Noname manuscript No.
(will be inserted by the editor)

A Pricing Scheme for QoS in Overlay Networks Based on First-Price Auctions and Reimbursement

Isabel Amigo · Pablo Belzarena · Sandrine Vaton

Abstract Providing Assured-Quality Services over data networks has been a key objective for the past few decades. Research and commercial activities have been focused on several aspects related to this main objective, such as implementing services over heterogeneous networks, providing scalable solutions and verifying network performance. However, less attention has been devoted to the interaction of these technical aspects with the business plane. Although several quality-based pricing schemes have been proposed, reimbursement proposals, while quite common in other scenarios as health, hotel reservation or airlines, are still rare in the field of Internet Economics. In this work, we propose a simple pricing scheme and study it in detail, in order to use Quality of Service monitoring information as feedback to the business plane, with the ultimate objective of improving the seller's revenue. In our framework, Assured-Quality Services are sold through first-price auctions, and in case of failure, a percentage of the price paid for the service is given back to the buyers. We derive the expression for the willingness to pay and we model the reimbursement problem through a zero-sum Stackelberg game. We show that the Nash equilibrium of such game implies reimbursing 100% in case of failures.

Keywords First-Price Auctions · Network Economics · Pricing · Reimbursement · Quality of Service · Stackelberg Games

I. Amigo, S.Vaton
Télécom Bretagne
Technopole Brest-Iroise - CS 83818
29238 Brest Cedex 3, France
Phone: (+33)(0)2 29 00 11 11
E-mail: {maria.amigo,sandrine.vaton}@telecom-bretagne.eu

P. Belzarena, I. Amigo
Facultad de Ingeniería, Universidad de la República
Julio Herrera y Reissig 565 - C.P. 11300
Montevideo, Uruguay
Phone: (+598) 2711 0974
E-mail: {belza,iamigo}@fing.edu.uy

1 Introduction

An expansion in service types and quality levels is expected in the near future [3]. Tele-presence, tele-medicine, multi-player network games and smart-grid services are a few examples of future enhanced services. In order to provide these services, in addition to intrinsic networking requirements such as scalability, confidentiality and technical aspects, market implications and end users' behaviour must be taken into account. Therefore, holistic and interdisciplinary approaches are needed. Such approaches have enriched the Networking and Internet Economics research fields over the past few decades, as recent research results show. The proposals range from interdomain Quality of Service (QoS) path composition and Service Level Agreement negotiations, such as in [9,32], to higher layers issues, such as modelling user reactions to changes in Internet pricing [30] or issues related to the net-neutrality debate, see for instance [7, 14,17,29]. The enlargement of the service offer also aims to create new market opportunities and target different kinds of user profiles. The real space for this new market has been identified as an issue to be studied, along with how users are expected to react to it. In this regard, quality of experience (QoE) and its influence on willingness to pay has gained importance and has put the end user back in stage [35,38].

In this context, traditional flat rates, where clients pay a single fee for Internet access regardless of usage, have to be revisited, not in order to eliminate them, but rather to identify enhanced services where special pricing (per-service, per-amount-of-bandwidth, per-level-of-quality, etc.) could be needed. Moreover, the mere existence of services with enhanced quality presupposes differentiated pricing, since otherwise every buyer would choose the highest level of quality, which is sustainable from neither a technical nor an economic point of view. In this regard, several pricing schemes for enhanced network services have been proposed (see [15] for a survey), including some based on QoS [33,47].

The justification of new pricing schemes for enhanced services is quite unquestionable from the point of view of the network service providers, as explained in the paragraph above. But would buyers accept differentiated pricing? Our intuition is that they would be ready to pay more for services that are assured without question to be delivered in high quality. The studies carried out throughout this paper show that this intuition, in the analysed context and with our model assumptions, is verified.

On the other hand, in today's networks, failures, while less and less frequent, still do occur, producing a negative effect on buyers' willingness to pay. Moreover, several studies have shown experimentally that user satisfaction has a positive effect on willingness to pay (e.g. [20]). A failure in this context could account for a QoS threshold violation, such as a bound on the delay, a jitter value, or even a service interruption. Intuition also says that, while potential failures have a negative impact on willingness to pay, reimbursement should have a positive one. These statements are intuitively true both in the case of end users and in the case of big customers or brokers, which would need to rely on a good quality service in order to, in turn, deliver properly their services to the final user.

It is in this context that we propose a pricing scheme for assured-quality service selling where buyers can make bids to obtain a quality-assured bit pipe, henceforth referred to as the service, and can be reimbursed if ultimately the service fails. We shall show that this reimbursement scheme, under the main assumption of symmetric buyers with private values, incentivizes buyers not to decrease their willingness to pay due to possible failures, which in the end results in an increase in the expected seller's revenue. Moreover, we show that reimbursing 100% overcomes problems like the *market for lemons* and *moral hazard*, which we show would arise when rational buyers are uncertain about service performance.

In particular, a first-price sealed auction mechanism is proposed to sell the services. Auctions make it possible both to find the market price of services that are not yet widely deployed, since services' market price is revealed as part of the mechanism, and to have guidelines to model the willingness to pay. Most bandwidth auctions in the literature propose the use of second-price auctions, or mechanism seeking to reveal the true valuation that buyers attach to the service. The motivation behind this is to be able to allocate the good in sale to those who value it the most, property usually referred to as incentive compatibility. In our scenario, the objective is rather to maximize seller's revenue, thus we estimate that the complexity that second-price auctions applied to networks imply, as we shall comment on the following section, is not justified.

The use of an auction mechanism rather than a fixed price presents several advantages compared to the later. While using a fixed price would provide the advantage, from the buyers' point of view, of not adding uncertainty about the price of the service, a fixed price would not incentive the adequate use of the scarce resources, which in our context is translated into the fact that QoS is more complex to provide, since there is no control access provided by the pricing mechanism. On the other hand, under an auction mechanism, while uncertainty about the price the buyer is going to pay is added, the buyer never pays a price higher than that one he is willing to pay for the service. In addition, auctions make it possible to provide QoS guarantees, acting as a fundamental building block for access control, and in our particular case they assign the scarce resources to those buyers who value it the most. Moreover, our reimbursement proposal reassures the buyer in the sense that if the quality is not attained, the money paid for the service will be given back to him. Is it based on these facts that we claim that the proposed mechanism achieves an adequate trade-off between prices uncertainties, guarantees to obtain the service with the correct quality and an efficient resource allocation.

In our framework, the seller could be, for example, an alliance of domains who sells a pipe for transit traffic with guaranteed quality. This is the case, for instance, in the solution that is being proposed by the ETICS project [1]. Such multidomain alliances are envisaged to provide a service catalogue where end-to-end pipes with quality guarantees are sold through differentiated pricing.

Buyers could be, for instance, other domains which need to buy transit, or content providers who need bandwidth with quality guarantees in order to properly deliver their enhanced services. Hereafter, we shall simply refer to sellers and buyers. We focus on one single alliance, and the competition that could arise among several alliance is considered out of the scope of the paper.

The proposed pricing scheme assumes the existence of a monitoring infrastructure, which would be the one triggering the reimbursement process. The research and industry community has somehow agreed that for QoS provisioning a monitoring infrastructure is essential. Examples of this are found in recent projects (e.g. [1]), in recent standardisation activities (e.g. [2]), and a wide variety of conducted research, for instance [11,39]. In this sense, our pricing scheme proposes using the existence of the monitoring infrastructure at the business plane, providing an economic justification for monitoring.

The remainder of this paper is organized as follows. In Section 2 we review the related literature and comment on the main contributions of the present work. In Section 3 we clearly state the model, assumptions and definitions required. In Section 4 we study buyers' willingness to pay in the context of first-price auctions, which is given by so-called best bidding strategies. In Section 5, we study the problem from the seller's standpoint in order to derive the best percentage of reimbursement. In particular we present the pricing game modelled through a Stackelberg game. Finally, concluding remarks and future work are presented in Section 6.

2 Related Work and Main Contributions

As aforementioned, though other allocation mechanisms could be used, in this work we focus on an auction mechanism, since we consider it is a flexible mechanism to determine the willingness to pay for services that are quite new in the market. In particular, we consider first-price auctions, since our ultimate objective is to provide an optimal mechanism, that is a mechanism maximizing the seller's revenue, and with low implementation complexity. Although in the present paper we focus on a single type of service, that is all services fail with the same probability, provide the same bandwidth and are delivered through the same single path, the implementation complexity becomes an important issue when extended to a multi-path and multi-service scenario, as is our will in future work. Thus, the work we shall now review focuses on auctions mechanisms and reimbursement policies. We shall not attempt, however, to provide an exhaustive review of allocation mechanisms nor reimbursement policies, which are quite abundant in fields other than the networking one.

