the expected value of $\rho_{BE}(\mathcal{P})$ (resp., $\rho_{ME}(\mathcal{P})$) over system profiles.

We consider a sequence of peer-to-peer systems indexed by N. The Nth system has N users and $K(N)$ files, where $K(N)$ is a nondecreasing function of N. The function $K(N)$ represents how the number of files scales with the number of users. We study an asymptotic regime where $N \to \infty$.

Since the number of users that cannot trade bilaterally is always greater than or equal to the number of users that cannot trade multilaterally, we have $\overline{\rho}_{BE}(K,N,s) - \overline{\rho}_{ME}(K,N,s) \geq 0$. The following propositions imply that in a large system $\overline{\rho}_{BE}(K,N,s) - \overline{\rho}_{ME}(K,N,s)$ may be significant when $s = 0$, but is always negligible when $s > 1$.

Proposition 18 *Assume $s = 0$, i.e., files are chosen uniformly. Moreover, $|S_i| = \sigma$ and $|T_i| = 1$ for all $i \in U$.*

1. *If $K(N) \geq \sigma\sqrt{N}$ for large N, then there exists $\overline{N}$ such that $\overline{\rho}_{BE}(K(N),N,0) \geq 1/e$ for $N \geq \overline{N}$.*

2. *If there exists $\varepsilon > 0$ such that $K(N) \leq \sigma N^{1/2-\varepsilon}$ for large N, then $\overline{\rho}_{BE}(K(N),N,0) \to 0$ as $N \to \infty$.*

3. *If there exists $\varepsilon > 0$ such that $K(N) \leq \sigma N^{1-\varepsilon}$ for large N, then $\overline{\rho}_{ME}(K(N),N,0) \to 0$ as $N \to \infty$.*

The case $K(N) \in [\sigma\sqrt{N}, \sigma N^{1-\varepsilon}]$ is of particular interest. According to Proposition 18 in this case $\overline{\rho}_{BE}(K(N),N,0) \geq 1/e \approx 0.37$ for large N, while $\overline{\rho}_{ME}(K(N),N,0) \to 0$ as $N \to \infty$. In words, when the system is large more than one third of the users cannot trade bilaterally, while almost all users can trade multilaterally. We conclude that if $\sigma N^{1-\varepsilon} \geq K(N) \geq \sigma\sqrt{N}$ and files are chosen uniformly, then multilateral exchange performs significantly better than bilateral exchange in terms of the number of users that can trade.

Proposition 19 *If $s > 1$, then $\overline{\rho}_{BE}(K(N),N,s) \to 0$ as $N \to \infty$ for any nondecreasing $K(N)$.*

Since $\overline{\rho}_{ME}(K(N),N,s) \leq \overline{\rho}_{BE}(K(N),N,s)$, we conclude that if files are chosen according to a Zipf distribution with $s > 1$ then both $\overline{\rho}_{BE}(K(N),N,s) \to 0$ and $\overline{\rho}_{ME}(K(N),N,s) \to 0$ as $N \to \infty$. We note that this result does not depend on the number of files that users possess. When $s > 1$, bilateral exchange performs very well asymptotically even if each user only possesses one file.

This is an interesting result: even though bilateral exchange significantly restricts trade compared to multilateral exchange, in expectation almost all users can trade under both types of exchange when the system is large and file popularity follows a Zipf distribution with exponent $s > 1$. The intuition behind this result is that when s is large, the popularity distribution is more concentrated, i.e., the most popular files are chosen with relatively high probability. As a result, for any user i both T_i and S_i probably consist of one of the most popular files, and it is more likely that there exists a user j, such that i and j have reciprocally desired files.

4.5.2 Data Analysis

In this section we quantitatively compare bilateral and multilateral exchange using a BitTorrent dataset collected by Piatek et. al. from the University of Washington [58].[3] We find that a significant percentage of users cannot trade bilaterally when each user is sharing one file; however, the percentage becomes negligible as peers share more files.

The dataset consists of 1,364,734 downloads, 679,523 users and 7,323 files. We use the number of downloads of each file in the dataset to estimate the probability that a given file is

[3] We are grateful to Piatek et. al. for providing the dataset.

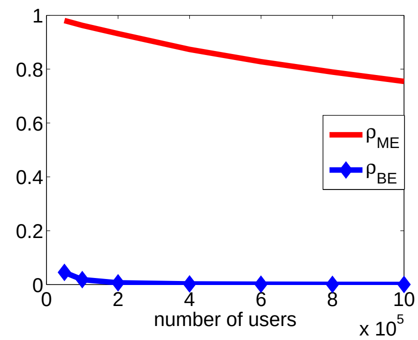

Figure 4.2: Percentages of users (from simulations) that cannot trade bilaterally and multilaterally when users desire and possess one file, i.e., $|T_i| = |S_i| = 1$ for all i. The horizontal axis shows the number of users in the system.

selected. We then use the estimated probabilities to generate system profiles and compute the percentages of users that cannot trade bilaterally and multilaterally. We assume that there are 7,323 files with the given distribution, and vary the number of users in the system.

We first assume that each user possesses and desires exactly one file, i.e., $|T_i| = |S_i| = 1$ for every $i \in U$. Figure 4.2 shows the percentages of users that cannot trade bilaterally and multilaterally from simulations for various numbers of users in the system.[4] We observe that the percentage of users that cannot trade bilaterally is significantly larger than the percentage of users that cannot trade multilaterally. Moreover, the latter is close to 0, indicating that most users are able to trade multilaterally. Finally, as the number of users increases, the percentages of users that can trade increase for both bilateral and multilateral exchange. For instance, when there are 200,000 users in the system, about 93% cannot trade bilaterally while only 0.7% cannot trade multilaterally. When the number of users increases to 1,000,000, about 75% cannot trade bilaterally and 0.06% cannot trade multilaterally.

We next assume each user desires one file and possesses multiple files. As the number of files that each user has increases, the number of possible trades increases, and as a result the percentage of users that can trade bilaterally increases. In Figure 4.3 we show the

[4] The algorithm we use to compute ρ_{BE} is exact: for every user i we check whether there is some user j such that i and j have reciprocally desired files. Computing the exact value of ρ_{ME} for a large system seems computationally intractable. Therefore, we use an approximation algorithm to compute ρ_{ME}: we recursively remove users that possess files not desired by others or desire files not possessed by others. Simulations for small numbers of users suggest that this algorithm provides a very good approximation for ρ_{ME}.

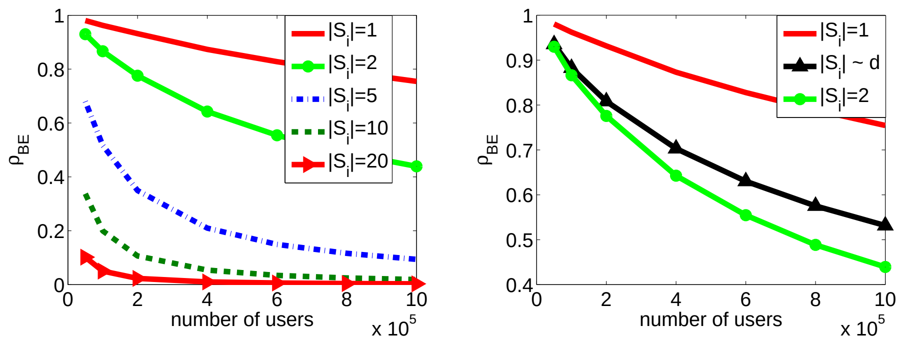

Figure 4.3: Percentages of users (from simulations) that cannot trade bilaterally when each user desires one file ($|T_i| = 1$) and possesses multiple files. The legend shows $|S_i|$ for each line. The horizontal axis shows the number of users in the system. In the left figure all users possess the same number of files ($|S_i| \in \{1, 2, 5, 10, 20\}$). In the right figure the number of files that a user possesses ($|S_i|$) is drawn from the dataset distribution (denoted by "d"). We also include the cases $|S_i| = 1$ and $|S_i| = 2$ for all users in the right figure.

percentages of users that cannot trade bilaterally when each user desires one file ($|T_i| = 1$) and possesses multiple files (from simulations).

