Pricing Cloud Computing: Inelasticity and Demand Discovery

June 7, 2013

Abstract

The recent growth of the cloud computing market has convinced many businesses and policy makers that cloud-based technologies will fundamentally change the way IT is built, bought and used. This paper studies how the unique features of cloud-based services alter the economics of IT pricing. In particular, we focus on two demand characteristics, namely (1) the presence of local demand inelasticity—a client's demand (i.e., quantity) for cloud-based services may not vary continuously with price; (2) the need to use pricing policy to discover the demand curve—new cloud-based services often face an unknown demand system. Our analysis indicates that these characteristics necessitate a reformulation of the traditional monopoly screening or nonlinear pricing problem. We present a first step towards this reformulation with the objective to provide practical guidance to cloud vendors facing such pricing challenges. We show that under a very weak ordering condition of customer types, the vendor's pricing problem is completely separable. Moreover, a vendor is able to collect more information about the demand system simply by proportionally increasing his entire myopic price menu.

1. Introduction

The recent growth of the cloud computing market has convinced many businesses and policy makers that cloud-based technologies will fundamentally change the way information technology (IT) is built, bought and used (Carr 2008). Cloud computing refers to a computing model in which dynamically scalable and often virtualized resources are provided as a service over the Internet. It includes the models of offering Infrastructure as a Service (IaaS), Platform as a Service (PaaS) as well as Software as a Service (SaaS). Despite the current economic downturn, the global cloud computing market is forecast to grow from about $41 billion in 2011 to $241 billion in 2020 (Forrester 2011). The Obama Administration announced the cloud computing initiative in 2009. As the world's largest IT consumer, the Federal government plans to migrate $20 billion of IT spending into the cloud (Federal Cloud Computing Strategy 2011). The optimal pricing of cloud computing is thus of substantial business importance, but also presents a number of new challenges to technology firms making the transition from the traditional provision of IT to supplying their software and services from the cloud.

There is a fairly extensive literature on how the unique characteristics of information technology alter the economics of IT pricing (for example, Bakos and Brynjolfsson 1999, Sundararajan 2004). Our paper focuses on the business-to-business IT market and expands prior literature by studying a set of new economic characteristics that distinguish the pricing of cloud computing from other IT-based (and non-IT) products and services, namely (1) the presence of local demand inelasticity; (2) the need to use pricing policy to discover the demand system.

Local demand inelasticity: A client firm cannot easily vary its cloud usage levels in response to a unit price change. A financial institution seeking to rent computing capacity to run a risk assessment program, for instance, may not be able to change the amount of computation that the program requires continuously with the cost of computing services. A firm that needs a cloud-based human resource (HR) application usually buys licenses for its entire HR team rather than varying the number of user licenses that it subscribes to according to price levels of the SaaS application. Thus, individual customer's utility from using a cloud service is less likely to vary smoothly with the quantity of serivces consumed. Instead, a customer is more likely to make a binary purchase decision: either buy the cloud service at the required quantity level, or do not buy at all, a phenomenon we label local demand inelasticity. This violates a critical assumption of prior literature that studies quantity-based pricing of software services. In this paper, we demonstrate that this local inelasticity has a *critical* impact on both how the pricing problem is modeled as well as on key efficiency results.

Demand discovery: A distinguishing characteristics of cloud-based services is that unlike the established corporate computing market, it may not be possible initially to describe the demand system that a cloud vendor faces. Thus, in the early stages of pricing, the choice of pricing policy is influenced by an objective beyond simple profit maximization: it might also be driven by what kind of pricing best informs the vendor about their demand system. Put simply, it may be optimal to explore early on towards more informed exploitation later.

Our analysis indicates that these characteristics necessitate a fundamental reformulation of the traditional monopoly screening or nonlinear pricing problem. We provide a first step towards this reformulation in a manner that is analytically tractable and based on parameters that are more easily measurable and likely to make empirical demand estimation and the real-world use of the analytical model more viable.