Auctions have been used in diverse forms in the networking field, mainly for allocating scarce network resources when QoS is needed. A first proposal appeared in an unpublished paper by Mac Kie-Mason [23], where second-price auctions are used at packet level in order to allocate resources in a multiservice network. More recent work includes [16,24,34], where second-price auctions are used as well. The main reason why such an auction mechanism is used is that it incentivizes buyers to reveal their true valuation. However, second-price auctions are rather hard to implement in the networking scenario in a distributed manner, as shown in [25]. On the contrary, and when seller's revenue maximization is sought, optimal mechanisms are preferred to efficient ones. In this case, first price auctions become a viable, attractive option, as proposed in [12] and used in [8], while revenue equivalence principles between first-price and second-price auctions can be proven (see [21] for more details).

In particular, in this work we model bidders' behaviour when a mechanism based on first-price auctions and reimbursement is in place. For such purpose economic literature provides a vast yet fertile field. Equilibrium or best bidding strategies have been studied under different assumptions. The basic bidding model was introduced in [43], where results for equilibrium bidding strategies with independent private values were shown for valuations drawn from a uniform distribution. More detail was later provided in [37]. Further results relaxing some assumptions were derived in [22,36] and instructive and complete summaries can be found in [28] and [21]. While none of these results consider either failures or reimbursements, in the present work we consider both.

Independently of the allocation mechanism, reimbursement policies proposals, while quite common in the revenue management literature, are rather limited in the networking field. Perhaps the closest work is that proposed by Tuffin *et al.* in [42]. In such work, a simple pricing model for communication networks is presented in which reimbursement occurs if a certain delay threshold is exceeded. Prices are fixed by the seller such that for a given amount of reimbursement, his or her own revenue is maximized. The authors model demand such that it is proportional to the probability that the utility exceeds a given cost. This cost is a function of the price paid, the cost of waiting and the negative cost in case the delay threshold is exceeded, which corresponds to a reimbursement. A certain shape for the utility's probability distribution function is assumed in order to draw conclusions and perform numerical experiments. The authors show through such studies that this mechanism increases the seller's revenue compared to the case with no reimbursement. The idea behind this method is the same as ours, though buyers' side is modelled in a very different way, since in that work the price is fixed by the seller, while in our work the price is determined through the auction mechanism.

Yet another approach is proposed in [13], where second-price auctions are used for buying one unit of a computing resource. Winning buyers pay and with a certain probability will indeed need the service. If winners do not use the service, they receive a refund of a percentage of the payment. The authors propose numerical studies to determine the best refund strategy, and define a correlation between the valuation of the object and the probability of not using it. They conclude that different correlations result in different percentages of optimal reimbursement. The scenario is quite different to ours, but the logic behind the mechanism is very similar to our proposal. However, in the present work we focus on analytical results, both for computing the dependency of the willingness to pay on the probability of failure and reimbursement, and to compute the seller's revenue and optimal percentage of reimbursement. In addition, we focus on first-price auctions.

Other reimbursement schemes have also been studied in another, non-networking context. A particular case refers to services that have the peculiarity that they can be returned afterwards and will still possess some value for the seller (see for instance [26]). This is not the case in the networking scenario, in which the service no longer has any value to the seller once it has been used.

Once we model willingness to pay, we shall focus on the seller's problem. We shall assess the seller's expected revenue and find the optimal percentage of reimbursement, that is, the percentage of reimbursement that maximizes the seller's expected revenue. We shall split the analysis into two different cases, namely when buyers perform some level of service monitoring or have some knowledge about service performance, which we shall refer to as the complete information case, and when they do not. The latter is an asymmetric information situation, where the seller has more knowledge about the service than buyers have. We model the reimbursement problem in this asymmetric case through a Stackelberg game [44], where the seller initially announces a percentage of reimbursement and buyers bid according to it. In a Stackelberg game, there is a *leader* that plays first and the *followers* that play afterwards, knowing the leader's move, to reach a Nash Equilibrium. Stackelberg games are very suitable for modelling pricing situations, where the network or the seller typically acts as *leader* and the users or buyers act as *followers.* This kind of game has been widely used in the literature to design revenue-maximizing network policies. For example, [10, 40] where an Internet packet-pricing scheme is proposed for monopolistic service providers and large numbers of users, or [6], where a pricing scheme for differentiated services is proposed, or yet [41], where a user loyalty model to Internet service providers is proposed and applied in a game-theoretical framework in order to derive optimal Internet access pricing strategies. In addition, this kind of hierarchical game has also been studied for pricing along with power control in wireless networks, see for instance [5, 18, 46] and for spectrum sharing in such networks [45].

In this work we provide insight into a pricing mechanism based on first-price bandwidth auctions with reimbursement, which to the best of our knowledge, has not been proposed before. The main contribution is twofold. We determine the best bidding strategy under a first-price auction for services that are prone to failures and with money-back guarantees; and we compute the optimal reimbursement value when bidders are aware of the reimbursement policy.

Regarding the best bidding strategy, we derive its mathematical expression for symmetric bidders with private values and show that it follows some intuition. Indeed, for values of reimbursement lower than 100%, the higher the probability of failure, the lower the bid, while the higher the percentage of reimbursement, the higher the bid.

With respect to the optimal percentage of reimbursement, we show that when there are no information asymmetries, that is, when both seller and buyers are sure about the performance of the service, the expected seller's revenue is not sensitive to the percentage of reimbursement. On the other hand, when there is asymmetric information, that is when buyers are uncertain about the performance of the service they plan to buy, and assuming they act rationally, setting the percentage of reimbursement equal to 100% maximizes the seller's revenue. In addition, in the latter case, we show that if percentage of reimbursement is to be set to a value smaller than 100%, the *market for lemons* phenomenon appears, where the bids decrease until market disappearance. Conversely, if the percentage of reimbursement is set to a value greater than 100%, the so-called *moral hazard* behaviour is observed, where buyers take the risk of assuming a very good performance, since in case of failure the seller would bear the costs. In both cases the seller's expected revenue diminishes. Setting the reimbursement equal to 100% overcomes these problems, and the resulting outcomes for seller and buyers are the same as when there is complete information.

3 The model

Let us begin by describing our working scenario and introducing the notations, definitions and assumptions needed to model it. We are studying a situation where quality-assured services are sold over an interdomain network. Such services could be, for instance, video on demand, a VPN service interconnecting two remote sites or a network game. In all cases, the service can be abstracted to a certain amount of bandwidth guaranteed between two sites through an overlay network, and with certain quality parameters associated with it. We shall call this abstraction an *object*. The quality parameters associated with the object could be given, for instance, by values of the delay, the jitter, the percentage of packet failures, the percentage of service availability, etc.

Objects are sold via a first-price sealed auction mechanism. The following common assumptions are made regarding this mechanism. We first assume a single-object case, that is, M bidders or buyers compete to buy one object. We then move to the case of multi-object, single-unit demand. In other words, M bidders compete to buy K identical objects, and each bidder is interested in buying one single unit of such objects. Each bidder i assigns a valuation X_i to the object and we assume that the X_is are independently and identically distributed according to a common distribution function F, which is known by all buyers. This is the so-called symmetric model, since all bidders' valuations are distributed according to the same distribution function. At the moment of bidding, bidder i knows the realization x_i of his or her valuation but does not know the valuation attached to the object by other bidders, and this knowledge would not affect his or her own valuation, which is the so-called private values model. Valuations in the private values model are as well usually referred to as types.

Conversely, we assume that the service has no value to the bidder if it fails. Please note that actually bidders could attach a negative value to the service when it fails, rather than a null one. This would be the case, for instance, if the failure causes losses to the buyer's business. This could be easily modelled by considering a negative deterministic valuation in case of failure. For clarity's sake, we shall not consider this artefact in the model, though doing so would not change the methodology applied to address the problem.

Bidders are assumed to be risk-neutral, as they seek to maximize their utilities, which we shall model through their expected payoff introduced in the following section. Bidder i's bid is denoted by b_i and it is obtained according to a bidding strategy called β_i. That is to say, bidder i's bid is determined as $b_i = \beta_i(x_i)$. Finally, we assume a discriminatory payment rule, which means that the winning bidder pays his bid. We shall generally simplify notation and refer to x as the realization of the valuation of any given bidder.

The service has an associated probability of failure, denoted by θ in our framework. If indeed the service fails, money is given back. The amount of money returned is proportional to what has been paid for the object and the coefficient of proportionality is represented by $q \geq 0$, which could *a priori* be greater than one. We shall as well refer to q as a percentage of reimbursement.

The percentage of reimbursement associated with the object is always announced to the bidders before they announce their bids, and it is the same value for all bidders. Bidders have their own estimations regarding the probability of failure of a service, which will have an impact on the value of the bid they submit. This estimation could, *a priori*, be based on service performance perception. Buyers could even perform their own measurements on historical observations in order to estimate the probability of failure, or they could infer it from the percentage of reimbursement announced. But we shall address this issue later on. For the moment let us denote the probability of failure assumed by the bidders at the moment of placing their bid as $\hat{\theta}$, which is not necessarily equal to θ.

