Algorithmica
DOI 10.1007/s00453-007-9107-5

Truth-Telling Reservations

Fang Wu · Li Zhang · Bernardo A. Huberman

Received: 8 February 2006 / Accepted: 10 March 2007

Abstract We present a mechanism for reservations of bursty resources that is both truthful and robust. It consists of option contracts whose pricing structure induces users to reveal the true likelihoods that they will purchase a given resource. Users are also allowed to adjust their options as their likelihood changes. This scheme helps users save cost and the providers to plan ahead so as to reduce the risk of under-utilization and overbooking. The mechanism extracts revenue similar to that of a monopoly provider practicing temporal pricing discrimination with a user population whose preference distribution is known in advance.

Keywords Mechanism design · Truth-telling · Reservation · Option · Contract · Incentive compatible

1 Introduction

A number of compute intensive applications often suffer from bursty usage patterns [1, 3, 6, 7, 10], whereby the demand for information technology (IT) resources, such as memory and bandwidth, can at times exceed the installed capacity within the organization. This problem can be addressed by providers of IT services who satisfy this peak demand for a given price, playing a role similar to utilities such as electricity or natural gas.

The emergence of a utility form of IT provisioning creates a number of problems for both providers and customers due to the uncertain nature of IT usage. On the provider side there is a need to design appropriate pricing schemes to encourage the use of such services and to gain better estimates of the usage pattern so as to enable effective statistical multiplexing. On the customer side, there needs to be a simple

F. Wu (✉) · L. Zhang · B.A. Huberman
HP Labs, Palo Alto, CA 94304, USA
e-mail: fang.wu@hp.com

way of figuring out how to anticipate and hedge the need for uncertain demand as well as the costs that it will add to the overall IT operations.

Recently, a proposal was made to use swing options for pricing the reservation of IT resources [3]. By purchasing a swing option the user pays an upfront premium to acquire the right, but not the obligation, to use a resource as defined in the option contract. As with the case with electricity, IT resources, such as bandwidth and CPU time, are non-storable and with volatile usage pattern. Thus, swing options provide flexibility in both the amount and the time period for which a resource is purchased, making them appealing to users whose bursty demand is hard to predict. From the point of view of the providers, if enough users purchase these options providers can offset the cost of providing peak capacity by multiplexing among many users.

Pricing a swing option for IT resources however, turns out to be difficult because of the complexity of the option contract and the lack of a good model of the spot market price. Moreover, there are two other important problems that need resolution. First, the user needs to be able to estimate the amount of resources that need to be reserved as well as their cost; and second, the provider needs to put in place a mechanism that will induce truth revelation on the part of the user when stating the likelihood that a given reservation or option will be exercised.

As was shown in [3], the first problem can be addressed by providing the user with a simulation tool for estimating the cost of a reservation from a set of historical data, as well as a provision for entering the user's assumptions about aggressive or conservative swings. Because the prices for swings are set ahead of time and not by market forces, the forecasting tool also provides a powerful "what-if" capability to both the resource provider and the customer for estimating outright costs and risks associated with fluctuations in customer demand.

As to the provider's problem with asymmetric information, it could be argued that a user's historical usage pattern allows to predict his future demand. But in many cases, such as with new users, the data may not be available or reflect unanticipated user needs. Even worse, users may intentionally misrepresent the likelihood of their needs in order to gain a pricing advantage, with the consequent loss to the provider. While in the swing option this was addressed by introducing a time dependent discount that induces early commitment to a contract, users can still misrepresent their likelihoods the first time they buy an option.

This paper presents a solution to the truth revelation problem in reservations by designing option contracts with a pricing structure that induces users to reveal their true likelihoods that they will purchase a given resource. A user is allowed to adjust his option later if his likelihood changes. Truthful revelation helps the provider to plan ahead to reduce the risk of under-utilization and overbooking, and also helps the users to save cost. In addition to its truthfulness and robustness, the mechanism extracts revenue similar to that of a monopoly provider practicing temporal pricing discrimination with a user population whose preference distribution is known in advance [2, 5, 6, 8–10, 12].

Besides its intrinsic appeal, the truth telling mechanism we propose may be useful in many areas. For example, airline seats, hotel reservations, network bandwidth and tickets for popular shows would benefit from a properly priced reservation system, leading to both more predictable use and revenue generation.

2 The Two Period Model

2.1 Motivating Example

Consider n users $\{1, 2, \ldots, n\}$ who live for two discrete periods. With probability p_i, user i may want to consume one unit of resource in period 2, which is worth to him a value v_i. With probability $1 - p_i$ he does not need it in period 2 (zero value). He can either reserve one unit of resource from the service provider in period 1 at a discount price, or buy it in period 2 at a higher last-minute price. His utility is his value minus his payment. Each user i knows his probability of usage p_i in period 1, and it is not until period 2 that he can be certain about his real need (unless $p = 0$ or 1). Assume it costs the service provider 1 to prepare one unit of resource in period 1, and $C \geq 1$ to produce it in period 2. Also assume $v_i \geq 1$ for all i. Suppose all users are rational and risk-neutral, and the distributions of their needs are independent.