For the first figure in Figure 4.3 we assume that all users possess the same number of files, i.e., $|S_i| = |S_j|$ for all $i, j \in U$. The figure shows the percentages of users that cannot trade multilaterally for various values of $|S_i|$. The case $|S_i| = 1$ has already been considered in Figure 4.2. In Figure 4.3 we also consider $|S_i| \in \{2, 5, 10, 20\}$. We observe a significant decrease in the percentage of users that cannot trade when $|S_i|$ increases from 1 to 20. We can illustrate this by considering the minimum required number of users in the system so that at most 10% are not able to trade: at least 1,000,000 users are required when each user has 5 files; at least 200,000 users are needed in the system when each possesses ten files; at least 50,000 users are needed when each user possesses 20 files. We observe that there is a significant decrease in the required number of users. Moreover, when each user possesses 20 files, the percentage of users that cannot trade bilaterally is very small when there are more than 200,000 users in the system: 2.3% cannot trade with 200,000 users; only 0.28% cannot trade with 1,000,000 users.

Our simulations up to now have assumed that all users in the system possess the same number of files, i.e., $|S_i| = |S_j|$ for all i, j. We next assume that the number of files that users possess vary across different users. We are interested in whether the percentage of users that can trade bilaterally increase as the variance of the distribution of $|S_i|$ increases (assuming that the mean remains the same). At first it may seem plausible that users with very large $|S_i|$ would be able to accommodate a lot of trades and as a result ρ_{BE} should increase as the $|S_i|$'s become more dispersed. However, this is not the case as we discuss next.

We infer the distribution of $|S_i|$ from the dataset.[5] The mean value of $|S_i|$ in the dataset is 2.0084. Therefore, we are interested in whether ρ_{BE} increases compared to the case that $|S_i| = 2$ for all i. The second figure in Figure 4.3 shows the percentage of users that cannot trade bilaterally when the number of files that users possess ($|S_i|$) follows the distribution from the dataset. For comparison, we also show the percentages that cannot trade when $|S_i|$ is equal to 1 or 2 for all users. We observe that when $|S_i|$ is drawn from the distribution, the percentages are between the case $|S_i| = 1$ and $|S_i| = 2$ — even though the expected value of $|S_i|$ is slightly greater than two. In particular, about 81% of users cannot trade bilaterally when $|S_i|$ is drawn from the distribution, while 77% cannot trade when $|S_i| = 2$ for all users. The percentage of users with large $|S_i|$ that cannot trade bilaterally significantly decreases: 26% of users with $|S_i| > 10$ cannot trade bilaterally. However, only 2% of the users belong in this category; 67 % of all users have $|S_i| = 1$ and 88% of them cannot trade bilaterally.

To further examine the conjecture that a more dispersed distribution of $|S_i|$ further restricts bilateral trade, we run simulations assuming that $|S_i| = 1$ with probability 0.5 and $|S_i| = 3$ with probability 0.5, so that the mean value of $|S_i|$ is equal to 2. This distribution is less dispersed than the distribution we obtain from the data, but more dispersed than when $|S_i| = 2$ for all users. Indeed, the percentage of users than cannot trade bilaterally when there are 200,000 users is 78%, which is greater than the percentage of users that cannot trade when $|S_2| = 1$ for all users (77%), but smaller than the percentage of users than cannot trade under the data distribution (80%).

4.6 Proofs

Proof of Proposition 15: We first define the concept of restricted BE and show that such an equilibrium always exists. We then use the exchange ratios of the restricted BE to construct

[5] We assume that the number of files that a user possesses is equal to the number of files he downloads.

a BE according to Definition 7.

The rate allocation $\vec{r}^*$ and the exchange ratios $\vec{\gamma}^*$ constitute a *restricted BE* if

1. $\gamma^*_{ij} = 0$ if $S_i \cap T_j = \emptyset$ or $S_j \cap T_i = \emptyset$;
   $\gamma^*_{ij} \cdot \gamma^*_{ij} = 1$ otherwise.

2. For each user i, $\vec{r}^*$ solves the Bilateral User Optimization problem given exchange ratios $\vec{\gamma}^*$.

At a restricted BE all exchange ratios between peers that cannot trade bilaterally are set to zero. We show that a restricted BE exists under Assumption 3.

Let $E = \{(i,j) : T_i \cap S_j \neq \emptyset, T_j \cap S_i \neq \emptyset\}$ be the set of tuples of users with reciprocally desired files. According to Definition 7 $\gamma_{ij} > 0$ if and only if $(i,j) \in E$ at a BE. We consider a price p_{ij} for every tuple $(i,j) \in E$. This is the price that user j pays to download from user i (see Section 4.2). For this proof, let $\vec{p} = (p_{ij}, (i,j) \in E)$. The exchange ratio between i and j is $\gamma_{ij} = p_{ij}/p_{ji}$. In particular, without loss of generality we assume that the budget constraint in the bilateral user optimization of user i is replaced by

$$p_{ji} \sum_f r_{jif} = p_{ij} \sum_f r_{ijf}.$$

We ignore pairs of users that are not in E (since by definition such users cannot trade bilaterally) and show that it is possible to have some $\vec{p} \gg 0$ such that the market clears.

For the purposes of this proof, let $\vec{r}^i(\vec{p})$ be the optimal solution for the bilateral optimization problem of user i when the exchange ratios are equal to $\gamma_{ij} = p_{ij}/p_{ji}$. If $\vec{r}$ and $\vec{p}$ constitute a BE, then $\vec{r} \in \vec{r}^i(\vec{p})$ for all $i \in U$. We note that each r^i_{ijf} is in general a correspondence. We define excess demand for each $(i,j) \in E$ as

$$z_{ij}(\vec{p}) = \sum_f r^j_{ijf}(\vec{p}) - \sum_f r^i_{ijf}(\vec{p}).$$

We first show that the excess demand $\vec{z}$ has the following properties:

(i) For every $\vec{p}$ and $\vec{z} \in \vec{z}(\vec{p})$, $\vec{p} \cdot \vec{z}(\vec{p}) = 0$.

(ii) $\vec{z}(\cdot)$ is convex-valued

(iii) $\vec{z}(\cdot)$ is homogeneous of degree 0

(iv) $\vec{z}(\cdot)$ is upper-hemicontinuous

(v) There is $s > R$ such that $z_{ij} > -s$ for any $\vec{z} \in \vec{z}(\vec{p})$ and $\vec{p}$.

(vi) If $\vec{p}^n \to \vec{p} \neq 0$, $\vec{z}^n \in \vec{z}(\vec{p}^n)$ and $p_{ij} = 0$, $p_{ji} > 0$ for some $(i, j) \in E$, then

$$\max\{z_{ij}^n : (i, j) \in E\} \to \infty.$$

By Assumption 3, the budget constraint of each user binds. The budget constraint of user i is

$$p_{ji} \sum_f r_{jif}^i(\vec{p}) = p_{ij} \sum_f r_{ijf}^i(\vec{p}).$$

By summing over all users, we obtain Property (i).

Fix a price vector $\vec{p} \gg 0$. By Assumption 3 $v(\cdot)$ is strictly concave; therefore $r_{ijf}^i(\vec{p})$ and $r_{ijf}^j(\vec{p})$ are convex-valued. Thus the aggregate excess demand $\vec{z}(\cdot)$ is a convex valued correspondence (Property (ii)).

Consider a price vector $\vec{p} \gg 0$, and fix a constant $t > 0$. It is clear that the feasible region of the bilateral user optimization problem remains unchanged if we replace the price vector $\vec{p}$ by $t\vec{p}$. Thus the aggregate excess demand is homogeneous of degree zero (Property (iii)).

By Assumption 3 $v(\cdot)$ is a continuous function. From the Theorem of the Maximum [10] it follows that $r_{ijf}^i(\vec{p})$ and $r_{jif}^i(\vec{p})$ are upper hemicontinuous correspondences. The aggregate excess demand for $(i, j) \in E$ is a linear combination of the rates $r_{ijf}^j(\vec{p})$ and $r_{ijf}^i(\vec{p})$, and therefore is also upper hemicontinuous (Property (iv)).

The upload rate of any user i is upper bounded by his upload rate constraint B_i, so the total supply is upper bounded and the excess demand is bounded from below (Property (v)).

Suppose that $\vec{p}^n \to \vec{p} \neq 0$, and $p_{ij} = 0$, $p_{ji} > 0$ for some $(i, j) \in E$. Let $f \in T_i \cap S_{i+1}$. As $\vec{p}^m \to \vec{p}$ and the amount of f that user i can afford goes to infinity. On the other hand, the total possible supply is bounded above by the upload rate constraint of user j. Thus $\max\{z_{ij}^m : (i, j) \in E\} \to \infty$, establishing Property (vi).