For simplicity in the presentation we shall focus on a one-shot auction, in which, as usually, services are allocated to the K highest bids. However, when considering multi-period allocations some considerations are worth noting. Since we are selling QoS services we shall not allow ongoing services, allocated in a previous period, to be interrupted in order to accept new higher bids in a subsequent period. Different proposals can be found in the literature that address the multi-period case and that could be adopted in our scenario. Examples of this are, for instance, allowing ongoing services to resubmit a bid [34], allocating services during their whole duration [12] and considering future bids assuming some statistical characterizations, or even allocating at each period the available capacity. Any of these approaches could be considered in the multi-period case.

4 The Optimal Bidding Strategy

In order to determine how the willingness to pay for a service is affected by the probability of failure $\hat{\theta}$ and the percentage of reimbursement q, we study the optimal bidding strategy under such conditions, under the assumptions of Section 3, first for a single object on sale and then for multiple objects.

First-price sealed auctions can be modelled as Bayesian games, which are strategic games with imperfect information (see for instance [31]). Indeed, buyers compete among them for getting the service and they all posses some private information (their types or valuations) which is not known for the rest of the buyers. This means that buyers do not know the realization of other buyers' valuations. A pure strategy for them is a function that maps a valuation into a bid. Thus, when we refer to an optimal bidding strategy, we are referring to the Bayesian Nash-equilibrium of the game defined by the auction mechanism. Throughout this section, references to the term equilibrium should be interpreted in that sense. In the following subsections this is formalized and the equilibrium strategies are derived.

4.1 The Single-Object Case

The single-object case models the situation in which the total available capacity, along with quality guarantees, is to be allocated to one single client, i.e. to the winning bidder. We shall show that in this case and under the assumptions of Section 3, a symmetric equilibrium exists, that is, an equilibrium where all bidders adopt the same best strategy. Theorem (1) formally states this along with the mathematical expression for the best bidding strategy. An outline of the proof is provided below, while the detailed proof can be found in Appendix A.

Let $Y_{M-1}^{(1)}$ be a random variable defined as the maximum over $M-1$ i.i.d. random values from distribution F and $D = \{\hat{\theta} \in [0,1), q \geq 0 : q\hat{\theta} < 1\}$.

Theorem 1 *The Symmetric Equilibrium, Single Object Case. Given a set of M symmetric bidders whose valuations X_i, $i = 1 \ldots M$ are identically and independently distributed (i.i.d.) from a probability distribution $F(x)$, the bidding strategy that maximizes each bidder's payoff in a first-price sealed auction mechanism for a single object which is assumed to fail with probability $\hat{\theta}$ and for which a percentage q of the amount paid is given back if it actually fails, is the same for all bidders and is given by:*

$$\beta(x) = \frac{1-\hat{\theta}}{1-q\hat{\theta}} E[Y_{M-1}^{(1)} | Y_{M-1}^{(1)} \leq x], \ (\hat{\theta}, q) \in D. \tag{1}$$

Proof Let us first assume that a symmetric equilibrium exists, meaning that all bidders follow the same strategy, $\beta_i = \beta \ \forall i \in M$. Any bidder's payoff p can thus be expressed as a function of bid b as in Equation (2), where $\beta(x) = b$ and $\mathbb{1}_e$ is equal to 1 if event e occurs, and 0 otherwise.

$$p = \mathbb{1}_{win}(x\mathbb{1}_{not\,failure} - b(1 - q\mathbb{1}_{failure})), \tag{2}$$

Now let G be the cumulative distribution function of the maximum valuation over $M-1$ valuations i.i.d. according to F, which we have denoted as

$Y_{M-1}^{(1)}$. The notation in $Y_{M-1}^{(1)}$ means that we are selecting the highest value, indicated by superscript 1, among a sample of size $M-1$, indicated by subscript $(M-1)$. Please note that $Y_{M-1}^{(1)}$ is then the $(M-1)-th$ order statistics of a sample of $(M-1)$ i.i.d. values according to F.

The expectation of a bidder i's payoff can be expressed as:

$$\tilde{P} = E\{p|X_i = x_i\} = G(\beta^{-1}(b_i))(x_i(1-\hat{\theta}) - b_i(1-q\hat{\theta})). \tag{3}$$

In Equation (3) we have used the fact that the probability of winning the auction is

$$\begin{aligned}\mathbb{P}(b_i > \max_{j\neq i} b_j) &= \mathbb{P}(\beta(x_i) > \max_{j\neq i} \beta(X_j) \\ &= \mathbb{P}(x_i > \max_{j\neq i} X_j) = G(x_i),\end{aligned} \tag{4}$$

where symmetric equilibrium is assumed, and in the last equality, the assumption is made that β is a strictly increasing function of x. Please note that in the reasoning above we have used $\hat{\theta}$ and not θ, since we are looking at the problem from the buyer's point of view. Since the previous reasoning is valid for any bidder, in what follows subscript notation is avoided.

Two different cases need to be addressed. When $q\hat{\theta} \geq 1$ the expected payoff does not present a maximum, but rather it continues to grow when the bid increases. It is easy to verify that such a situation is not an equilibrium. Consequently, we are not interested in this case. In the following we focus on the second case, i.e. $q\hat{\theta} < 1$.

Finding the bidding strategy β that maximizes Equation (3) reduces to setting its derivative with respect to b equal to zero and imposing $b = \beta(x)$. The derivative of the expected payoff with respect to b is shown in Equation (5), where we have introduced the notation $g(x) = G'(x)$.

$$xg(x)(1-\hat{\theta}) - [g(x)\beta(x) + G(x)\beta'(x)](1-q\hat{\theta}) = 0 \tag{5}$$

Equation (5) is a first order differential equation, whose solution, derived in Appendix A, is:

$$\beta(x) = \frac{1-\hat{\theta}}{1-q\hat{\theta}} \frac{1}{G(x)} \int_0^x vg(v)dv, \, \forall\, (\hat{\theta}, q) \in D. \tag{6}$$

Directly applying the definition of conditional expectation Equation (6), for $(\hat{\theta}, q) \in D$ can be written as

$$\beta(x) = E[Y_{M-1}^{(1)} | Y_{M-1}^{(1)} \leq x] \frac{1-\hat{\theta}}{1-q\hat{\theta}}. \tag{7}$$

Equation (7) possesses a closed form expression for several distributions, as we shall see in an example of later on.

Finally, it remains to be shown that the assumption of symmetric equilibrium holds. We again refer to Appendix A for the complete proof. Without loss of generality, consider that all bidders but bidder 1 use the same strategy given by Equation (6). It must be checked whether in that case it is also optimal for bidder 1 to follow such a strategy.

Let bidder 1's expected payoff $\tilde{P}$ be expressed as a function of his valuation and his bid. If bidder 1 bids $\beta(z)$, when the value he attaches to the service is actually x, his payoff is:

$$\tilde{P}(\beta(z), x) = G(z)[x(1-\hat{\theta}) - \beta(z)(1-q\hat{\theta})]. \tag{8}$$

Hence, the difference with his expected payoff if he were to bid $\beta(x)$ is:

$$\begin{aligned}
&\tilde{P}(\beta(z), x) - \tilde{P}(\beta(x), x) = \\
&\quad G(z)[x(1-\hat{\theta}) - \beta(z)(1-q\hat{\theta})] \\
&\quad - G(x)[x(1-\hat{\theta}) - \beta(x)(1-q\hat{\theta})] \\
&\quad = (1-\hat{\theta})\left[G(z)(x-z) + \int_x^z G(v)dv\right],
\end{aligned} \tag{9}$$

where the last equality comes from applying integration by parts, as detailed in Appendix A.

It remains to be determined if Equation (9) is negative for any value of z. If it were to be negative, then bidder 1's expected payoff by bidding something different from $\beta(x)$ would not be greater than the payoff obtained if he were to bid $\beta(x)$.

We graphically show that indeed, Equation (9) is strictly negative $\forall z \neq x$. Consider Fig. 1, where two different cases are distinguished. When $z < x$, as illustrated in Fig. 1a, the shaded area corresponds to $-\int_x^z G(v)dv$ and the dotted area corresponds to $G(z)(x-z)$. The shaded area is greater than the dotted area for any G (we recall that G is a distribution function, thus it is always non-decreasing). Hence, Equation (9) is negative for any $z < x$.

Likewise, when $z > x$, which corresponds to the situation in Fig. 1b, the dotted area corresponds to the term $\int_x^z G(v)dv$ and the shaded area to the opposite of $G(z)(x-z)$. The shaded area is always greater than the dotted one, thus Equation (9) is again strictly negative. All together, we obtain that Equation (9) is negative for any value of $z \neq x$, which means that there is no other bid leading to a greater payoff than bidding $\beta(x)$. We conclude that the assumption of symmetric equilibrium holds for the case $(\hat{\theta}, q) \in D$.