To motivate our idea, let us consider a concrete example: $C = 5$, $v_i = 4$, $p_i = 0.1$. In this setting the user will not reserve in period 1 because his expected value of one unit then is $v_i p_i = 0.4$, less than the service provider's cost 1. He definitely will not buy it in period 2 either, when his value, 4, is less than the service provider's cost 5. There can be no business.

The reason for this is that the user cannot afford the resource in period 2, yet he is unwilling to reserve it in period 1 when it is cheap because of uncertainty. Although he values the unit more than the reservation price, he will make a reservation only when he is relatively certain about his later usage. When he is "not that certain" (p small), there is no way that he can make a reservation at a lower price.

In what follows we describe a reservation mechanism that allows the user to pay a small premium that guarantees him one unit of resource whenever he needs it in period 2, at a price not much higher than the discount price 1. In addition, the mechanism makes the user truthfully reveal his probability of using the resource to the provider, who can then accurately anticipate user demand and benefit from planning. At a later stage, we show how this mechanism can be thought of as an option.

Remark In our model, we assume that the cost at period 2 is fixed at $C > 1$. This is to capture the fact that for many resources, it is expensive to prepare, and therefore expensive to acquire, on spot. Either under-provisioning or over-provisioning at period 1 can cause a loss to the provider. Thus, a truth-telling mechanism helps the provider plan in period 1, as we shall see later in this section. In addition, it provides a better incentive to the users compared to the fixed price selling scheme, also shown in this section. One may argue that for a monopoly seller, if he knows the users utility and probability of usage, he may come up with the optimum, likely a temporally discriminated, fixed price selling scheme that maximizes his profit. For resources with volatile and bursty demand, such as IT resources, it is often difficult to obtain the users' usage information. Even when such information is available, the reservation mechanism may extract more profits as we show by example in Sect. 4.

2.2 The Coordinator Game

To better illustrate the benefit of this mechanism, we introduce a third agent, the coordinator, who aggregates the users' probabilities and makes a profit while absorbing the users' risk. The coordinator can be either a real middle-agent (e.g. third-party travel agents such as Expedia and Travelocity) or the service provider himself. We assume that the coordinator can obtain one unit of resource from the provider at prices 1 and C, in period 1 and 2, respectively. (When the coordinator is the service provider himself this is obvious.) We first propose a two-stage truth-telling reservation mechanism which is beneficial to the coordinator and the users. Next, we formulate the profit maximization problem for the coordinator.

Remark While the introduction of a coordinator is for analysis convenience, there are also good reasons for such middle agents to exist in many situations when they can be more efficient to aggregate information or absorb market risks, especially when there are multiple sellers.

1. (Period 1) The coordinator asks each user to submit a probability q_i.
2. (Period 1) The coordinator reserves $\sum q_i$ units of resource from the resource provider (at price 1), ready to be consumed in period 2.
3. (Period 2) The coordinator delivers the reserved resource units to users who claim them. If the amount he reserved is not enough to satisfy the demand, he buys more resource from the provider (at the higher unit price C) to meet the demand.
4. (Period 2) User i pays

$$\begin{cases} f(q_i) & \text{if he needs one unit of resource,} \\ g(q_i) & \text{if he does not need it,} \end{cases} \tag{1}$$

   where $f, g : [0, 1] \to \mathbb{R}^+$ are two functions whose forms will be specified later.

These terms are publicly announced to everyone, before step 1.

For the coordinator to profit, the following two conditions have to be satisfied:

Condition A *The coordinator can make a profit by providing this service.*

Condition B *Each user prefers to use the new service provided by the coordinator.*

The next two truth-telling conditions, although not absolutely necessary, are useful for conditions A and B to hold.

Condition T1 (Step 1 truth-telling) *Each user submits his true probability p_i in step* 1, *so that he expects to pay the least later in step* 4.

Condition T2 (Step 3 truth-telling) *In step* 3, *when a user does not need a resource in period* 2, *he reports it to the coordinator.*

From Condition T1, user i expects to pay $w(q_i) \equiv p_i f(q_i) + (1 - p_i)g(q_i)$ in period 2. His optimal submission q_i^* is determined by the first-order condition $w'(q_i^*) = p_i f'(q_i^*) + (1 - p_i)g'(q_i^*) = 0$. Truth-telling requires that $q_i^* = p_i$, or $p_i f'(p_i) + (1 - p_i)g'(p_i) = 0$. Condition T2 simply requires that $f(p) \geq g(p)$ for all $p \in [0, 1]$.

Now we study Condition A when all users submit their true probabilities $\{p_i\}$. Let U be the total resource usage of all users in period 2, and let W be the their total payment. Both U and W are random variables. Clearly, $\mathbb{E}\,U = \sum p_i$, and $\mathbb{E}\,W = \sum w(p_i)$.