Using properties (i)-(vi) we show that there exists a BE. Let

$$\Delta = \{\vec{p} \in R_{+}^{|E|} : p_{ij} + p_{ji} = 1, (i,j) \in E\}$$

$$\Delta^n = \{\vec{p} \in \Delta : p_{ij} \geq 1/n, (i,j) \in E\}$$

We observe that Δ^n is compact. Then (from property (iv)) for each n, there exists $\vec{r}^n > 0$ such that $\vec{z}(\vec{p}) \subset [-\vec{r}^n, \vec{r}^n]^{|E|}$. For each n, define $\vec{f}^n : \Delta^n \times [-\vec{r}^n, \vec{r}^n]^{|E|} \to \Delta^n \times [-\vec{r}^n, \vec{r}^n]^{|E|}$ by

$$\vec{f}^n(\vec{p}, \vec{z}) = \{\vec{q} \in \Delta^n : \vec{z} \cdot \vec{q} \geq \vec{z} \cdot \vec{q}', \forall \vec{q}' \in \Delta^n\} \times \vec{z}(\vec{p}).$$

For each n, the correspondence $\vec{f}^n$ is convex-valued and upper-hemicontinuous. We can now apply Kakutani's theorem to conclude that for each n, $\vec{f}^n(\cdot)$ has a fixed point, which we denote by $(\vec{p}^n, \vec{z}^n)$.

The sequence $\vec{p}^n$ in Δ has a subsequence that converges, because Δ is compact. By (v) and the fact that $\vec{z}^n$ is bounded, the limit must be in the interior of Δ. Therefore, by taking a subsequence if necessary, we can assume that $\vec{p}^n \to \vec{p}^*$ and $\vec{z}^n \to \vec{z}^*$, where $\vec{p}^*$ is in the interior of Δ. The limit $\vec{p}^*$ is a BE price vector.

We have shown that a restricted BE exists under Assumption 3. We know show how to construct a BE (according to Definition 7) when Assumption 8 holds. Suppose $\vec{\tilde{r}}$ and $\vec{\tilde{\gamma}}$ constitute a restricted BE. Let $\vec{r}^* = \vec{\tilde{r}}$. For pairs of users i,j such that $S_i \cap T_j \neq \emptyset$ and $S_j \cap T_i \neq \emptyset$, set $\gamma^*_{ij} = \tilde{\gamma}_{ij}$. Having set exchange ratios for all pairs of users that can trade bilaterally, we now consider users that cannot trade bilaterally. If $S_i \cap T_j = \emptyset$ and $S_j \cap T_i = \emptyset$, set $\gamma^*_{ij} = \gamma^*_{ji} = 1$. If $S_i \cap T_j \neq \emptyset$ and $S_j \cap T_i = \emptyset$, set

$$\gamma^*_{ij} = \varepsilon + \max_{k: S_k \cap T_i \neq \emptyset, S_i \cap T_k \neq \emptyset} \{\gamma^*_{kj}\},$$

and $\gamma^*_{ji} = 1/\gamma_{ij}$. If Assumption 8 holds, then $\vec{r}^*$ solves the bilateral optimization problem of every user with respect to exchange ratios $\vec{\gamma}^*$. In particular, user i can find every file in T_i through bilateral trade at the same exchange ratios as in the restricted BE. Exchange ratios with users that i cannot trade with bilaterally are set so that they do not affect i's optimization problem.

$\blacksquare$

Proof of Proposition 16: If Assumption 5 holds, then either the user graph is strongly connected, or the system can be decomposed to subsystems for which the user graphs are strongly connected. Therefore, without loss of generality, in this proof we assume that the user graph is strongly connected.

For the purposes of this proof, let $\vec{r}^i(\vec{p})$ be the optimal solution of the multilateral optimization problem of user i when the price vector is $\vec{p} = (p_i, i \in U)$. If $\vec{r}$ and $\vec{p}$ constitute a CE, then $\vec{r} \in \vec{r}^i(\vec{p})$ for all $i \in U$.

We define excess demand for the upload rate of each user $i \in U$ as

$$z_i(\vec{p}) = \sum_{f,j} r^j_{ijf}(\vec{p}) - \sum_{f,j} r^i_{ijf}(\vec{p}).$$

We show that the aggregate excess demand correspondence $\vec{z}(\cdot)$ defined on $(0,\infty)^{|U|}$ satisfies the following properties:

1. For every $\vec{p} \gg 0$ and $\vec{z} \in \vec{z}(\vec{p})$, $\vec{p} \cdot \vec{z} = 0$.

2. $\vec{z}(\cdot)$ is convex-valued.

3. $\vec{z}(\cdot)$ is homogeneous of degree 0.

4. $\vec{z}(\cdot)$ is upper hemicontinuous.

5. There is an $s > 0$ such that $z_j > -s$ for any $\vec{z} \in \vec{z}(\vec{p})$, for every file $j \in F$ and every price vector $\vec{p} \gg 0$.

6. If $\vec{p}^m \to \vec{p} \neq \vec{0}$, $\vec{z}^m \in \vec{z}(\vec{p}^m)$ and $p_j = 0$ for some j, then $\max\{z^m_j : j \in F\} \to \infty$.

Then the existence of a CE follows from standard results in microeconomics; see, e.g., [50], Exercise 17.C.2.

By Assumption 3, the budget constraint of each user binds. The budget constraint of user i is

$$\sum_{j,f} p_j r^i_{jif}(\vec{p}) = p_i \sum_{j,f} r^i_{ijf}(\vec{p}).$$

By summing over all users, we obtain Property 1.

Fix a price vector $\vec{p} \gg 0$. By Assumption 3 $v(\cdot)$ is strictly concave; therefore $r^i_{ijf}(\vec{p})$ and $r^j_{ijf}(\vec{p})$ are convex-valued. Thus the aggregate excess demand $\vec{z}(\cdot)$ is a convex valued correspondence (Property 2).

Consider a price vector $\vec{p} \gg 0$, and fix a constant $t > 0$. It is clear that the feasible region of the multilateral user optimization problem remains unchanged if we replace the price vector $\vec{p}$ by $t\vec{p}$. Thus the aggregate excess demand is also homogeneous of degree zero (Property 3).

We now show that the aggregate excess demand correspondence is upper hemicontinuous. By Assumption 3 $v(\cdot)$ is a continuous function. From the Theorem of the Maximum [10] it follows that $r^i_{ijf}(\vec{p})$ and $r^i_{jif}(\vec{p})$ are upper hemicontinuous correspondences. The aggregate excess demand for user i is a linear combination of the rates $r^j_{ijf}(\vec{p})$ and $r^i_{ijf}(\vec{p})$, and therefore is also upper hemicontinuous (Property 4).

The upload rate of any user i is upper bounded by his upload rate constraint B_i, so the total supply is upper bounded and the excess demand is bounded from below (Property 5).

If $\vec{p}^m \to \vec{p} \neq 0$ and $p_j = 0$, then $p_k > 0$ for some k. Because of Assumption 5, there is a sequence of users $1, 2, \ldots, n \in U$ such that $T_i \cap S_{i+1} \neq \emptyset$. Thus, there is a user i such that p_i approaches a strictly positive limit and p_{i+1} approaches zero. Let $f \in T_i \cap S_{i+1}$. The budget of user i approaches a strictly positive limit as $\vec{p}^m \to \vec{p}$ and the amount of f he can afford goes to infinity. On the other hand, the total possible supply is bounded above by the upload rate constraints of user $i+1$. Thus $\max\{z^m_j : j \in F\} \to \infty$, establishing Property 6. ■

Proof of Theorem 6: Suppose that $\vec{r} \in \mathcal{X}$ is a Pareto improvement. Then some user i strictly prefers $\vec{r}$ to $\vec{r}^*$. Since $\vec{r}$ is not an optimal solution for user i under $\vec{p}$, it must be that

$$\sum_{j,f} p_j r_{jif} > p_i \sum_{j,f} r_{ijf}.$$

All users $k \neq i$ are at least as well off under $\vec{r}$ as under $\vec{r}^*$. This implies that

$$\sum_{j,f} p_j r_{jkf} \geq p_k \sum_{j,f} r_{kjf},$$

because the utilities are increasing in the total rates of files that users are interested in. In particular, consider a user k who gets exactly the same utility under $\vec{r}$ and $\vec{r}^*$: if $\sum_{j,f} p_k r_{jkf} < p_k \sum_{j,f} r_{kjf}$, then there is a rate allocation that satisfies k's budget constraint and k strictly prefers to $\vec{r}$, which would implies that $\vec{r}^*$ is not optimal.