Finally, it is easy to check from Equation (7) that β is an increasing function of x. □

4.2 Multi-Object, Single-Demand Case

The multi-object scenario corresponds to the case where the available capacity is split into several identical bandwidth chunks, each with certain quality guarantees, and each to be assigned to a different client. Single-demand means that each client is interested in exactly one of these chunks or services. Let us

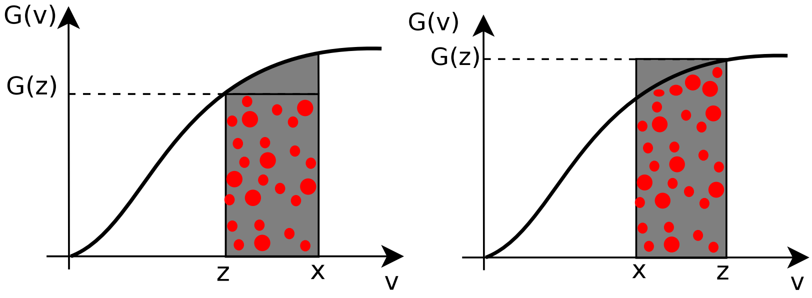

(a) $z < x$ (b) $z > x$

Fig. 1: Illustration of the components of Equation (9).

assume that the auction mechanism is launched for selling K objects. The only thing that changes with respect to the single-object case is the probability of winning the auction, which in this case corresponds to the probability of being among the K highest bids.

Let $Y_{M-1}^{(K)}$ be a random variable defined as the K-th highest value among $(M-1)$ i.i.d. values according to F, and where $D = \{\hat{\theta} \in [0, 1), q \geq 0 : q\hat{\theta} < 1\}$.

Theorem 2 *The Symmetric Equilibrium, Multiple-Object Single-Demand Case. Consider a multi-unit single-demand first-price sealed auction mechanism to sell K objects among a set of M symmetric bidders whose valuations X_i, $i = 1 \ldots M$ are i.i.d. from a probability distribution $F(x)$.*

Consider buyers assume services fail with a probability $\hat{\theta}$. If the service ultimately fails, winning buyers receive back a percentage q of the money paid for the service.

In such conditions, the bidding strategy that maximizes each bidder's payoff is the same for all the bidders and is given by:

$$\beta(x) = E[Y_{M-1}^{(K)} | Y_{M-1}^{(K)} \leq x] \frac{1 - \hat{\theta}}{1 - q\hat{\theta}}, \hat{\theta}, q \in D. \tag{10}$$

Proof Since β is assumed to be an increasing function of x_i, and we assume a symmetric equilibrium, i.e. $b_i = \beta(x_i) \forall i$, the probability of winning for bidder i is equal to the probability that his valuation x_i is among the K highest valuations.

Considering this probability of winning the auction the proof is analogous to the single object case. Hence, the symmetric equilibrium for the optimal bidding strategy is readily generalized to the multi-object scenario and is given by Equation (10). □

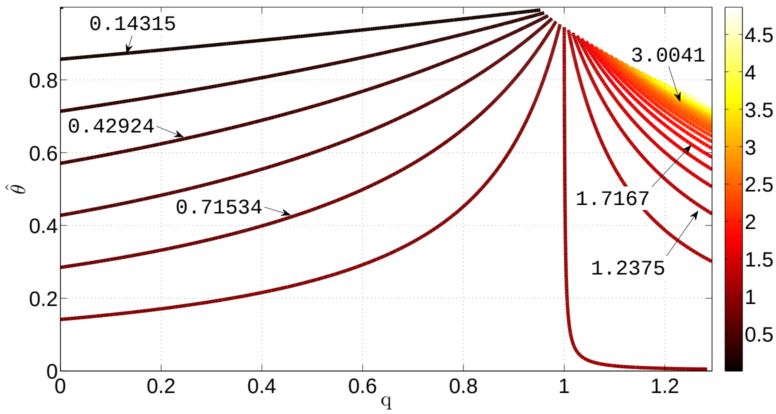

Fig. 2: Contour Lines of function α, the multiplying factor of the best bidding strategy, for different assumed probabilities of failure $\hat{\theta}$ and percentage of reimbursement q.

4.3 Bidding Behaviour Remarks

The best bidding strategy deserves a closer look. First, it is interesting to note that it follows the intuitions stated above in Section 1. Indeed, according to Theorem (1) and Theorem (2), β increases when the percentage of reimbursement does so, and when the percentage of reimbursement is less than 100%, β decreases when the probability of failure increases. This means that buyers decrease their bids when they assume that services fail frequently unless the percentage of reimbursement is greater than or equal to 100%, which is quite intuitive. It means as well that for the same level of failures, the higher the percentage of reimbursement, the higher the bid, which is also consistent with intuition.

This behaviour is shown in Fig. 2 for different values of $(\hat{\theta}, q) \in D$, where α, defined as $\alpha = \frac{1-\hat{\theta}}{1-q\hat{\theta}}$, is the multiplying factor on the best bidding strategy.

Let us finalize this section with an illustrative example. Consider a situation where valuations are uniformly distributed on $[0, 1]$ and there is a single service for sale. In this case $F(x) = x$, $G(x) = x^{M-1}$ and the optimal bidding strategy is given by Equation (11), which is valid for $(\hat{\theta}, q) \in D$.

$$\beta(x) = \frac{1-\hat{\theta}}{1-q\hat{\theta}} \frac{M-1}{M} x. \tag{11}$$

It is worth noting that for different values of q and $\hat{\theta}$, bidders either shade their valuations, announce their true valuation, or overbid. We say that they shade their valuations when the bid is smaller than the valuation, as occurs in the case with no reimbursements and no failures. Conversely, we say that bidders overbid when their bids are greater than their valuations. In contrast, in cases with no failures, bidders always shade their bids.

5 Expected Seller's Revenue

We shall now study the problem from the seller's standpoint, with the ultimate objective of finding the optimal value of the percentage of reimbursement q. For these purposes, we shall study the seller's outcome. Since there are some uncertainties at each transaction, namely whether the service will fail or not, as well as the intrinsic uncertainty of the selling price due to the auction mechanism in place, we shall model the seller's outcome through his or her expected revenue. We recall that there are K units of the same object for sale. There are $M \geq K$ bidders, who participate in a first-price auction to obtain one object. Let us order their bids as:

$$b^{(1)} \geq b^{(2)} \geq \cdots \geq b^{(M)}, \tag{12}$$

which are obtained as $b^{(i)} = \beta(x^{(i)})$, where $x^{(i)}$ represents the ordered bidders' valuation and β is given by Theorem (2).

The services are allocated to the K highest bids; thus the seller's revenue can be expressed as a function of K and the valuations $\mathbf{x} = (x^{(1)}, \ldots, x^{(M)})$ as:

$$\begin{aligned} I(K, \mathbf{x}) &= \sum_{i=1}^{K} b^{(i)} = \sum_{i=1}^{K} \beta(x^{(i)}) \\ &= \sum_{i=1}^{K} E[Y_{M-1}^{(K)} | Y_{M-1}^{(K)} \leq x^{(i)}] \frac{1-\hat{\theta}}{1-q\hat{\theta}}. \end{aligned} \tag{13}$$

In order to simplify the notations henceforth let us introduce function $u(K, \mathbf{x})$ defined as:

$$u(K, \mathbf{x}) = \sum_{i=1}^{K} E[Y_{M-1}^{(K)} | Y_{M-1}^{(K)} \leq x^{(i)}]. \tag{14}$$

According to our proposed pricing scheme, if failures take place the seller will give money back. We shall thus compute the net seller's revenue, that is, the revenue earned from selling the services minus the money that is given back. For brevity, we shall from now on simply refer to it as revenue. Let us assume that failures occur for all services at the same time. This model accounts for failures given by an equipment failure or congestion for example. Then the seller's revenue given that the bidders' valuations are $\mathbf{x} = (x^{(1)}, \ldots, x^{(M)})$, is:

$$\begin{aligned} &R(K, \mathbf{x}) \\ =&I(K, \mathbf{x}) \mathbb{1}_{no\,failure} + (1-q) I(K, \mathbf{x}) \mathbb{1}_{failure} \end{aligned} \tag{15}$$

Since valuations and failure events are independent, the mean seller's revenue given the bidders' valuations $\mathbf{x}$ is:

$$\begin{aligned}&\bar{R}(K,\mathbf{x})\\ =&I(K,\mathbf{x})(1-q\theta)=u(K,\mathbf{x})\frac{1-\hat{\theta}}{1-q\hat{\theta}}(1-q\theta).\end{aligned} \tag{16}$$

Please note that if we were to assume that failures do not occur for all objects at the same time, it is easy to check that we would also obtain Equation (16) for the expected seller's revenue.