Lemma 1 *If there exists an arbitrarily small* $\epsilon > 0$ *such that* $w(p) \geq p + \epsilon$ *for all* $p \in [0, 1]$, *then* $W - U \to \infty$ *a.s. as* $n \to \infty$. *That is, by charging an arbitrarily small premium, the coordinator makes profit when there are many users* (*Condition* A).

Proof This follows directly from the "X^4-strong law". (See e.g. [11]. The random usage of each user does not have to be identically distributed.) □

The small number ϵ is merely a technical device. In what follows we will neglect it and use a weakened condition of Lemma 1 as a sufficient condition of Condition A (not rigorous): $w(p) \geq p$ for all $p \in [0, 1]$.

Last, Condition B says two things. First, the user should not lose utility: $w(p) \leq vp$. Second, the user prefers to use the coordinator's service rather than to deal with the resource provider directly: $w(p) \leq \min(1, Cp)$.

To summarize, the following conditions on concave f and g are sufficient for the truth-telling mechanism to work:

$$pf'(p) + (1-p)g'(p) = 0, \tag{2}$$

$$f(p) \geq g(p), \tag{3}$$

$$p \leq pf(p) + (1-p)g(p) \leq \min(1, Cp, vp), \tag{4}$$

for all $p \in [0, 1]$. Under these conditions a non-fictitious third-party coordinator is profitable. Also, as we have seen in the motivation example, when $v \leq C$ and $vp \leq 1$ the service provider himself would have incentive to play the coordinator. We leave the more subtle $v > C$ case to Sect. 4.

Consider the following choice[1]

$$f(p) = 1 + \frac{k}{2} - kp + \frac{kp^2}{2}, \tag{5}$$

$$g(p) = \frac{kp^2}{2}, \tag{6}$$

which satisfies (2). To check (3) and (4), we first calculate

$$w(p) = \left(1 + \frac{k}{2}\right) p - \frac{k}{2} p^2. \tag{7}$$

[1] This choice is not unique, but is analytically simple. For example, we could have chosen $g \propto p^a$ ($a > 1$), which will have no essential impact on the rest of the paper.

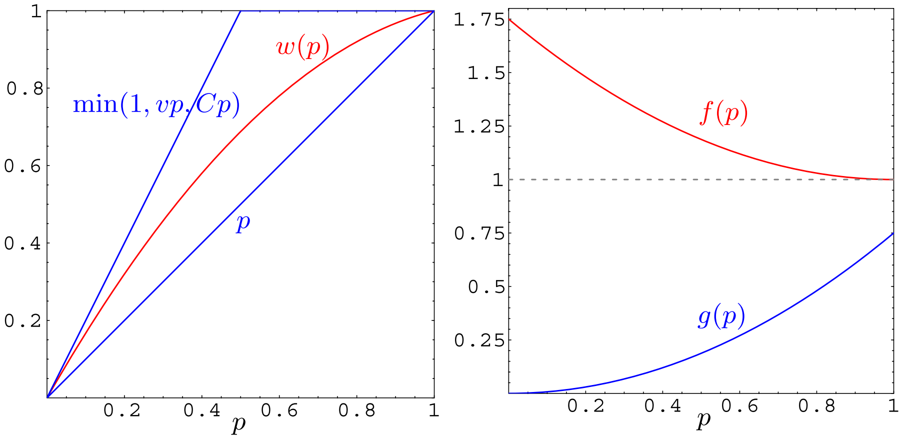

Fig. 1 Both figures are plotted under the choice $v = C = 2$ and $k = 1.5$. **a** Plot of $w(p)$. The curve $w(p)$ lies completely in the triangular region. **b** Payment curves of the user. If the user needs one unit of resource in period 2, he pays according to the *upper curve*. Otherwise he pays according to the *lower curve*

And then it is not hard to show

Lemma 2 *For the choice of f and g in* (5) *and* (6), *conditions* (3) *and* (4) *are satisfied for* $k \in [0, \min\{2(v-1), 2(C-1), 2\}]$.

Proof Equation (3) is satisfied because

$$f(p) - g(p) = 1 + k\left(\frac{1}{2} - p\right) \geq 1 - \frac{k}{2} \geq 0. \tag{8}$$

To verify (4), we write

$$w(p) = p + \frac{k}{2}p(1-p) \geq p, \tag{9}$$

$$w(p) = 1 - (1-p)\left(1 - \frac{k}{2}p\right) \leq 1, \tag{10}$$

$$w(p) \leq \min(v, C)p - \frac{k}{2}p^2 \leq \min(v, C)p. \tag{11}$$

□

Figure 1a shows the special case $v = C = 2$ and $k = 1.5$. As can be seen, the curve $w(p)$ lies completely in the triangular region. The difference between the upper blue curve and the red curve is the amount of money the user saves (varying with different p). The difference between the red curve and the lower blue line is the coordinator's expected payoff from one user. Note that his payoff is larger for values of p's lying in the middle of the range, and is zero for $p = 0$ and $p = 1$. This result is hardly surprising, for when there is no uncertainty the user does not need a