Summing over all users,

$$\sum_k \sum_{j,f} p_j r_{jkf} > \sum_k p_k \sum_{j,f} r_{kjf},$$

which is a contradiction. We conclude that a CE allocation is Pareto efficient. ■

Proof of Proposition 17: Define $r^*_{ij} \equiv \sum_f r^*_{ijf}$ (the total rate that user i sends to user j). We define the matrix $\vec{Q}$ such that $Q_{ij} = r^*_{ij}$ if $i \neq j$; and $Q_{ii} = -\sum_j r^*_{ij}$. By construction, $\vec{Q}$ is a transition rate matrix with no transient subclasses, since $r^*_{ij} > 0$ implies that $r^*_{ji} > 0$ (by the definition of BE). In what follows we consider the communicating classes of $\vec{Q}$: if $r^*_{ij} > 0$, then users i and j are in the same communicating class. For the purposes of this proof, let $\mathcal{N}_i(\vec{r}^*)$ be the set of peers with which i trades under $\vec{r}^*$, i.e., $\mathcal{N}_i(\vec{r}^*) = \{j \in U : r^*_{ji} > 0\}$.

As noted above, we first observe that $\vec{\pi}\vec{Q} = \vec{0}$ implies that the multilateral budget constraint is satisfied; therefore for any invariant distribution $\vec{\pi}$, $\vec{r}^*$ is feasible for the multilateral optimization problem of every user when prices are equal to $\vec{\pi}$. We show that for some invariant distribution of $\vec{Q}$ (say $\vec{p}$), $\vec{r}^*$ and $\vec{p}$ constitute a CE. In particular, we show that for each user i, $\vec{r}^*$ solves the multilateral optimization problem under $\vec{p}$.

This is done in three steps. First, we show that if $\vec{r}^*$ is Pareto efficient, then $\vec{Q}$ corresponds to a reversible Markov chain. This implies that if $\vec{\pi}$ is a strictly positive invariant distribution of $\vec{Q}$, then $\gamma^*_{ij} = \pi_i/\pi_j$ whenever $r^*_{ij} > 0$, and as a result $\vec{r}^*$ solves the multilateral optimization problem of user i given prices $\vec{\pi}$ if user i is restricted only with users in $\mathcal{N}_i(\vec{r}^*)$ (Step 1). We then show that if user i is restricted to trade with users in the same communicating class under prices $\vec{\pi}$, then $\vec{r}^*$ is an optimal solution of the multilateral optimization problem (Step 2). Step 2 completes the proof if $\vec{Q}$ consists of one communicating class. Finally, we show that if there are multiple communicating classes, there exists an invariant distribution $\vec{p}$ (where the invariant distribution of each communicating class is scaled appropriately) such that $\vec{r}^*$ is an optimal solution of the multilateral optimization problem of each user (Step 3).

We show each of these steps by demonstrating that otherwise there exists a rate vector

$\vec{r}$ that Pareto improves $\vec{r}^*$. Suppose $\vec{r}^*$ solves the bilateral optimization problem of user i under $\vec{\gamma}^*$. Let $(x^*_{if}, f \in T_i)$ and y_i be the corresponding download and upload rates for user i. Consider a rate allocation $\vec{r}$ where $(x_{if}, f \in T_i)$ and y_i are the corresponding download and upload rates for user i. Assuming that $x_{if} - x^*_{if}$ and $y_i - y^*_i$ are sufficiently small, we can use Taylor's approximation to conclude that user i is strictly better off under $\vec{r}$ if

$$(x_{if} - x^*_{if}) \frac{\partial v_i(\vec{x}^*_i, y^*_i)}{\partial x_{if}} > (y_i - y^*_i) \frac{\partial v_i(\vec{x}^*_i, y^*_i)}{\partial y_i}. \tag{4.1}$$

Suppose that $r^*_{jif} > 0$ for some j, which implies that $x^*_{if} > 0$. The optimality conditions for the bilateral optimization problem of user i give that

$$\frac{\partial v_i(x^*_{ig}, g \in T_i)}{\partial x_{if}} = \frac{1}{\gamma^*_{ij}} \frac{\partial v_i(\vec{x}^*_i, y^*_i)}{\partial y_i}.$$

Combining this with (4.1), we see that user strictly prefers $\vec{r}$ to $\vec{r}^*$ if

$$\frac{x_{if} - x^*_{if}}{y_i - y^*_i} > \gamma^*_{ij}. \tag{4.2}$$

Step 1: Let $\vec{\pi}$ be a strictly positive invariant distribution of $\vec{Q}$, i.e., $\vec{\pi} \gg 0$ and $\vec{\pi} \cdot \vec{Q} = \vec{0}$. Then, for every user i, $\sum_k (\pi_k/\pi_i) r^*_{ki} = \sum_k r^*_{ik}$. On the other hand, the budget constraint of the bilateral optimization problem of user i implies that $\gamma^*_{ki} r^*_{ki} = r^*_{ik}$. Summing over k and substituting, we conclude that

$$\sum_k \gamma^*_{ki} r^*_{ki} = \sum_k \frac{p_k}{p_i} r^*_{ki}. \tag{4.3}$$

If $\vec{Q}$ is reversible, then the detailed balance equations hold for every $i, j \in U$, i.e., $\pi_i r^*_{ij} = \pi_j r^*_{ji}$. We note that the detailed balance equations trivially hold if $r^*_{ij} = 0$, because then also $r^*_{ji} = 0$. We show that Pareto efficiency of $\vec{r}^*$ implies reversibility of $\vec{Q}$.

Assume that $\vec{Q}$ is not reversible. Then $\pi_i r^*_{ij} < \pi_j r^*_{ji}$ for some i, j with $r^*_{ij} > 0$. Moreover, since $\vec{\gamma}^*$ and $\vec{r}^*$ constitute a BE, we have $\gamma^*_{ji} = r^*_{ij}/r^*_{ji}$ whenever $r^*_{ji} > 0$. Thus, if $\vec{Q}$ is not reversible, then $\pi_j/\pi_i > \gamma^*_{ji}$ for some i, j with $r^*_{ij} > 0$. Without loss of generality, we relabel i to be $j+1$. Then, by (4.3), there exists some user k such that $\pi_{j+1}/\pi_k > \gamma^*_{j+1,k}$ and $r^*_{j+1,k} > 0$. We relabel k to be user $j+2$, and then $\pi_{j+1}/\pi_{j+2} > \gamma^*_{j+1,j+2}$. Applying this reasoning inductively, we can find a sequence of users $1,2,...,K,K+1$ such that $1 \equiv K+1$ and $\pi_k/\pi_{k+1} > \gamma^*_{k,k+1}$ for all k.

We show how the utility of each user in $S = \{1,2,...,K\}$ can increase while the rate allocation to users outside S remains the same. In particular, we increase $r^*_{k,k-1}$ and y^*_k by a_k for all $k \in S$, as illustrated in the first part of Figure 4.4 (for $K = 3$). We note that users' upload capacity constraints do not bind at the BE, a consequence of the assumption that $v_i(\vec{x}_i, y_i) \to \infty$ as $y_i \to B_i$. Therefore, it is feasible to slightly increase the upload rates of all users. Applying (4.2), user k is better off if

$$\frac{a_{k+1}}{a_k} > \gamma^*_{k,k+1}.$$

Since $\pi_k/\pi_{k+1} > \gamma^*_{k,k+1}$, it follows $\prod_k \gamma^*_{k,k+1} < 1$. Then, it is possible to make all users in the set better off by choosing δ and ε small enough, and setting $a_1 = \delta$; $a_{k+1} = \gamma_{k,k+1} \cdot a_k + \varepsilon$, for all $k \in S$.

We conclude that if $\vec{r}^*$ is the rate allocation of a BE and is Pareto efficient, then $\vec{Q}$ is reversible, and $\gamma^*_{ij} = \pi_i/\pi_j$ whenever $r^*_{ij} > 0$. This means that $\vec{r}^*$ solves the multilateral optimization problem of user i given prices $\vec{\pi}$ if he is restricted to trade with peers in $\mathcal{N}_i(\vec{r}^*)$. The remainder of the proof shows that for some invariant distribution $\vec{p}$, $\vec{r}^*$ is optimal for the multilateral optimization problem of every user under $\vec{p}$.

Step 2: Let $\vec{\pi}$ be a strictly positive invariant distribution of $\vec{Q}$, and consider the multilateral user optimization problems when prices are given by $\vec{\pi}$. We already showed in Step 1 that $\vec{r}^*$ is feasible. Suppose that $\vec{r}^*$ is not optimal for the multilateral optimization problem of some user i. Then by Step 1 there must exist a user j such that $r^*_{j,i} = 0$ with which i wants to exchange under $\vec{\pi}$.