Finally, the *ex ante* expected seller's revenue, that is the seller's expected revenue taking into account the randomness of valuation's vector $\mathbf{X}$, is obtained as

$$E\{\bar{R}(K,\mathbf{X})\}=E\{u(K,\mathbf{X})\}\frac{1-\hat{\theta}}{1-q\hat{\theta}}(1-q\theta), \tag{17}$$

where the expectation is over the valuations X_i, and $\mathbf{X}=(X^{(1)}\ldots X^{(M)})$ is the vector of valuations X_i, $i\in\{1,\ldots,M\}$ sorted in non-increasing order.

Hence, the seller's expected revenue can be tuned through the value of q. However, the value of $\hat{\theta}$, which also influences the seller's expected revenue, is determined by the buyers. We shall divide the following study into three parts, each of which makes different assumptions about the buyers' behaviour.

5.1 Complete Information

We shall first assume that buyers have perfect information about the services' performance. The information can be obtained, for instance, through a monitoring infrastructure available for buyers' consultation, or from knowledge obtained through previous observations. In our model, this situation is translated into $\hat{\theta}=\theta$. We shall refer to this scenario as complete information because it supposes that seller and buyers have the same information regarding the probability of failures. Note that it is an assumption of this scenario that the monitoring infrastructure has run long enough to provide accurate information. We shall relax this assumption in the following subsections.

Hence, the expected seller's revenue becomes, regardless the value of q, equal to:

$$E\{\bar{R}(K,\mathbf{X})\}=E\{u(K,\mathbf{X})\}(1-\theta). \tag{18}$$

According to Equation (18), the seller has incentives to keep the probability of failure low, which is quite intuitive.

5.2 Asymmetric Information with *Naive* Buyers

We shall now consider the situation where buyers have no means of determining the probability of failure on their own. We refer to this situation as asymmetric with respect to the information, since the seller has more knowledge about service performance than buyers do. In this situation, the seller can announce the probability of failure, along with the percentage of reimbursement. We shall assume that buyers will take the value of the probability of failure announced by the seller as granted. We have called buyers in this situation as *naive*, risking to be using a too strong characterization. This should be rather interpreted as opposed to the rational buyers case, which we shall introduce following in Subsection 5.3, and the term naive should only be interpreted as buyers being seller's announcements takers.

It can be readily derived from the seller's expected revenue in Equation (17), that this expected revenue increases with q when $\hat{\theta}$ is greater than θ.

Indeed, let us define $B_\theta : D \to \mathbb{R}^+$ as:

$$B_\theta(q,\hat{\theta}) = \frac{1-\hat{\theta}}{1-q\hat{\theta}}(1-q\theta), \tag{19}$$

where $D = \{\hat{\theta} \in [0,1), q \geq 0 : q\hat{\theta} < 1\}$.

According to Equation (17), the behaviour of the seller's revenue is driven by Equation (19), which is shown in Fig. 3, where as an illustrative example we have plotted B_θ when $\theta = 10\%$ and for different values of q. The seller could take advantage of this behaviour by announcing a probability of failure higher than the real one, and setting reimbursement at a value greater than 100%. In other words, negative marketing could be used, with a higher probability of failure being announced than the real one, with the goal of fooling naive buyers for the seller's benefit.

However, the negative marketing policy could be disadvantageous for the seller as well, for at least two reasons. First, if in the end buyers disregard sellers announcement with respect to the probability of service and assume some other value convenient for them, seller's revenue could diminish, to a value even lower than the revenue obtained when reimbursing 100%. Second, the seller could be judged by buyers as untrustworthy, which could lead to losses not captured in our model. More formally, in the former, the buyers would thus act as rational, which is the case addressed in the following section. If in addition, reputation effects should be studied, in order to draw conclusions about the impact on the seller's revenue when announcing a misleading value of the probability of failure, a repeated game should be studied and a number of hypothesis should be established. For instance, hypothesis should be made with respect to whether buyers and seller have a finite-horizon or a long-term horizon for maximizing their utilities (see e.g. [19]) or whether the outcomes of the game are observable for all buyers or not. However, we shall not deepen in this but rather show in the following subsection that there is a seller's strategy, namely $q = 1$, that renders his expected revenue independent of the buyer's assumptions.

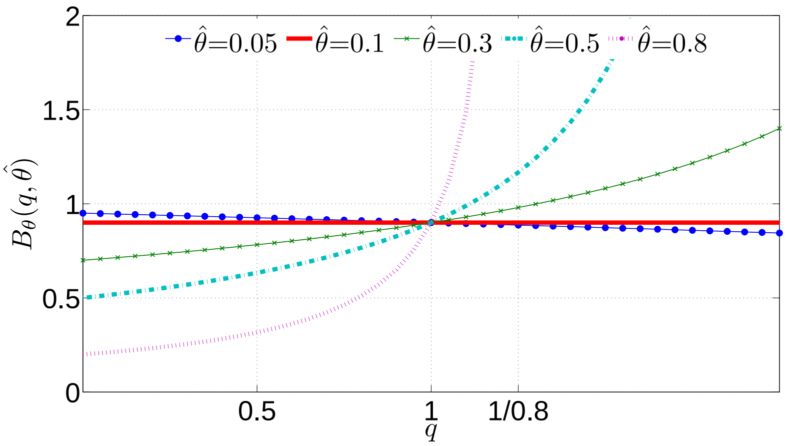

Fig. 3: Variation of the seller's expected revenue as a function of reimbursement q for a real probability of failure $\theta = 0.1$ and for different values of probability of failure assumed by the buyers.

5.3 Asymmetric Information with Rational Buyers

Let us now consider the case where buyers are uncertain about the probability of failure of the service they wish to buy and where they act rationally, seeking to maximize their payoffs. The seller's ultimate objective is still to set the value of q such that his or her revenue is maximized. We shall show that these two objectives, namely maximizing seller's expected revenue and buyers' expected payoff, are conflicting, thus rather than finding an optimum we shall formulate the problem as a Stackelberg game and determine the percentage of reimbursement as the Stackelberg strategy of the seller, i.e. the value of q at which seller's revenue is maximized considering the reaction of rational buyers. Let us formalize this in what follows.

We recall that the dynamics of the proposed pricing mechanism implies that the seller announces a percentage of reimbursement q for a service which fails with probability θ. After the announcement, buyers bid to obtain this service, assuming that the probability of failure is $\hat{\theta}$, *a priori* not necessarily equal to θ, and being aware of the value of q. The dynamics of service selling naturally impose an order: the seller announces a value of q and the buyers follow. Each side of the market, seller and buyers, take an action seeking to maximize their own outcome.

This kind of interaction is conveniently modelled by Stackelberg games [44], introduced by von Stackelberg in 1934. In a two-sided Stackelberg game there is a *leader* that plays first, in our case the seller, and the *follower*, in our case the buyers, who plays next knowing the leader's move. We have already introduced the seller's utility through Equation (17). Let us now introduce the utilities of the buyers.

5.3.1 Bidders' Expected Payoff

We now derive the buyer's expected payoff considering their payoff when bidding according to the best bidding strategy, given by Theorem (2). Following the same reasoning as for deriving the best bidding strategy in Section 4 a bidder's payoff is:

$$P = \mathbb{1}_{win}(x\mathbb{1}_{not\, failure} - \beta(x)(1 - q\mathbb{1}_{failure})), \tag{20}$$

where β is expressed in Equation (10) and the *failure* event refers to the real event of failure.

Computing expectations over the event of winning or not and the event of real failures, and replacing β by its definition we obtain the following expression for the expected payoff of each bidder:

$$\begin{aligned} E\{P|X=x\} &= G(x) \\ &\cdot \left[x(1-\theta) - E[Y_{M-1}^{(K)}|Y_{M-1}^{(K)} \leq x] \frac{1-\hat{\theta}}{1-q\hat{\theta}}(1-q\theta) \right]. \end{aligned} \tag{21}$$

In Equation (21) we have considered a given realization of X. Let us now consider the expected payoff prior to having knowledge of this realization, by computing the so-called *ex ante* expected payoff as:

$$\begin{aligned} E\{P\} = E\{E\{P|X=x\}\} &= E\{G(X)X\} \cdot (1-\theta) \\ &- E\{\int_0^X vg(v)dv\} \frac{1-\hat{\theta}}{1-q\hat{\theta}}(1-q\theta), \end{aligned} \tag{22}$$

where the expectation is over the valuations.

5.3.2 The Stackelberg Reimbursement Game

It can be readily noticed from the expected seller's revenue and buyers' payoff in Equations (17) and (22) respectively, which we have reproduced in Table 1 for convenience, that seller's objective and buyers' objective comes to respectively maximizing and minimizing $B_\theta(q, \hat{\theta})$, respectively. They thus have opposing objectives. As aforementioned, because of the nature of the process and since in this subsection buyers are assumed as rational, the situation can be conveniently modelled through a Stackelberg game. More precisely, in our Stackelberg reimbursement game, the leader is the seller and the buyers are followers. The leader's set of available actions is $\{q : q \in \mathbb{R}^+\}$ and the follower's set of available actions is the set $\{\hat{\theta} \in [0,1) : 0 \leq \hat{\theta} < \frac{1}{q}\}$. Finally, the leader's utility is $B_\theta(q, \hat{\theta})$ and the follower's utility is $-B_\theta(q, \hat{\theta})$. Therefore, it is a zero-sum Stackelberg game, and thus the leader's Stackelberg strategy coincides with that of the Nash-equilibrium of the zero-sum game (see e.g. [19]).