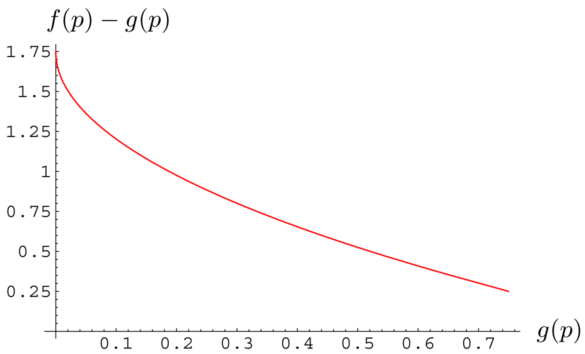

Fig. 2 Price-premium curve for the user. The *horizontal axis* is the premium (value of contract, value of option). The *vertical axis* is the price the user pays for one unit of resource in period 2

coordinator at all. Thus, the coordinator makes a profit out of uncertainties in user behavior.

Figure 1b plots the two payment curves, $f(p)$ and $g(p)$, for the same choice of parameters. After signing a contract, a user agrees to pay later either the upper curve for one unit of resource, or the lower curve for nothing. Note that $f(p)$ is strictly decreasing, a feature essential for the user to be truth-telling. A user with a high p is more likely to pay the upper curve rather than the lower curve. Knowing this, he has an incentive to submit a high probability of use and thus not to cheat.

2.3 The Coordinator's Optimization Problem

So far we have not discussed what particular k to choose in the interval specified in Lemma 2. Suppose the coordinator believes in a prior distribution $f(v, p)$ of the users' (v, p). If he seeks to maximize his profit, he would choose the k that solves

$$\max_{k\in[0,\min\{2(C-1),2\}]} \int\int \frac{k}{2}(p-p^2) I\left(v \geq \frac{k}{2}+1\right) f(v,p)\, dv\, dp, \tag{12}$$

where $k(p-p^2)/2 = w(p) - p$ is his profit from one user, and $v \geq k/2+1$ is the condition for the user to participate.

2.4 The Reservation Contract as an Option

The contract discussed in previous sections can be equivalently regarded as an "option". Because $g(p)$ is the minimal amount the user has to pay in any event, we can ask him to pay it in period 1, and only to pay $f(p) - g(p)$ in period 2 if he needs one unit of resource at that time. Hence, by paying an amount $g(p)$, the user achieves the right but no the obligation to buy one unit of resource at price $f(p)-g(p)$ in period 2. Naturally, we may call $g(p)$ the *premium* or the *price of option*, and $f(p) - g(p)$ the *price of the resource*.

Figure 2 shows the parametric plot of resource price versus option price, for $p \in [0, 1]$. Instead of submitting an explicit p, the user can equivalently choose one point on this curve and pay accordingly. His probability p can then be inferred from his choice (using (5) or (6)). This alternative method may be more user-friendly because people tend to be more sensitive to monetary values rather than probabilities.

Fig. 3 The information structure for three periods

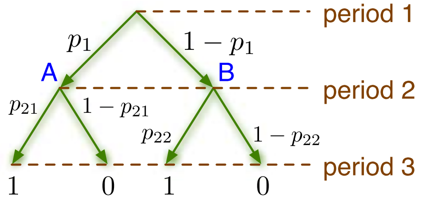

We can even further simplify the curve by providing the user with a table with the values of a few discrete points along the curve.

3 A Multi Period Truth-Telling Reservation

In the previous 2-period mechanism, if a user learns more in time about the likelihood of his needing the resource, it is impossible for him to modify the original contract. To solve this issue we extend our mechanism so that the user can both submit early for a larger discount and update his probability afterwards to a more accurate one. We thus consider a dynamic extension of the problem in which the user is allowed to change his probability of future use some time after his initial submission.

3.1 The Information Structure

Assume that everyone lives for m periods. In period m the user might need to consume one unit of resource. He can reserve/buy one unit in period i at price C^{i-1}, for $i = 1, \dots, m$.[2] The intermediate periods are introduced to exploit the user's information gaining process. We assume that at each period i, the user can always make a "best guess" of his probability of usage, given his information up to period i. Formally, his "best guess" can be described by a random process p_t adapted to an information filtration [4], satisfying the property $\mathbb{E}_t p_{t+1} = p_t$.

As an illustrative example, consider the three period information structure depicted in Fig. 3. The user enters state A with probability p_1 and state B with probability $1 - p_1$. If he enters state A, with probability p_{21} he will need the resource in period 3. If he enters state B, he will need the resource with probability p_{22}. Under our requirement, the user is able to make the best guess $p = p_1 p_{21} + (1 - p_1) p_{22}$ in period 1, which will change to either p_{21} or p_{22} in period 2.

3.2 The Three Period Coordinator Game

Again we describe a mechanism used by a coordinator to make profit by aggregating the user's uncertainty. A user may submit a probability in period 1, as in the 2-period setting. Additionally, when he enters period 2 he is allowed to update his probability based on his new information gained at that time. This way the user can enjoy the

[2] The $(1, C, \dots, C^{m-1})$ assumption is not essential. We could have assumed $(1, C_2, \dots, C_m)$ instead and the main result of this section will continue to hold, just that the maths would become considerably messier.