In this step we consider the case that i and j are in the same communicating class. Then, there exists a sequence of users between i and j such that each two consecutive users trade at the BE. Without loss of generality we relabel user i by K, user j by 1, and the users in the sequence by $2,3,...,K-1$. Then, $r^*_{j,j-1} > 0$ for $j = 2,3,...,K$. We show that there is a Pareto improvement, where the utilities of all users in the set $S = \{1,2,...,K\}$ strictly increase, while utilities of users outside S remain the same.

Let a_j be the amount by which we increase rate $r_{j,j-1}$. We assume that all users in

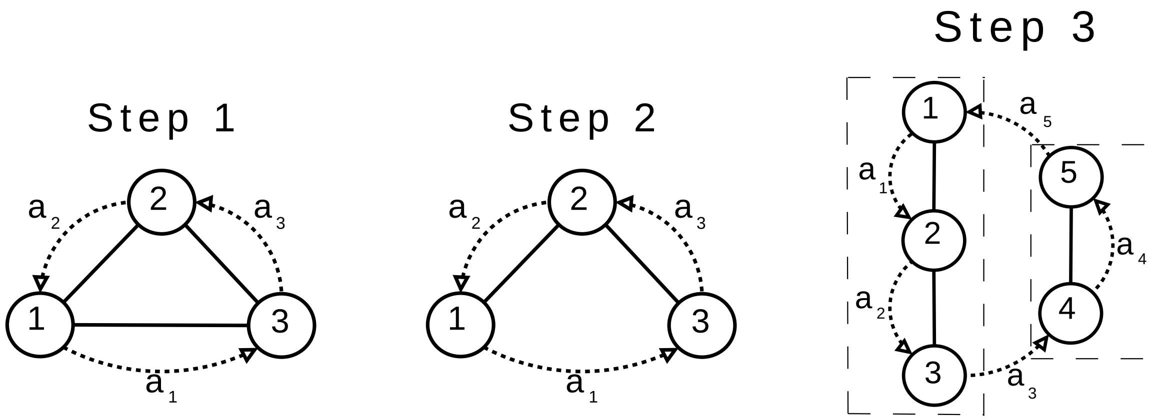

Figure 4.4: Pareto improvements when the BE allocation does not correspond to a CE allocation for Steps 1, 2, and 3 of the proof of Proposition 17 respectively. A pair of users that trade at the BE is connected with a solid line. Dotted arrows show the rates that increase for the Pareto improvement: user i increases his upload rate and the rate he sends to user $i-1$ by a_i. In the third figure (Step 3) there are two communicating classes — each class is included in a dashed box.

the set increase this rate by increasing their upload rates. In particular, user j increases his upload rate by a_j, and gets a_{j+1} more from user $j+1$. This is illustrated for $K=3$ in the second part of Figure 4.4. Applying (4.2), user $j \neq K$ is better off if $a_{j+1}/a_j > \gamma^*_{j,j+1} \equiv \pi_j/\pi_{j+1}$ (the last part follows from the reversibility of $\vec{Q}$). To conclude this step, we show that user K is better off is $a_1/a_K > \pi_K/\pi_1$. Then, as in Step 1, it is possible to find a_i for $i \in S$, such that all users in S are better off.

Now consider user K. Let $f \in T_K$ be a file that user K wants to get from user 1 under prices $\vec{\pi}$. There are two cases to consider, depending on whether user K downloads file f at the BE.

- $r^*_{jKf} > 0$ for some j. Then, by (4.2), we conclude that user K is better off if $a_1/a_K > \gamma_{Kj}$. Moreover, since K prefers to get f from 1 under $\vec{\pi}$ it must be that $\pi_K/\pi_1 > \gamma_{Kj}$. Thus, it suffices that $a_1/a_K > \pi_K/\pi_1$.

- $\sum_j r^*_{jKf} = 0$, i.e., K does not download file f at rate allocation $\vec{r}^*$. Under $\vec{\pi}$, user K is strictly better off downloading a positive amount of f from user 1. This implies that

$$\frac{\partial v_i(x^*_{ig}, g \in T_i)}{\partial x_{if}} > \frac{\pi_1}{\pi_K} \frac{\partial v_i(\vec{x}^*_i, y^*_i)}{\partial y_i}.$$

Combining this with (4.1) we conclude that user K is better off if $a_1/a_K > \pi_K/\pi_1$.

In either case, user K is better off if $a_1/a_K > \pi_K/\pi_1$. This shows that at any optimal solution of the multilateral optimization problem for user i under $\vec{\pi}$, $r_{jif} = 0$ if $r^*_{jif} = 0$ and i, j are in the same communicating class.

Step 3: We now extend the result of Step 2 across communicating classes. Let $\vec{\pi}_c$ be the invariant distribution for communicating class c. We show that there exist coefficients ρ_c such that $\vec{r}^*$ is optimal for the multilateral optimization problem of each user under $\vec{p} \equiv \sum_c \rho_c \vec{\pi}_c$.

We start by deriving the conditions that the coefficients ρ_c need to satisfy. Consider two communicating classes c and c'. If $(\cup_{i \in c} T_i) \cap (\cup_{j \in c'} S_j) \neq \emptyset$, then some users from class c are interested in files that are possessed by users in class c'. To ensure that $\vec{r}^*$ is optimal for the multilateral optimization problems of these users, the ratio $\rho_{c'}/\rho_c$ should be sufficiently large. We denote this lower bound by $\xi_{c',c}$.

Suppose that there do not exist coefficients ρ_c such that $\vec{r}^*$ is an optimal solution of the multilateral optimization problem of each peer. Then, there exists a directed cycle of classes such that (1) $(\cup_{i \in c} T_i) \cap (\cup_{j \in c'} S_j) \neq \emptyset$ for each two consecutive classes in the cycle, and (2) the product of $\xi_{c',c}$ along the cycle is strictly greater than 1. This implies the existence of a vector $\vec{\rho}$ such that $\rho_{c'}/\rho_c < \xi_{c',c}$ for every pair of consecutive classes along the directed cycle. In particular, when prices are $\vec{p} \equiv \sum_c \rho_c \vec{\pi}_c$, for each pair of consecutive classes along the cycle c and c', there is a user n_c in class c that wants to trade with user $m_{c'}$ from class c'. We construct a set S that includes users n_c, m_c as well as the users between them, i.e., users $i_{c1}, \ldots, i_{cl}$ such that $n_c \equiv i_{c1}$, $m_c \equiv i_{cl}$ and $r^*_{i_{cj}, i_{c,j+1}} > 0$. We relabel users in S by $\{1, 2, \ldots, K\}$ such that if i and $i+1$ are in different communicating classes (say c and c') then $i = n_c$ and $i+1 = m_{c'}$, i.e., user i wants to trade with user $i+1$.

We demonstrate a Pareto improvement where user $i \in S$ increases his upload rate and the rate he sends to user $i-1$ by a_i. In Figure 4.4 we illustrate an example with two communicating classes. We demonstrate that it is possible to reallocate rates in a way that strictly increases the utilities of all users in S and does not change the utilities of users outside S. From (4.2) we see that a user $j \neq n_c$ can be made better off if $a_{j+1}/a_j > p_j/p_{j+1}$. A user $j \equiv n_c$ for some i can be made better off if $a_{j+1}/a_j > p_j/p_{j+1}$ (this can be shown by applying the same argument we used for user K in Step 2). As in Steps 1 and 2, since the product of all left hand sides is equal to 1 while the product of all right hand sides is strictly less than 1, it is possible to find a vector $\vec{a}$ that satisfies all these inequalities. ■

Proof of Proposition 18: If $s = 0$, files are chosen uniformly. If there are K files in the system, the probability that a given user has files $\{f_1, ..., f_\sigma\}$ and wants file g is equal to $1/(\binom{K}{\sigma}(K-\sigma))$ for any set of distinct files $\{f_1, ..., f_\sigma, g\}$. A given user i can trade bilaterally with user j with probability

$$\frac{\sigma\binom{K-2}{\sigma-1}}{\binom{K}{\sigma}(K-\sigma)} = \frac{\sigma^2}{K(K-1)}$$

Thus a user cannot trade bilaterally with probability

$$\overline{\rho}_{BE}(K,N,0) = \left(1 - \frac{\sigma^2}{K(K-1)}\right)^{N-1}.$$

We observe that $\overline{\rho}_{BE}(K,N,0)$ is increasing in K. If $K(N) = \sigma\sqrt{N}$, then $\overline{\rho}_{BE}(K(N),N,0) \to 1/e$ as $N \to \infty$. Thus, if $K(N) \geq \sigma\sqrt{N}$, then $\overline{\rho}_{BE}(K(N),N,0) \geq 1/e$ for large N. On the other hand, if there exists $\varepsilon > 0$ such that $K(N) \leq \sigma N^{1/2-\varepsilon}$ for large N, then $\overline{\rho}_{BE}(K(N),N,0) \to 0$ as $N \to \infty$.