Best Bidding Strategy	$\beta(x) = E[Y_{M-1}^{(K)} \mid Y_{M-1}^{(K)} \leq x] \frac{1-\hat{\theta}}{1-q\hat{\theta}}$
Seller's Expected Revenue	$E\{\bar{R}\} = E\{u(K, \mathbf{X})\} B_\theta(q, \hat{\theta})$
Buyers' Expected Payoff	$E\{P\} = E\{G(X)X\} \cdot (1-\theta)$ $-E\{\int_0^X vg(v)dv\} B_\theta(q, \hat{\theta})$.
$B_\theta(q, \hat{\theta}) = \frac{1-\hat{\theta}}{1-q\hat{\theta}}(1-q\theta)$, $(q, \hat{\theta}) \in \{\hat{\theta} \in [0,1), q \geq 0 : q\hat{\theta} < 1\}$	

Table 1: Summary of derived expressions.

In the Stackelberg reimbursement game we have assumed that all buyers would play the same $\hat{\theta}$. This comes directly from the fact that buyers are symmetric.

We shall solve the game through the so-called backward induction method, i.e. the maximization is solved first at the follower's level and this result is in turn used to solve the problem at the leader's level. Let us formalize this solution in the following Theorem.

Theorem 3 *The Stackelberg reimbursement game as defined above has as a solution the set* $\{(q, \hat{\theta}) \in \mathbb{R} \times \mathbb{R} : q = 1,\ 0 \leq \hat{\theta} < 1\}$.

Proof Note that the Stackelberg reimbursement game can be reformulated as the following bi-level optimization problem.

$$\max_q B_\theta(q, \hat{\theta})$$
$$\text{s.t. } q \geq 0,\ \hat{\theta} \in \operatorname*{argmin}_{\hat{\theta}' \in [0,1): q\hat{\theta}' < 1} B_\theta(q, \hat{\theta}')$$

In order to solve Problem (23), the backward induction method is applied. Hence, we first solve the second level optimization. As usual, in order to solve

$$\min_{\hat{\theta} \in [0,1)} B_\theta(q, \hat{\theta}) \tag{23}$$
$$\text{s.t. } 0 \leq \hat{\theta} < 1/q,\ q \geq 0$$

the minimum is found at the values of $\hat{\theta}$ where the first derivative of $B_\theta(q, \hat{\theta})$ with respect to $\hat{\theta}$ is equal to zero or at the border of B_θ's domain. The first derivative of B_θ with respect to $\hat{\theta}$ is

$$\frac{\partial B_\theta(q, \hat{\theta})}{\partial \hat{\theta}} = \frac{q-1}{(1-q\hat{\theta})^2}(1-q\theta), \tag{24}$$

and there is no value of $\hat{\theta}$ that renders it equal to zero. Hence, given that B_θ is a continuous function, the minimum, or infimum, must be reached at the border of its domain. Three cases must be distinguished, namely:

– $0 \leq q < 1$*:* The infimum is attained at $\hat{\theta} = 1$

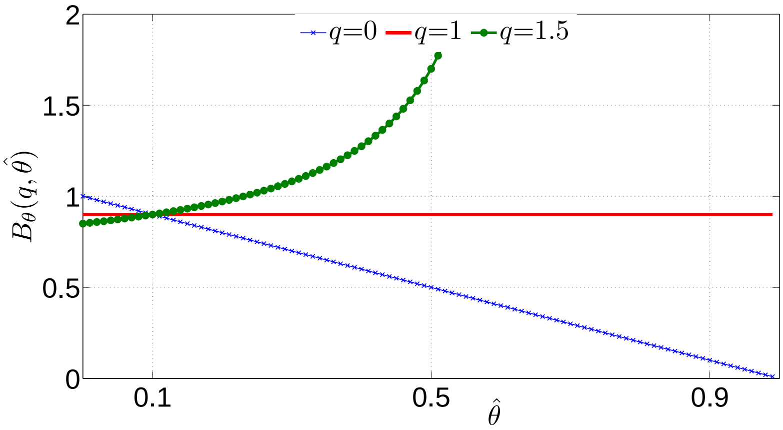

Fig. 4: Variation of B_θ as a function of $\hat{\theta}$ for a real probability of failure $\theta = 0.1$ and for different values of reimbursement. Rational buyers select $\hat{\theta}$ such that it minimizes B_θ.

- $q = 1$: Function B_θ is constant for all $\hat{\theta} \in [0, 1)$
- $1 < q$: The infimum is attained at $\hat{\theta} = 0$ and it is a minimum

This behaviour is shown in Fig. 4, where B_θ is plotted as a function of $\hat{\theta}$ for different values of q.

Finally, the solution to the second level optimization is incorporated to Problem (23) obtaining the following equivalent problem

$$\max_q \{B_\theta(q, 1-\epsilon)\mathbb{1}_{q<1} + B_\theta(1, \hat{\theta})\mathbb{1}_{q=1} + B_\theta(q, 0)\mathbb{1}_{1<q}\},$$

which, evaluating B_θ, can be expressed as

$$\max_q \{\frac{\epsilon}{1 - q(1-\epsilon)}(1 - q\theta)\mathbb{1}_{0\leq q<1} + (1 - q\theta)\mathbb{1}_{1\leq q}\} \tag{25}$$

and where ϵ is an arbitrarily small positive real number. It is easy to see that the solution to the seller's problem is attained at $q = 1$, which concludes the proof. This is the so-called leader's Stackelberg strategy. Fig. 5 shows the behaviour of B_θ for the different cases considered in Problem (25), which illustrates this result. □

5.3.3 Remarks and Interpretations

Interesting interpretations can be derived from the analytical results obtained above.

First, let us highlight the intuition behind the obtained results. If the seller announces a rather small percentage of reimbursement, buyers will, to some

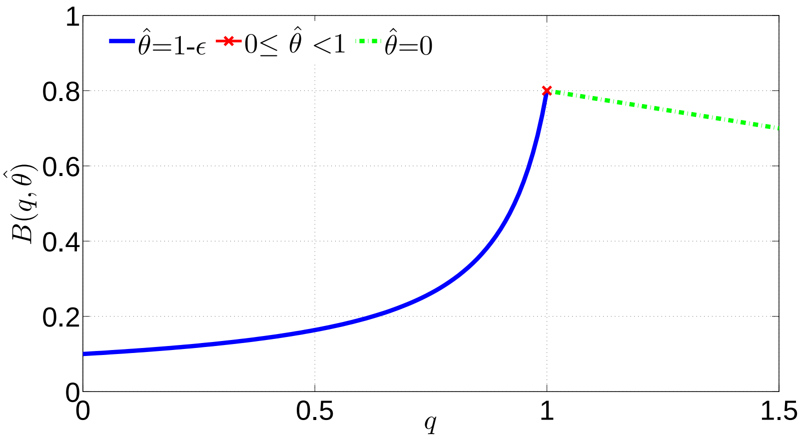

Fig. 5: Variation of the seller's expected revenue as a function of reimbursement q for a real probability of failure $\theta = 0.1$ and a probability of failure estimated by rational buyer's $\hat{\theta}$ equal to their best response for each q. The seller selects q such that it maximizes B_θ.

extent, tend to believe that the service fails a lot, and estimate the probability of failure $\hat{\theta}$ close to 1. This is the so-called *market for lemons* phenomenon, introduced by Akerlof in 1970 [4]. The *market for lemons* states that when buyers are uncertain about the quality of the goods to buy, the market for high quality goods is reduced until it disappears. Indeed, this is what happens according to the theoretical analysis presented above: buyers assume that quality is very bad, which causes the value of the bids to approach zero.

Conversely, if the seller announces a high percentage of reimbursement, greater than 100%, buyers would intuitively assume that failures are not frequent, and thus estimate the probability of failure $\hat{\theta}$ close to 0. We obtain here the so-called *moral hazard* behaviour. That is to say, the buyers take a risk, by considering $\hat{\theta}$ small (theoretically equal to zero), because if a failure were to occur it would be the seller who would bear the cost, through a high reimbursement. This behaviour is observed in many contexts where one of the players taking a decision is not the one bearing the responsibility for this decision. See, for instance, [27] for details on this phenomenon.

All in all, a reimbursement of 100% overcomes the problems that arise when there is asymmetric information. In addition, setting $q = 1$ provides the following three properties, worth highlighting.

Credibility. When the percentage of reimbursement is 100%, and this value is announced to the buyers along with a given probability of failure, then buyers can safely trust the announced probability of failure. Indeed, according to the expression of the seller's expected revenue shown in Equation (17) and illustrated through Fig. 3, when q is set to 1, the seller's expected revenue is constant. The seller thus has no incentives to announce a misleading value for the probability of failure in order to take advantage of naive buyers.