Table 1 The user's payment table. The columns represent his three possible submission patterns

	q_1 not q_2	q_2 not q_1	Both q_1 and q_2
Uses one unit	$f_1(q_1)$	$f_2(q_2)$	$f_1(q_1) - \alpha f_2(q_1) + \alpha f_2(q_2)$
Does not use	$g_1(q_1)$	$g_2(q_2)$	$g_1(q_1) - \alpha g_2(q_1) + \alpha g_2(q_2)$

full discount while simultaneously utilize maximum information. His final payment in period 3 is determined by the one or two probabilities he submitted. The whole mechanism is described more rigorously as follows.

1. (Period 1) The user may submit a probability q_1.
2. (Period 1) The coordinator reserves q_1 units of resource from the resource provider (at price 1).
3. (Period 2) The user may submit a probability q_2.
4. (Period 2) The coordinator adjusts his holdings to match the new probability q_2.
5. (Period 3) If the user claims the need of one unit of resource, the coordinator delivers one reserved unit to him. If his reservation pool is not large enough, he buys more resource from the provider (at the higher unit price C^2) to meet the demand.
6. (Period 3) The user pays according to Table 1.

In Table 1, (f_1, g_1) and (f_2, g_2) are two sets of 2-period truth-telling functions solved in Sect. 2.2.

$$f_1(p) = 1 + \frac{k_1}{2} - k_1 p + \frac{k_1 p^2}{2}, \qquad g_1(p) = \frac{k_1 p^2}{2}, \tag{13}$$

$$f_2(p) = C + \frac{k_2}{2} - k_2 p + \frac{k_2 p^2}{2}, \qquad g_2(p) = \frac{k_2 p^2}{2}, \tag{14}$$

where $k_1 \in [0, \min\{2(C^2 - 1), 2\}]$ and $k_2 \in [0, \min\{2(C^2 - C), 2C\}]$. To make the mathematical analysis easier, we will choose $k_1 = k \in [0, \min\{2(C-1), 2\}]$ and $k_2 = Ck \in [0, \min\{2(C^2 - C), 2C\}]$, so that $f_2(p) = Cf_1(p)$ and $g_2(p) = Cg_1(p)$. We require that $f_2 > f_1$ and $g_2 > g_1$, so that the user pays more when he reserves late. In order to focus on truth-telling properties rather than incentive issues, we assume $v \geq C^2$ so the user will always buy the resource from the provider or the coordinator if he needs it in period 2.

Theorem 1 *Suppose $\alpha \in (0, 1/C)$. The user's optimal strategy is to submit a probability in period 1 and to adjust it in period 2. Each probability he submits is his true probability in that period. In addition, the coordinator is profitable.*

Proof Follows from the four lemmas in the Appendix. □

3.3 Three Period Options

As for the 2-period problem, there is an equivalent "option" form of the 3-period contract, which we now describe. Assume $f_2(p) = Cf_1(p)$ and $g_2(p) = Cg_1(p)$.

1. (Period 1) There are various options that the user can buy, with option price $g_1(p)$ and resource price $f_1(p) - g_1(p)$, for all $p \in [0, 1]$. The user buys one share of q_1-option at price $g_1(q_1)$.
2. (Period 2) The user can swap αC (remember $\alpha C < 1$) share of his q_1-option for a q_2-option, by paying the difference price $\alpha C(g_1(q_2) - g_1(q_1))$. Then he holds a share $(1 - \alpha C)$ of q_1-options and a share αC of q_2-options.
3. (Period 3) If the user needs one unit of resource, he executes his options. That is, he pays $(1-\alpha C)(f_1(q_1) - g_1(q_1))$ using his q_1 option, plus $\alpha C(f_1(q_2) - g_1(q_2))$ using his q_2 option.

It is easy to verify that this option payment plan is equivalent to Table 1.

3.4 Multi Period Options

The option form of the 3-period contract can be easily extrapolated to an m-period contract ($m > 3$). Assume β is a positive number such that $\beta + \cdots + \beta^{m-2} < 1$. Such a β certainly exists. For example $0 \le \beta \le 1/2$ is enough for the condition to hold for all m. The contract now says:

1. (Period 1) There are various options that the user can buy, with option price $g_1(p)$ and resource price $f_1(p) - g_1(p)$, for all $p \in [0, 1]$. The user buys one share of q_1-option at price $g_1(q_1)$.

i. (Period i: $i = 2, \ldots, m - 1$) The user can swap β^{i-1} share of his q_1-option for a q_i-option, by paying the difference price $\beta^{i-1}(g_1(q_i) - g_1(q_1))$.

m. (Period m) If the user needs one unit of resource, he executes his options. That is, he pays

$$\left(1 - \frac{\beta - \beta^{m-1}}{1 - \beta}\right)(f_1(q_1) - g_1(q_1)) + \sum_{i=2}^{m-1} \beta^{i-1}(f_1(q_i) - g_1(q_i)). \tag{15}$$

4 Mechanism Behavior

We have seen that the truth-telling reservation mechanism helps the user to save money and the coordinator to make money, so they both have an incentive to use it. An interesting question to ask now is when $v \ge C$, whether the resource provider himself would want to use the reservation mechanism, playing both roles of seller and coordinator.