We now consider multilateral exchange. We do not have a closed form formula to compute $\overline{\rho}_{ME}(K,N,0)$. Instead, we reduce this to a random graph problem, and we use the results from [20], which studies the size of a strongly connected component of a random digraph. We consider the user graph that was defined in Section 4.3. Recall that this is the directed graph $G = (V,E)$ with $V = U$, and $E = \{(i,j) : S_i \cap T_j \neq \emptyset\}$. If this graph is strongly connected, then all users participate in the multilateral exchange. When users choose file uniformly, there is a directed edge from user i to user j with probability σ/K; this is the probability that user i has the file user j wants. On the other hand, there are N nodes in this graph. Applying the results of [20], the size of the strongly connected component is $\approx K$ if and only if $\equiv (1-\sigma/K)^N \to 0$. This is the case if $N^{1-\varepsilon} \geq \sigma K$ for some small $\varepsilon > 0$. ■

Proof of Proposition 19:

The expected percentage of users that cannot trade bilaterally when there are K files and N users is

$$\overline{\rho}_{BE}(K,N,s) = \sum_{i\neq j:i,j\in\{1,\ldots,K(N)\}} \frac{(ij)^{-s}}{\sum_{i\neq j:i,j\in\{1,\ldots,K(N)\}}(ij)^{-s}} \left(1 - \frac{(ij)^{-s}}{\sum_{i\neq j:i,j\in\{1,\ldots,K(N)\}}(ij)^{-s}}\right)^{N-1}.$$

Let $A_N \equiv \overline{\rho}_{BE}(K(N),N,s)$. We are interested in the limit of A_N as $N \to \infty$.

We observe that

$$1 - \frac{(ij)^{-s}}{\sum_{i\neq j:i,j\in\{1,\ldots,K(N)\}}(ij)^{-s}} < 1,$$

and thus each term in the sum approaches 0 as $N \to \infty$.

We first assume that $K(N) \not\to \infty$ as $N \to \infty$. Then A_N is the sum of a finite number of terms, each of which $\to 0$ as $N \to \infty$. Thus, $A_N \to 0$ as $N \to \infty$.

Now assume that $K(N) \to \infty$ as $N \to \infty$ and let

$$\sigma(s) \equiv \sum_{i\neq j:i,j\in\{1,2,\ldots\}} (ij)^{-s}.$$

Since $s > 1$, $\sigma(s)$ is finite.

$$A_N \leq \frac{1}{\sum_{i\neq j:i,j\in\{1,\ldots,K(N)\}}(ij)^{-s}} \sum_{i\neq j:i,j\in\{1,\ldots,K(N)\}} (ij)^{-s} \left(1 - \frac{(ij)^{-s}}{\sigma(s)}\right)^{N-1}$$

Since

$$\frac{1}{\sum_{i\neq j:i,j\in\{1,\ldots,K(N)\}}(ij)^{-s}} \to \frac{1}{\sigma(s)} < \infty,$$

it suffices to show that

$$B_N \equiv \sum_{i\neq j:i,j\in\{1,\ldots,K(N)\}} (ij)^{-s} \left(1 - \frac{(ij)^{-s}}{\sigma(s)}\right)^{N-1} \to 0 \text{ as } N \to \infty.$$

In particular, it suffices to show that for every $\varepsilon > 0$ there exists $\overline{N}$ such that $B_N \leq \varepsilon$ for all $N < \overline{N}$. We observe that for any N_1,

$$B_N < \sum_{i\neq j:i\cdot j\leq N_1} \left(1 - \frac{(ij)^{-s}}{\sigma(s)}\right)^{N-1} + \sum_{i\neq j:i\cdot j>N_1} (ij)^{-s}.$$

For $\varepsilon > 0$, we choose $N_1(\varepsilon)$ such that

$$\sum_{i \neq j : i \cdot j > N_1} (ij)^{-s} < \varepsilon/2.$$

We choose $\overline{N}$ such that

$$|\{(i,j) : i \cdot j \leq N_1\}| \cdot \left(1 - \frac{N_1^{-s}}{\sigma(s)}\right)^{\overline{N}-1} < \varepsilon/2.$$

Then $B_N < \varepsilon$ for all $N \geq \overline{N}$. ■

Chapter 5

Conclusions

This thesis studies mechanisms that provide incentives for efficient outcomes in large scale online systems. We consider two classes of incentive mechanisms. *Aggregation mechanisms* provide aggregate information on the past behavior of a user to other users in the system. When there is such a mechanism in place, a user should expect that bad behavior now affects his future interactions within the system, and may be incentivized to act in a way that is beneficial for the system. *Market mechanisms* can be used to incentivize contribution to the system by using prices to identify value, and associating a budget with each user; the budget increases when the user contributes to the system and decreases when he uses system resources. By requiring that users have non-negative budgets, users can only use the system in return for high value contributions.

We conclude by briefly discussing the choices made in this thesis with respect to three central modeling issues: the objective, the design space, and dynamics. We mention other reasonable choices and open problems.

5.1 Objective

Throughout this thesis the objective is to design an incentive mechanism that achieves Pareto efficiency. An allocation is Pareto efficient if there does not exist a reallocation of commodities at which some individual is better off and no individual is worse off. This is a natural objective; however, there are other objectives which are reasonable to consider in certain cases.

In the context of an electronic marketplace, efficient trade is closely related to truthful description of the item by the seller. In particular, the outcome is efficient if the item is allocated to the individual that values it most, and since only the seller observes the item before it is sold, it is important that he describes it accurately. This motivates looking for aggregation mechanisms that incentivize sellers to always describe their items truthfully. Since it is not always possible to incentivize the seller to be truthful (e.g., if he does not value future payments), we take the following design approach in Chapter 2: find an aggregation mechanism that maximizes the range of parameters for which it is optimal for the seller to *always* describe his items truthfully to potential buyers. Alternatively, one could relax the requirement that the seller always describes his items truthfully, and instead try to maximize the number of periods or the proportion of time that he is truthful, or minimize the maximum amount by which the seller exaggerates the description of the item.

Even though maximizing truthfulness seems like a natural goal, it may not be optimal for the electronic marketplace in terms of maximizing profits. Sellers pay the marketplace for posting an item for sale (listing fee) and for selling the item (closing fee). Naturally, the marketplace would like to maximize its revenue from these fees. In the long run it seems plausible that a marketplace will want to have a reputation mechanism that achieves efficient trade, because that increases buyers' trust in the market and potentially also the volume of trades. However, in the short run, the marketplace may be better off suppressing some bad information (e.g., by making it less accessible) in order to increase buyers' trust and the number of trades in the near future. In particular, the owner of the marketplace faces the following tradeoff: suppressing bad information increases the volume of trades in the near future, but decreases profits in the long run since buyers may eventually lose their trust in the market. It is an open question to understand how well incentivizing truthfulness is aligned with maximizing profits for the marketplace.

In the peer-to-peer system context, this thesis considers Pareto efficient rate allocations. We model the system as a market associating prices with peers in order to identify value and reward users who are uploading the most valuable content. This is a novel aspect of our approach, since prior work only tries to make users contribute without considering what is the most valuable contribution. We note however, that other objectives related to system performance may also be reasonable, such as maximizing utilization of upload capacities of peers, or minimizing the completion time of some percentile. Even though these objectives do not directly incorporate users' utilities, they are simple and do incorporate the fact that users prefer faster downloads.

5.2 Design Space

A user's behavior may influence his future interaction within the system in two ways. First, certain privileges may apply because of the system design. Second, if the past behavior of users is aggregated into scores, users with high scores may have an advantage when interacting with other users in the system. The system designer can obviously affect the former (by design). Moreover, it can affect the latter by appropriately designing an aggregation mechanism that maps information on the users' past behavior into scores. This thesis studies these two dimensions of the system designer's choices separately; the former in the context of peer-to-peer systems, and the latter in the context of electronic marketplaces. The design space becomes richer when we consider both dimensions simultaneously. This approach would be more relevant for electronic marketplaces, as we discuss below.

In Chapter 3, we show how to view a peer-to-peer system as an exchange economy. In this context the system design specifies that peers need to satisfy a budget constraint, which ensures that peers upload valuable content in return for downloading. This is specified by the system, and is not decided by other peers. It could be the case that if peers were shown some kind of scores of others in the system, they would prioritize peers that contributed more to the system. However, this is not a very strong incentive since the motivation is fairness, but not direct benefit.