Insensitivity to the buyers' network performance assumption. At the setting $q = 1$, the seller's expected revenue is insensitive to the probability of failure assumed by the buyers. This can be directly seen setting q equal to 1 in Equation (17), which thus renders $E\{\bar{R}\} = E\{u(K, \mathbf{X})\}(1 - \theta)$, which is constant for any value of $\hat{\theta}$. In particular, the seller's expected revenue is the same that he would obtain in the complete information case.

The analogous interpretation from the buyer's standpoint is translated into the following statement.

No value of information. At the setting $q = 1$, knowing the real probability of failure has no value to the buyer from the point of view of his or her expected payoff. Each buyer's expected payoff is the same as when having complete information. This is readily derived from the buyers' expected payoff in Equation (22), which shows that when setting $q = 1$, the buyer's expected payoff is not affected by the assumed probability of failure $\hat{\theta}$. Of course this knowledge could be valuable for the buyers for further reasons not captured in the model.

6 Conclusion

We have proposed a pricing scheme where Assured-Quality Services over data networks are sold via first-price auctions and where in case of failures buyers are reimbursed a certain percentage of what they have paid to obtain the service. The percentage of reimbursement is announced by the seller before the service is sold. Under these conditions and with certain symmetry assumptions among buyers, we have analytically derived the best bidding strategy, which presents an intuitive behaviour. Indeed, for the same level of assumed probability of failure, the higher the percentage of reimbursement, the higher the bid. In addition, for any given percentage of reimbursement lower than 100%, the higher the assumed probability of failure, the lower the bid.

We have addressed the problem of where to set the percentage of reimbursement assuming different situations and buyer behaviours. Namely, we have studied the case of complete information, where buyers perform service monitoring and the case of asymmetric information, where only the seller has knowledge about the real probability of failure of the service on sale.

In the asymmetric information scenario with rational buyers, we have modelled the reimbursement problem as a zero-sum Stackelberg game and shown that the leader's Stackelberg strategy, i.e. the strategy at which seller's revenue is maximized considering the reaction of rational buyers, is given by a reimbursement equal to 100%. In such setting, the *market for lemons* effect and the *moral hazard* one are overcome and seller's and buyers' expected payoff are the same as when having complete information.

In particular our results show that, under the same level of failures, reimbursing 100% provides more revenue in expectation than no reimbursement at all. It must be noted that this output should be compared to the cost of the monitoring infrastructure in order to conclude if the overall balance is positive.

We are currently studying more complex scenarios such as asymmetric bidders, which immediately leads to situations where no close forms can be derived for the best bidding strategies. For such situations, we are studying numerical approaches, in order to draw similar or novel conclusions. We are also working to relax other assumptions such as the independence among bidders in order to take into account collusion effects. Future work will also consider probability of failure as a function of the accepted bandwidth, in order to add feedback from the accepted traffic into pricing.

A Proof of Theorem 1

Bidder's payoff is given by

$$p = \mathbb{1}_{win}(x\mathbb{1}_{not\,failure} - b(1 - q\mathbb{1}_{failure})). \tag{A.1}$$

Assume first that equilibrium occurs when all bidders use the same bidding function $\beta(x)$.

The expected payoff of any given bidder is:

$$\tilde{P} = E\{p|X = x\} = G(\beta^{-1}(b))(x(1 - \hat{\theta}) - b(1 - q\hat{\theta})). \tag{A.2}$$

Computing the derivative of Equation (A.2) and making it equal to zero we obtain:

$$\frac{g(\beta^{-1}(b))}{\beta'(\beta^{-1}(b))}(x(1 - \hat{\theta}) - b(1 - q\hat{\theta})) - G(\beta^{-1}(b))(1 - q\hat{\theta}) = 0, \tag{A.3}$$

where we have applied the well-known formula for the derivative of the inverse function.

Under the assumption of symmetric equilibrium $\beta^{-1}(b) = x$ holds. Applying this equality to Equation (A.3) we obtain:

$$xg(x)(1 - \hat{\theta}) - g(x)\beta(x)(1 - q\hat{\theta}) - G(x)\beta'(x)(1 - q\hat{\theta}) = 0 \tag{A.4}$$

The study must be divided into two cases, namely $q\hat{\theta} < 1$ and $q\hat{\theta} \geq 1$. We assume as well that $\beta(0) = 0$.

$q\hat{\theta} < 1$. First consider $\hat{\theta} \neq 1$. In this case Equation (A.4) can be rewritten as:

$$\beta'(x) + \beta(x)\frac{g(x)}{G(x)} - x\frac{g(x)}{G(x)}\frac{1 - \tilde{\theta}}{1 - q\hat{\theta}} = 0 \tag{A.5}$$

whose solution is

$$\beta(x) = -e^{-S(x)} \int_0^x e^{S(z)} T(z) dz, \tag{A.6}$$

where:

$$S(x) = \int_0^x \frac{g(z)}{G(z)} dz = \log G(x) \text{ and } T(x) = -x\frac{g(x)}{G(x)}\frac{1 - \tilde{\theta}}{1 - q\hat{\theta}}.$$

Hence, operating we obtain:

$$\begin{aligned} \beta(x) &= \frac{1}{G(x)} \int_0^x zg(z)dz \frac{1 - \hat{\theta}}{1 - q\hat{\theta}} \\ &= E[Y_{M-1}^{(1)}|Y_{M-1}^{(1)} \leq x] \frac{1 - \hat{\theta}}{1 - q\hat{\theta}}, \end{aligned} \tag{A.7}$$

where the last equality comes directly from the definition of conditional expectation.

Consider now $\hat{\theta} = 1$. Equation (A.3) reduces to

$$\frac{g(\beta^{-1}(b))}{\beta'(\beta^{-1}(b))}(-b(1-q\hat{\theta})) - G(\beta^{-1}(b))(1-q\hat{\theta}) = 0, \tag{A.8}$$

which results in

$$\frac{-g(x)}{G(x)} = \frac{\beta'(x)}{\beta(x)}. \tag{A.9}$$

Integrating Equation (A.9) on both sides we obtain

$$\beta(x) = \frac{\kappa}{G(x)}, \tag{A.10}$$

where κ is a real constant of integration.

In order to verify whether the assumption of symmetric bidding functions holds, we suppose, without loss of generality, that all bidders but one bid with the same optimal bidding function found above. We shall check if it is also optimal for the remaining bidder to bid according to this function.

Consider first the case of $\hat{\theta} \neq 1$. Bidder 1's expected payoff ($\tilde{P}$) if he or she bids $\beta(z)$ when his or her value is actually x is:

$$\tilde{P}(\beta(z), x) = G(z)(x(1-\hat{\theta}) - \beta(z)(1-q\hat{\theta})). \tag{A.11}$$

Hence, the difference with the bidder's expected payoff if bidding $\beta(x)$ is:

$$\begin{aligned}
&\tilde{P}(\beta(z), x) - \tilde{P}(\beta(x), x) = \\
&\quad G(z)(x(1-\hat{\theta}) - \beta(z)(1-q\hat{\theta})) - \\
&\quad G(x)(x(1-\hat{\theta}) - \beta(x)(1-q\hat{\theta})) \\
&\quad = (1-\hat{\theta})x(G(z) - \\
&\quad - G(x)) + (1-\hat{\theta}) \int_z^x vg(v)dv \\
&\quad = (1-\hat{\theta})x(G(z) - G(x)) + (1-\hat{\theta}) \left[G(v)v \mid_z^x - \int_z^x G(v)dv \right] \\
&\quad = (1-\hat{\theta}) \left[G(z)(x-z) + \int_x^z G(v)dv \right],
\end{aligned} \tag{A.12}$$

where we have applied integration by parts. Equation (A.12) is negative for any value of $z \neq x$, as detailed in Section 4.

Consider now the case $q < 1$ and $\hat{\theta} = 1$. β is given by Equation (A.10), where κ is such that $\beta(0) = 0$. But since $G(0) = 0$, it follows that there is no value of κ, either than the trivial $\kappa = 0$, verifying $\beta(0) = 0$. We have obtained a contradiction, hence we conclude that the assumption of a symmetric equilibrium is not possible, and no equilibrium exists for this case.

$q\hat{\theta} \geq 1$. In this case we notice that the expected payoff, given in Equation (A.2), is always increasing with b, thus there is no equilibrium. □

Acknowledgments

This work was funded by the ETICS project (EC FP7 248567)[1], the Uruguayan Agency for Research and Innovation (PR-POS-2008-003) and the SticAmSud project 11STIC-02 MAITRE.