4.1 The User's Utility

Suppose the user has an expected utility in the form $u = vp - c \equiv v_1 - c$. Here, c is the minimum expected price he has to pay for one unit of resource, estimated in period 1. v_1 is the value of the unit to him in period 1, scaled to $v_1 \in [0, 1]$. If the user does not buy the resource when he needs it, his utility is zero. The user is risk-neutral.

4.2 The Seller's Problem

4.2.1 Direct Selling

Suppose the resource provider wants to maximize his revenue. When the cost is considerably less than the user's value this is equivalent to maximizing his profit. Say he believes in a prior distribution $f(v_1, p)$ of the users, where $f(v_1, p)\,dv_1\,dp$ is the fraction of users whose (v_1, p) lie in the small rectangle $(v_1, v_1 + dv_1) \times (p, p + dp)$. Without using the truth-telling reservation mechanism, he chooses the optimal reservation price C_1 and spot price C_2 that solve[3]:

$$\max_{0 \le C_1 \le C_2 \le 1} \iint dv_1\,dp\,f(v_1, p) I(v_1 \ge C_1 \wedge C_2 p) C_1 \wedge C_2 p. \tag{16}$$

It is easy to calculate that for the uniform prior $f(v_1, p) = 1$ the maximal revenue $R_{\max} = 5/24$ is achieved at $C_1 = 1/2$ and $C_2 = 1$.

4.2.2 Options

Within the truth-telling reservation framework, the seller sets two prices, $f(p)$ and $g(p)$, by choosing the parameters C_1, C_2 and k. Note that C_2 does not appear explicitly in the prices, but only appears implicitly in the constraint $k \le 2(C_2 - 1)$. Thus the seller can choose a sufficiently large C_2.[4] In the many-user limit, his optimization problem becomes

$$\max_{0 \le C_1 \le 1 \le k \le 2} \iint dv_1\,dp\,f(v_1, p) I(v_1 \ge w(C_1, k, p)) w(C_1, k, p), \tag{17}$$

where

$$w(C_1, k, p) = C_1 \left[\left(1 + \frac{k}{2}\right) p - \frac{k}{2} p^2 \right] \tag{18}$$

is the expected revenue he collects from a user whose expected value exceeds the expected cost.

For the uniform prior it can be shown that the seller's maximal revenue is again $R_{\max} = 5/24$, achieved at $C_1 = 5/8$. Hence in this case the option mechanism and direct-selling yield the same maximal revenue. While this is coincidental, as we shall see in the next section, it does suggest that the two revenues are comparable.

4.3 Other Distributions

We will now compare the two pricing schemes for other probability distributions. Again assume that v_1 and p are independent, and v_1 is uniform on $[0, 1]$. Assume now that p is uniformly distributed on $[a, b]$, where $0 \le a \le b \le 1$.

---

[3] As in standard probability texts, here $a \wedge b$ denotes the minimum of a and b, and $I(\cdot)$ is the indicator function.

[4] This may seem surprising, but remember that the user never pays the on-spot price when he buys an option! In fact, C_2 can be set greater than 1 in this case.

Table 2 The seller's revenue per person, using direct selling or options. For example, when the users' p is uniformly distributed over $(0, 1/2)$, the seller's revenue per person when using the option mechanism is 0.197

p	$(0, 1)$	$(0, 1/2)$	$(1/2, 1)$	$(0, 1/3)$	$(1/3, 2/3)$	$(2/3, 1)$	$(0, 1/5)$	$(2/5, 3/5)$	$(4/5, 1)$
Direct	0.208	0.167	0.250	0.130	0.245	0.250	0.087	0.248	0.250
Options	0.208	0.197	0.248	0.183	0.246	0.250	0.141	0.249	0.250

We optimize the seller's revenue for the two schemes with multiple choices of a and b. The numerical results are shown in Table 2. It can be seen that in most cases the option mechanism performs better than the direct mechanism. In particular when the users' probabilities are concentrated at the small end (row $(0, 1/2)$, $(0, 1/3)$ and $(0, 1/5)$ in the table), the option mechanism significantly beats direct selling. This is because in the direct selling scheme, the seller has to compromise for a low C_1 for small p, therefore losing considerable profit. On the other hand, by selling options he can settle on a much higher C_1 and profit from the premium.

We thus conclude that the truth-telling mechanism is particularly efficient for reservations of peak demands and rare events (small p).

5 Conclusion

In this paper we presented a solution to the truth revelation problem in reservations by designing option contracts with a pricing structure that induces users to reveal their true likelihoods that they will purchase a given resource. Truthful revelation helps the provider to plan ahead to reduce the risk of under-utilization and overbooking. In addition to its truthfulness and robustness, the scheme can extract similar revenue to that of a monopoly provider who has accurate information about the population's probability distribution and uses temporal discrimination pricing.