According to empirical studies, buyers in electronic marketplaces are willing to pay more in order to interact with sellers with higher scores. The reason is that buyers believe they will directly benefit by this, since a seller with a high score is probably more reliable. In Chapter 2, we study how the marketplace can incentivize sellers to be truthful by appropriately aggregating ratings into scores. The incentivization of sellers can be strengthened by providing privileges to sellers with high scores on the system's behalf. For instance, sellers with high scores may enjoy privileges, such as being able to sell items in certain categories or list multiple items at the same time, or being recommended to buyers. On the other hand, sellers with low scores may be banned from the marketplace.

An open design issue in the area of e-commerce is the optimal level of intervention by the owner of the marketplace. This thesis assumes that the electronic marketplace only provides the platform through which buyers and sellers can trade electronically and uses a reputation mechanism to promote trust. In particular, a key assumption is that the owner of the marketplace designs the aggregation mechanism, but does not otherwise intervene in the interactions of its users. However, the marketplace could try to enhance trust by providing guarantees to potential buyers: if the buyer receives an item that is materially different than what the seller described, then the buyer is covered by the marketplace up to some amount. For instance, this is done by the A-to-z Guarantee program at the Amazon Marketplace and the Paypal Buyer Protection at eBay. Many interesting questions arise in the setting. What guarantees should the marketplace provide? How should the guarantees depend on the seller's score? Is it better if a neutral, third party provides the guarantees instead of the marketplace? How do incentives change when guarantees are provided?

5.3 Dynamics

Chapter 2 studies aggregation mechanisms for electronic marketplaces in a static setting: the payment function, that captures the aggregate behavior of buyers, is assumed to be fixed over time. Moreover, a key assumption in our analysis is that the payment is a function of the seller's score and the aggregation mechanism, but does not depend on the seller's strategy. Our approach is to study the seller's best response with respect to the payment function. We take this non-equilibrium approach, because we believe that the large and dynamic set of participants in the major online markets makes the rationality, knowledge, and coordination required for equilibrium improbable. However, the payment function may change over time (e.g., because buyers partially learn). It would be interesting to study the dynamics of his process in order to understand the implications for the design problem. Do buyers learn the sellers' strategies? Why and how does the payment function change over time? How does the seller learn the new payment function? How quickly does the seller adjust his strategy to changes of the payment? How can the marketplace learn the payment function efficiently?

Dynamics are particularly important when there is a change in the aggregation mechanism. When a change occurs, we expect that there is an adjustment period in which the payment function evolves rapidly. This is something that the designer could take into account when considering whether to change the mechanism. Changing the mechanism too frequently is probably not good for the marketplace, because of cognitive difficulties on the buyers side. Moreover, frequent changes may indicate that the designer is not able to find a good mechanism and as a result buyers may lose their trust in the market. On the other hand, sellers may consider certain mechanisms as more fair and may not be satisfied if an "unfair" mechanism is used. For instance, when eBay added the Detailed Seller Ratings in 2007, many sellers thought that it would hurt them. This is something that the mechanism designer needs to consider, because unsatisfied sellers may stop participating in the marketplace.

In the peer-to-peer setting, this thesis focuses on the analysis and comparison of equilibria, but also considers dynamics. In Section 3.3 we study the tâtonnement price adjustment process, and analyze the rate of convergence. In Section 3.5.2 we discuss peer discovery and price discovery. However, our theoretical analysis is predominantly static: we consider a timescale where peers' preferences do not change, and a fixed population of peers. To better capture reality, we need to consider that a peer's preferences change over time. For instance, when a peer completes a file download, he is no longer interested in that file. Moreover, new files may arrive in the system, and new peers may join. Finally, some peers may leave the system either temporarily of permanently.

The management of the money supply in a peer-to-peer economy becomes an important issue when we consider dynamics. In our equilibrium analysis prices are clearing the market. However, in practice currency would be needed to facilitate trade (as discussed in Section 3.5). As users join and leave, and the network size changes, the system must ensure that an appropriate amount of currency remains. While the adaptation of prices according to excess demand makes the system somewhat robust to moderate inflation and deflation, excessive inflation may cause price convergence to take too long, while excessive deflation can lead to insufficient liquidity. The money supply can be managed through rebate policies or by decreasing the value of all peers' old currency, either proportionally or progressively, and giving an additional direct rebate at a regular rate.

The peer-to-peer system economy of Chapter 3, as well as any digital economy, has the following very interesting property: in a closed system prices eventually become zero. In particular, once a file is downloaded by all peers interested in having it, there is no more demand for the file (while there is supply) and the equilibrium price is zero. In this setting, minimizing the completion time of all peers is equivalent to minimizing the period for which the equilibrium price is strictly positive. An open question is whether simple price update rules perform well with respect to the latter objective.

Bibliography

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

[2] George A. Akerlof. The market for lemons: Quality uncertainty and the market mechanism. *Quart. J. of Econom.*, 84(3):488–500, 1970.

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

[4] Norman H. Anderson. *Foundations of Information Integration Theory*. New York: Academic Press, 1981.

[5] Christina Aperjis, Michael J. Freedman, and Ramesh Johari. Peer-assisted content distribution with prices. In *ACM CoNEXT*, 2008.

[6] Christina Aperjis and Ramesh Johari. A peer-to-peer system as an exchange economy. In *GameNets*, 2006.

[7] Christina Aperjis and Ramesh Johari. Reputation mechanisms and efficient trade. Technical report, Stanford University, 2008.

[8] Christina Aperjis and Ramesh Johari. Optimal windows for aggregating ratings in electronic marketplaces. Technical report, Stanford University, 2009.

[9] Kenneth Arrow, H. D. Block, and Leonid Hurwicz. On the stability of the competitive equilibrium II. *Econometrica*, 27(1):82–109, 1959.

[10] Lawrence M. Ausubel and Raymond J. Deneckere. A generalized theorem of the maximum. Technical Report 899, Northwestern University, Center for Mathematical Studies in Economics and Management Science, 1990.

[11] Dimitri Bertsekas. *Dynamic Programming and Optimal Control*. Athena Scientific, 2005.

[12] Rajat Bhattacharjee and Ashish Goel. Avoiding ballot stuffing in eBay-like reputation systems. In *P2PECON*, pages 133–137, 2005.

[13] Ruchir Bindal and Pei Cao. Can self-organizing P2P file distribution provide QoS guarantees? *OSR, Self-Organizing Systems*, 2006.

[14] Gary Bolton, Ben Greiner, and Axel Ockenfels. Engineering trust - reciprocity in the production of reputation information. Working paper series in economics, University of Cologne, Department of Economics, February 2009.

[15] Chiranjeeb Buragohain, Divyakant Agrawal, and Subhash Suri. A game theoretic framework for incentives in P2P systems. In *Third International Conference on Peer-to-Peer Computing*, pages 48–56, 2003.

[16] Luis Cabral and Ali Hortacsu. Dynamics of seller reputation: Theory and evidence from eBay. *J. of Industr. Econom. (to appear)*, 2009.

[17] Meeyoung Cha, Haewoon Kwak, Pablo Rodriguez, Yong-Yeol Ahn, and Sue Moon. I tube, you tube, everybody tubes: analyzing the world's largest user generated content video system. In *IMC '07: Proceedings of the 7th ACM SIGCOMM conference on Internet measurement*, pages 1–14, New York, NY, USA, 2007. ACM.

[18] Paul Chwelos and Tirtha Dhar. Differences in "truthiness" across online reputation mechanisms. Working Paper, Sauder School of Business. 2008.

[19] Bram Cohen. Incentives build robustness in BitTorrent. In *Workshop on Economics of Peer-to-Peer Systems*, 2003.

[20] Colin Cooper and Alan Frieze. The size of the largest strongly connected component of a random digraph with a given degree sequence. *Combinatorics, Probability and Computing*, 2004.

[21] Martin Cripps, George J. Mailath, and Larry Samuelson. Imperfect monitoring and impermanent reputations. *Econometrica*, 72:407–432, 2004.

[22] Chrysanthos Dellarocas. Reputation mechanism design in online trading environments with pure moral hazard. *Inform. Systems Res.*, 16(2):209–230, 2005.

[23] Chrysanthos Dellarocas and Charles A. Wood. The sound of silence in online feedback: Estimating trading risks in the presence of reporting bias. *Management Sci.*, 54(3):460–476, 2008.