References

1. ETICS: Economics and Technologies for Inter-carrier Services. European FP7 Research Project [Online] Available: http://www.ict-etics.eu.
2. RFC 6812: Cisco Service-Level Assurance Protocol [Online] Available: http://www.rfc-editor.org/rfc/rfc6812.txt.
3. Cisco Systems: Entering the Zettabyte Era. Technical report, May 2011.
4. G. A. Akerlof. The Market for Lemons: Quality Uncertainty and the Market Mechanism. *The Quarterly Journal of Economics*, 84(3):488–500, 1970.
5. A. Al Daoud, T. Alpcan, S. Agarwal, and M. Alanyali. A Stackelberg game for pricing uplink power in wide-band cognitive radio networks. In *47th IEEE Conference on Decision and Control*, pages 1422–1427. IEEE, December 2008.
6. E. Altman, R. Marquez, R. El-Azouzi, D. Ros, and B. Tuffin. Stackelberg approach for pricing differentiated services. In *Proceedings of the 2nd international conference on Performance evaluation methodologies and tools*, ValueTools '07.
7. E. Altman, J. Rojas, S. Wong, M. Hanawal, and Y. Xu. Net Neutrality and Quality of Service. In R. Jain and R. Kannan, editors, *Game Theory for Networks*, volume 75 of *Lecture Notes of the Institute for Computer Sciences, Social Informatics and Telecommunications Engineering*, pages 137–152. Springer Berlin Heidelberg, 2012.
8. I. Amigo, P. Belzarena, F. Larroca, and S. Vaton. Network Bandwidth allocation with end-to-end QoS constraints and Revenue Sharing in Multidomain Federations. In *ICQT 2011*, pages 50–62, Berlin, Heidelberg, 2011. Springer-Verlag.
9. D. Barth, L. Echabbi, and C. Hamlaoui. Stock provisioning strategies for QoS at the interdomain level. *Telecommunication Systems*, 52(2):623-632, February 2013.
10. T. Basar and R. Srikant. A Stackelberg Network Game with a Large Number of Followers. *J. Optimization Theory and Applications*, 115(3):479–490, December 2002.
11. A. Belghith, S. Said, B. Cousin, and S. Lahoud. Proactive and reactive collaboration schemes for multi-domain networks monitoring. In *Computing, Communications and Applications Conference*, pages 150–157, January 2012.
12. P. Belzarena, A. Ferragut, and F. Paganini. Bandwidth Allocation via Distributed Auctions with Time Reservations. In *IEEE INFOCOM, Procedings. Rio de Janeiro, Brazil*, 2009.
13. H. K. Bhargava and S. Sundaresan. Contingent Bids in Auctions: Availability, Commitment and Pricing of Computing as Utility. In *System Sciences, 2004. Proceedings of the 37th Annual Hawaii International Conference on*, page 10. IEEE, January 2004.
14. F. Boussion, P. Maille, and B. Tuffin. Net neutrality debate: Impact of competition among ISPs. In *Communication Systems and Networks (COMSNETS), 2012 Fourth International Conference on*, pages 1–8, 2012.
15. C. Courcoubetis and R. Weber. *Pricing and Communications Networks.* John Wiley & Sons, Ltd, 2003.
16. M. Dramitinos, G. D. Stamoulis, and C. Courcoubetis. An Auction Mechanism for Allocating the Bandwidth of Networks to their Users. *Computer Networks*, 51:4979–4996, December 2007.
17. N. Economides and J. Tåg. Network neutrality on the Internet: A two-sided market analysis. *Information Economics and Policy*, 24(2):91–104, 2012.
18. N. Feng, S. Mau, and N. B. Mandayam. Pricing and Power Control for Joint Network-centric and User-centric Radio Resource Management. *Communications, IEEE Transactions on*, 52(9):1547–1557, september 2004.
19. D. Fundenberg and J. Tirole. *Game Theory.* Mit Press, 1991.
20. C. Homburg, N. Koschate, and W. D. Hoyer. Do Satisfied Customers Really Pay More? A Study of the Relationship Between Customer Satisfaction and Willingness to Pay. *Journal of Marketing*, 69(2):84–96, Avril 2005.
21. V. Krishna. *Auction Theory.* Academic Press/Elsevier, 2009.
22. H. Li and J. G. Riley. Auction choice. *International Journal of Industrial Organization*, 25(6):1269–1298, December 2007.
23. J. MacKie-Mason. A Smart Market for Resource Reservation in Multiple Quality of Service Information Networks. Technical report, 1997. University of Michigan, USA.

24. P. Maillé and B. Tuffin. Pricing the Internet with Multibid Auctions. *IEEE/ACM Transactions on Networking*, 14:992–1004, October 2006.
25. P. Maillé and B. Tuffin. Why VCG Auctions can Hardly be Applied to the Pricing of Inter-domain and Ad hoc Networks. In *3rd EuroNGI conference, Trondheim, Norway*, pages 36 – 39, 2007.
26. D. P. Mann and J. P. Wissink. Money-Back Contracts with Double Moral Hazard. *The RAND Journal of Economics*, 19(2):285-292, 1988.
27. A. Mas-Colell, M. Whinston, and J. Green. *Microeconomic Theory.* Oxford student edition. Oxford University Press, 1995.
28. P. R. Milgrom and R. J. Weber. A Theory of Auctions and Competitive Bidding. *Econometrica*, 50(5):1089–1122, 1982.
29. J. Musacchio, G. Schwartz, and J. Walrand. A Two-Sided Market Analysis of Provider Investment Incentives with an Application to the Net-Neutrality Issue. *Review of Network Economics*, 8(1):3, 2009.
30. P. S. Neelakanta and R. C. Sardenberg. Consumer benefit versus price elasticity of demand: a nonlinear complex system model of pricing internet services on QoS-centric architecture. *NETNOMICS: Economic Research and Electronic Networking*, 12:31–60, 2011.
31. M. J. Osborne and A. Rubinstein. *A Course in Game Theory.* The MIT Press, first edition, July 1994.
32. H. Pouyllau and G. Carofiglio. Inter-carrier SLA negotiation using Q-learning. *Telecommunication Systems*, 52(2):611-622, February 2013.
33. P. Reichl. From Charging for Quality of Service to Charging for Quality of Experience. *Annals of Telecommunications*, 65(3):189–199, April 2010.
34. P. Reichl, S. Bessler, and B. Stiller. Second-Chance Auctions for Multimedia Session Pricing. In G. Ventre and R. Canonico, editors, *Interactive Multimedia on Next Generation Networks*, volume 2899 of *LNCS*, pages 306–318. Springer, 2003.
35. P. Reichl, B. Tuffin, and R. Schatz. Logarithmic laws in service quality perception: where microeconomics meets psychophysics and quality of experience. *Telecommunication Systems*, 52(2):587-600, February 2013.
36. P. J. Reny. On the Existence of Pure and Mixed Strategy Nash Equilibria in Discontinuous Games. *Econometrica*, 67(5):1029–1056, 1999.
37. J. G. Riley and W. F. Samuelson. Optimal Auctions. *American Economic Review*, 71(3):381–92, June 1981.
38. A. Sackl, P. Zwickl, and P. Reichl. From Quality of Experience to Willingness to Pay for Interconnection Service Quality. In *NETWORKING 2012 Workshops*, volume 7291 of *LNCS*, pages 89–96. Springer Berlin Heidelberg, 2012.
39. E. Salhi, S. Lahoud, and B. Cousin. Global Vs. Per-domain monitoring of multi-domain networks. In *Proceedings of the 2011 IEEE 36th Conference on Local Computer Networks*, LCN '11, pages 391–398, Washington, DC, USA, 2011. IEEE Computer Society.
40. H. Shen and T. Basar. Differentiated Internet pricing using a hierarchical network game model. In *Proceedings of the American Control Conference*, volume 3, pages 2322–2327 vol.3. IEEE, June 2004.
41. T. Trinh, L. Gyarmati, and G. Sallai. Understanding the impact of loyal user behaviour on Internet access pricing: a game-theoretic framework. *Telecommunication Systems*, 48:43-61, 2011.
42. B. Tuffin, H. Le Cadre, and M. Bouhtou. Optimal Pricing Strategy with Compensation when QoS is not Satisfied. *Telecommunication Annals*, 65(3-4):163–170, 2010.
43. W. Vickrey. Counterspeculation, Auctions, and Competitive Sealed Tenders. *The Journal of Finance*, 16(1):8–37, 1961.
44. H. von Stackelberg. *Market Structure and Equilibrium.* Springer, 1st edition, December 2010.
45. B. Wang, Y. Wu, and K. J. R. Liu. Game theory for cognitive radio networks: An overview. *Computer Networks*, 54(14):2537–2561, October 2010.
46. D. Wang and Z. Li. On Traffic Pricing Model Based on Incentive Stackelberg Strategies for Communication Networks. In *Communications and Information Processing*, volume 288 of *Communications in Computer and Information Science*, pages 498–505. Springer Berlin Heidelberg, 2012.
47. F. Zhang and P. Verma. Pricing multi-class network services using the Shapley Value. *NETNOMICS: Economic Research and Electronic Networking*, 12:61-75, 2011.