This mechanism can be applied to any resource that exhibits bursty usage, from IT provisioning and network bandwidth, to conference rooms and airline and hotel reservations, and solves an information asymmetry problem for the provider that has traditionally led to inefficient over or under provision.

This approach can be extended in a number of ways so as to become useful in a number of realistic situations. With the addition of a simulation tool developed in the context of swing options [3], for example, users can anticipate their future needs for resources at given times and price them accordingly before committing to a reservation contract. Yet another extension would allow for the reservation of single units of a resource (airline seats or conference rooms, for example) over a time interval, as opposed to a particular date.

Given the rather inefficient way through which most bursty resources are now allocated, we believe that this mechanism will contribute to a more useful and profitable way of allocating them to those who need them, while giving the provider essential information on future demand that he can then use to rationally plan its provisioning.

Acknowledgement We thank Andrew Byde for valuable suggestions.

Appendix

In this appendix we define $\delta = 1/C$ to simplify the expressions.

Lemma 3 *If a user submits in period* 1, *it is weakly better for him to adjust in period* 2.

Proof Consider a user who has arrived at period 2. He already submitted q_1 in period 1, and now can either adjust his probability to q_2, or do nothing. If he chooses to adjust, he will have to pay the adjustment fee in period 3, expected to be

$$p_2\alpha[f_2(q_2) - f_2(q_1)] + (1 - p_2)\alpha[g_2(q_2) - g_2(q_1)], \tag{19}$$

where $p_2 = p_{21}$ or p_{22} is his real probability of using the resource in period 3, which he now knows. It can be easily checked that, no matter what he submitted in period 1, it is always weakly better for his to adjust q_1 to p_2 (truth-telling). Letting $q_2 = p_2$ in (19), we have

$$\begin{aligned}
&p_2\alpha[f_2(p_2) - f_2(q_1)] + (1 - p_2)\alpha[g_2(p_2) - g_2(q_1)] \\
&\quad = \alpha\{[p_2 f_2(p_2) + (1 - p_2)g_2(p_2)] - [p_2 f_2(q_1) + (1 - p_2)g_2(q_1)]\} \le 0.
\end{aligned} \tag{20}$$

The last step follows from Condition T1. □

Lemma 4 (Period 1 truth-telling) *Suppose $\alpha < \delta$. If a user submits in period* 1, *he submits his real probability.*

Proof From Lemma 3 we know that the user will adjust his probability to p_2 in period 2. As a result his expected cost is

$$\begin{aligned}
c_{12} &= pf_1(q_1) + (1 - p)g_1(q_1) \\
&\quad + p_1 p_{21}\alpha[f_2(p_{21}) - f_2(q_1)] + p_1(1 - p_{21})\alpha[g_2(p_{21}) - g_2(q_1)] \\
&\quad + (1 - p_1)p_{22}\alpha[f_2(p_{22}) - f_2(q_1)] + (1 - p_1)(1 - p_{22})\alpha[g_2(p_{22}) - g_2(q_1)] \\
&= p[f_1(q_1) - \alpha f_2(q_1)] + (1 - p)[g_1(q_1) - \alpha g_2(q_1)] + \text{function}(p_1, p_{21}, p_{22}) \\
&= (1 - \alpha C)[pf_1(q_1) + (1 - p)g_1(q_1)] + \text{function}(p_1, p_{21}, p_{22}).
\end{aligned} \tag{21}$$

Here the notation c_{12} means the user submits both in period 1 and 2. By assumption $1 - \alpha C > 0$. Because (f_1, g_1) is truth-telling, the last equation is minimized when $q_1 = p$. Thus, if the user submits a likelihood, he better submit p, the true probability (estimated in period 1) that he will use one unit of resource in period 3. □

Note that, for the special choice $f_2 = Cf_1$, when the user submits two probabilities and uses the resource, his payment can be written as

$$f_1(q_1) - \alpha f_2(q_1) + \alpha f_2(q_2) = (1 - \alpha C)f_1(q_1) + \alpha C f_1(q_2). \tag{22}$$

Lemma 4 assumes that $\alpha C < 1$. Then the mechanism can be understood as having the user buy $(1-\alpha C)$ fraction of the q_1 contract to take advantage of the large discount, and buys αC fraction of the q_2 contract to take advantage of his increased level of information.

Lemma 5 *The user prefers to submit a rough estimation in period* 1 *and then adjust it in period* 2, *rather than to ignore period* 1 *and only submit in period* 2.

Proof We compare the user's cost in both cases. If he only submits in period 2, he would of course submit the true probability (in period 2). Thus his expected cost is

$$\begin{aligned} c_2 &= p_1 w_2(p_{21}) + (1-p_1) w_2(p_{22}) \\ &= p_1 C w_1(p_{21}) + (1-p_1) C w_1(p_{22}). \end{aligned} \tag{23}$$

If he submits in both periods, his expected payoff is

$$\begin{aligned} c_{12} &= (1-\alpha C)[p f_1(p) + (1-p) g_1(p)] + \alpha c_2 \\ &= (1-\alpha C) w_1(p) + \alpha c_2, \end{aligned} \tag{24}$$

where the first "=" is obtained by using the result of Lemma 4 to replace q_1 by p in (21).