[24] Debojyoti Dutta, Ashish Goel, Ramesh Govindan, and Hui Zhang. The design of a distributed rating scheme for peer-to-peer systems. In *Workshop on Economics of Peer-to-Peer Systems*, 2003.

[25] Mehmet Elmekci. Sustainable reputations with rating systems. Working Paper, Northwestern University. 2008.

[26] Ming Fan, Yong Tan, and Andrew B. Whinston. Evaluation and design of online cooperative feedback mechanisms for reputation management. *IEEE Trans. on Knowl. and Data Eng.*, 17(2):244–254, 2005.

[27] Allan M. Feldman. Bilateral trading processes, pairwise optimality, and pareto optimality. *The Review of Economic Studies*, 40(4):463–473, 1973.

[28] Michal Feldman, Christos Papadimitriou, John Chuang, and Ion Stoica. Free-riding and whitewashing in peer-to-peer systems. In *PINS '04: Proceedings of the ACM SIGCOMM workshop on Practice and theory of incentives in networked systems*, pages 228–236, New York, NY, USA, 2004. ACM Press.

[29] Eric J. Friedman, Joseph Y. Halpern, and Ian Kash. Efficiency and Nash equilibria in a scrip system for P2P networks. In *ACM Conference on Electronic Commerce*, 2006.

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

[31] Anindya Ghose, Panagiotis G. Ipeirotis, and Arun Sundararajan. Reputation premiums in electronic peer-to-peer markets: Analyzing textual feedback and network structure. In *P2PECON*, pages 150–154, 2005.

[32] Philippe Golle, Kevin Leyton-Brown, and Ilya Mironov. Incentives for sharing in peer-to-peer networks. In *ACM Conference on Electronic Commerce*, pages 264–267, 2001.

[33] Ramanathan Guha, Ravi Kumar, Prabhakar Raghavan, and Andrew Tomkins. Propagation of trust and distrust. In *International World Wide Web Conference (WWW2004)*, pages 403–412, 2004.

[34] Minaxi Gupta, Paul Judge, and Mostafa Ammar. A reputation system for peer-to-peer networks. In *NOSSDAV '03: Proceedings of the 13th international workshop on Network and operating systems support for digital audio and video*, pages 144–152, New York, NY, USA, 2003. ACM Press.

[35] Frank Hahn and Takashi Negishi. A theorem on non-tatonnement stability. *Econometrica*, 30(3):463–469, 1962.

[36] Robin M. Hogarth and Hillel J. Einhorn. Order effects in belief updating: The belief-adjustment model. *Cognitive Psychology*, 24(1):1–55, 1992.

[37] Daniel Houser and John Wooders. Reputation in auctions: Theory, and evidence from eBay. *J. Econom. & Management Str.*, 15(2):353–369, 2006.

[38] Daniel Hughes, Geoff Coulson, and James Walkerdine. Free riding on Gnutella revisited: The bell tolls? *IEEE Distributed Systems Online*, 6(6), 2005.

[39] John Ioannidis, Sotiris Ioannidis, Angelos D. Keromytis, and Vassilis Prevelakis. Fileteller: Paying and getting paid for file storage. In *Financial Cryptography*, pages 282–299, 2002.

[40] William Stanley Jevons. *Money and the Mechanism of Exchange*. Kegan Paul, Trench, 1885.

[41] Ramesh Johari and John N. Tsitsiklis. Efficiency loss in a network resource allocation game. *Mathematics of Operations Research*, 29(3):407–435, 2004.

[42] Sham M. Kakade, Michael Kearns, and Luis E. Ortiz. Graphical economics. In *Learning Theory*, volume 3120/2004 of *Lecture Notes in Computer Science*, pages 17–32. Springer Berlin / Heidelberg, 2004.

[43] Kirthi Kalyanam and Shelby McIntyre. Return on reputation in online auction markets. Working Paper, Santa Clara University. June 2001.

[44] Yoshihisa Kashima and Andrew R. Z. Kerekes. A distributed memory model of averaging phenomena in person impression formation. *Journal of Experimental Social Psychology*, 30(5):407 – 455, 1994.

[45] Lingfang Li. Reputation, trust, and rebates: How online auction markets can improve their feedback systems. *Journal of Economics and Management Strategy (forthcoming)*.

[46] Qignmin Liu. Information acquisition and reputation dynamics. SIEPR Discussion Paper 06-30, Stanford University, 2006.

[47] Jeffrey A. Livingston. How valuable is a good reputation? A sample selection model of internet auctions. *The Review of Economics and Statistics*, 87(3):453–465, December 2005.

[48] David Lucking-Reiley, Doug Bryan, Naghi Prasad, and Daniel Reeves. Pennies from eBay: The determinants of price in online auctions. *J. Industrial Econom.*, 55(2):223–233, 2007.

[49] George J. Mailath and Larry Samuelson. *Repeated Games and Reputations*. Oxford University Press, 2006.

[50] Andreu Mas-Colell, Michael D. Whinston, and Jerry R. Green. *Microeconomic Theory*. Oxford University Press, Oxford, United Kingdom, 1995.

[51] Paul Milgrom and Chris Shannon. Monotone comparative statics. *Econometrica*, 62(1), January 1994.

[52] Nolan Miller, Paul Resnick, and Richard Zeckhauser. Eliciting informative feedback: The peer-prediction method. *Management Sci.*, 51(9):1359–1373, 2005.

[53] Tomasz Obloj and Laurence Capron. Whether and when sellers can extract value from reputation advantage: The case of internet auctions. In *Atlanta Competitive Advantage Conference Paper*, 2008.

[54] Joseph M. Ostroy and Ross M. Starr. The transactions role of money. In B. M. Friedman and F. H. Hahn, editors, *Handbook of Monetary Economics*.

[55] Joseph M. Ostroy and Ross M. Starr. Money and the decentralization of exchange. *Econometrica*, 42(6):1093–1113, 1974.

[56] Paul A. Pavlou and Angelika Dimoka. The nature and role of feedback text comments in online marketplaces: Implications for trust building, price premiums, and seller differentiation. *Inform. Systems Res.*, 17(4):392–414, 2006.

[57] Michael Piatek, Tomas Isdal, Thomas Anderson, Arvind Krishnamurthy, and Arun Venkataramani. Do incentives build robustness in BitTorrent? In *NSDI*, 2007.

[58] Michael Piatek, Tomas Isdal, Arvind Krishnamurthy, and Thomas Anderson. One hop reputations for peer-to-peer file sharing workloads. In *NSDI*, 2008.

[59] Paul Resnick, Ko Kuwabara, Richard Zeckhauser, and Eric Friedman. Reputation systems. *Commun. ACM*, 43(12):45–48, 2000.

[60] Paul Resnick and Richard Zeckhauser. *The Economics of the Internet and E-Commerce*, volume 11, chapter Trust among strangers in internet transactions: Empirical analysis of eBay's reputation system, pages 127–157. Elsevier Science Ltd., 2002.

[61] Paul Resnick, Richard Zeckhauser, John Swanson, and Kate Lockwood. The value of reputation on eBay: A controlled experiment. *Experimental Econom.*, 9(2):79–101, June 2006.

[62] J. Rosen. Existence and uniqueness of equilibrium points for concave N-person games. *Econometrica*, 33(3):520–534, 1965.

[63] Siddnartha Sahi and Shuntian Yao. The noncooperative equilibria of a trading economy with complete markets and consistent prices. *Journal of Mathematical Economics*, 18:315–346, 1989.

[64] Carl Shapiro. Premium for high quality products as returns to reputations. *Quart. J. Econom.*, 98:659 – 679, 1983.

[65] Lloyd Shapley and Martin Shubik. Trade using one commodity as a means of payment. *The Journal of Political Economy*, 85(5):937–968, 1977.

[66] Michael Sirivianos, Jong Han Park, Xiaowei Yang, and Stanislaw Jarecki. Dandelion: Cooperative content distribution with robust incentives. In *USENIX Technical*, 2007.

[67] Ross M. Starr. The structure of exchange in barter and monetary economies. *The Quarterly Journal of Economics*, 86(2):290–302, 1972.

[68] Donald M. Topkis. *Supermodularily and complementarity*. Princeton University Press, 1998.

[69] Amos Tversky and Daniel Kahneman. Judgment under uncertainty: Heuristics and biases. *Science*, 185:1124–1131, 1974.

[70] Vivek Vishnumurthy, Sangeeth Chandrakumar, and Emin Gun Sirer. Karma: A secure economic framework for peer-to-peer resource sharing. In *Workshop on Economics of Peer-to-Peer Systems*, 2003.