We want to show that $c_{12} < c_2$, so that the user wants to submit twice. It can be found after some algebra that the condition is equivalent to having

$$\frac{p_1 w_1(p_{21}) + (1-p_1) w_1(p_{22})}{w_1(p)} > \frac{\delta - \alpha}{1-\alpha}. \tag{25}$$

If $p_{21} = p_{22}$ the left hand side is 1, so the inequality is satisfied. If $p_{21} \neq p_{22}$, then without loss of generosity we can assume that $p_{21} < p_{22}$, and (25) can be written as

$$\frac{p_{22} - p}{p_{22} - p_{21}} \frac{w_1(p_{21})}{w_1(p)} + \frac{p - p_{21}}{p_{22} - p_{21}} \frac{w_1(p_{22})}{w_1(p)} > \frac{\delta - \alpha}{1-\alpha}, \tag{26}$$

where $p = p_1 p_{21} + (1-p_1) p_{22} \in [p_{21}, p_{22}]$. Note that for fixed p, the left-hand side of (26) is increasing in p_{21} and decreasing in p_{22}, so we can let $p_{21} = 0$ and $p_{22} = 1$ to obtain the stronger condition

$$\frac{w_1(p)}{p} < \frac{1-\alpha}{\delta - \alpha}. \tag{27}$$

If (27) holds for all p, then $c_{12} < c_2$ for all (p_1, p_{21}, p_{22}).

At this stage we take into account the specific form of $w_1(p)$:

$$w_1(p) = \left(1 + \frac{k}{2}\right) p - \frac{k^2}{2} p^2. \tag{28}$$

We then have

$$\frac{w_1(p)}{p} = 1 + \frac{k}{2} - \frac{k^2}{2} p \le 1 + \frac{k}{2} \le C = \frac{1}{\delta} < \frac{1-\alpha}{\delta - \alpha}, \tag{29}$$

where the second "$\leq$" from the fact that $k \in [1, \min\{2(C-1), 2\}]$. Hence (27) indeed holds for all p, and $c_{12} < c_2$. □

Lemma 6 *The coordinator makes profit when* $0 < \alpha < \delta$.

Proof The coordinator expects to collect from the user

$$\begin{aligned} c_{12} &= (1-\alpha C)w_1(p) + \alpha C[p_1 w_1(p_{21}) + (1-p_1)w_1(p_{22})] \\ &\geq (1-\alpha C)p + \alpha C[p_1 p_{21} + (1-p_1)p_{22}] \\ &= (1-\alpha C)p + \alpha C p = p, \end{aligned} \tag{30}$$

where we have used the fact $w_1(p) \geq p$ for all $p \in [0, 1]$. Thus he expects to collect $\geq p$ from each user who has probability p, so he makes profit. □

References

1. Beran, J., Sherman, R., Taqqu, M.S., Willinger, W.: Long-range dependence in variable-bit-rate video traffic. IEEE Trans. Commun. **43**(2–4), 1566–1579 (1995)
2. Conlisk, J., Gerstner, E., Sobel, J.: Cyclic pricing by a durable goods monopolist. Q. J. Econ. **99**, 489–505 (1984)
3. Clearwater, S.H., Huberman, B.A.: Swing options: a mechanism for pricing IT peak demand, http://www.hpl.hp.com/research/idl/papers/swings/ (2005)
4. Durrett, R.: Probability: Theory and Examples, 3rd edn. Brooks/Cole (2005), Chap. 4
5. Gale, I.: Advance-purchase discounts and monopoly allocation of capacity. Am. Econ. Rev. **83**(1), 135–46 (1993)
6. Gribble, S.D., Manku, G.S., Roselli, D.S., Brewer, E.A., Gibson, T.J., Miller, E.L.: Self-similarity in file systems. Meas. Model. Comput. Syst. 141–150 (1998)
7. Huberman, B.A., Lukose, R.: Social dilemmas and Internet congestion. Science **277**, 535–537 (1997)
8. Sobel, J.: Durable goods monopoly with entry of new customers. Econometrica **59**, 1455–1485 (1991)
9. Varian, H.: Price discrimination. In: Schmalansee, R., Willing, R. (eds.) The Handbook of Industrial Organization, vol. 1, pp. 597–654. Elsevier Science/North-Holland, Amsterdam/New York (1989), Chap. 10
10. Voldman, J., Mandelbrot, B.B., Hoevel, L.W., Knight, J., Rosenfeld, P.: Fractal Nature of Software-Cache Interaction. IBM J. Res. Dev. **27**(6), 164–170 (1981)
11. Williams, D.: Probability with Martingales, pp. 72–73. Cambridge University Press, Cambridge (2001)
12. Wilson, R.B.: Nonlinear Pricing, pp. 377–379. Oxford University Press, New York (1993)