Our paper makes two main contributions. First, we formulate and characterize the solution to a complicated pricing problem with discontinuous and inelastic individual demand functions, with virtually no restrictions on the distribution of customer types, and no single-crossing restrictions on the utility functions. We show that under a very weak ordering condition of customer types, the monopolist's pricing problem is completely separable–the monopolist can treat customers with different demand quantities as separate markets, and his optimal price schedule composes of revenue-maximizing prices for each market.

Managerially, it describes a solution that requires relatively little information about customers – an assesment of anticipated consumption levels and the fraction of customers who place different value levels on each of these consumption levels. In contrast, typical screening problems require one to specify a complete utility function for a continuum of customer types (see, for example, Maskin and Riley 1984). Therefore, our findings extend the pricing theory in a direction that admits more practical applications.

Second, we extend the baseline model and study a monopolist's pricing problem when the demand is initially unknown. We adopt Blackwell's (1953) notion of informativeness and demonstrate that a monopolist is able to collect more information about the demand simply by proportionally increasing his entire price menu. Any other price menues, in general, do not guarantee "more information", or a higher expected future payoff for the monopolist than a myopic price menu. The vendor's optimal pricing policy depends on the tradeoff between his gain in expected future payoff through improved knowledge about the demand system and his loss from the current period payoff.

2. Model

Consider an infinite period model in which a monopoly vendor sells a cloud-based service in each period that may be used by customers in varying quantities, $q \geqslant 0$. At the beginning of each period t, one customer enters the market. The monopolist announces a price schedule $p_t(q)$ that specifies the price for supplying quantity q of services to customers in period t. The customer observes the prices and makes a purchase decision. Trade takes place. The customer exits the market at the end of the period.

Customers are heterogeneous and indexed by a type parameter $\theta \in \Theta$. A type-θ customer demands a fixed quantity q_θ of services and is willing to pay up to v_θ for this quantity. Thus, the utility that a customer of type θ derives from using q units of services is defined as:

Assumption 1:

$$U(q;\theta) = \begin{cases} 0 & \text{for } q < q_\theta \\ v_\theta & \text{for } q \geqslant q_\theta \end{cases} . \tag{2.1}$$

A customer does not treat a consumption level greater than q_θ as inferior to a consumption level of q_θ if she does not need to pay more for the additional units of services. Neither does she treat it as superior if it costs at least as much as the consumption level q_θ. Figure 2.1 illustrates an example of a customer's utility function. For simplicity, we assume that the seller's variable costs are constant at zero.

We first assume that the monopolist does not observe each customer's type, but knows the

Figure 2.1: Customers' utility function

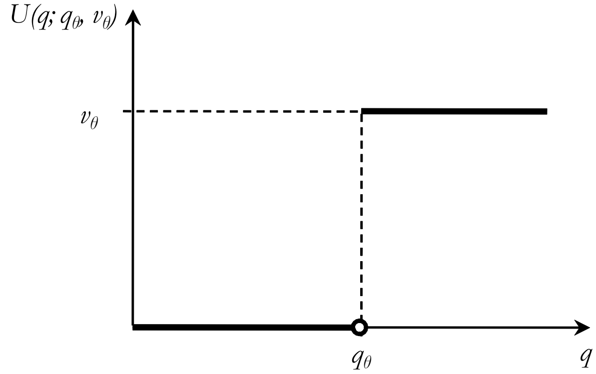

distribution of types within the population from which each customer is randomly drawn (Section 2.1). Denote by $F\left(q_\theta, v_\theta\right)$ the cumulative density function of this distribution. (q_θ, v_θ) is distributed in the rectangle $[q_L, q_H] \times [0, \overline{v}]$.

Assumption 2. $F(q_\theta, v_\theta)$ is absolutely continuous and differentiable in both arguments.

Denote by $f\left(q_\theta, v_\theta\right)$ the probability density function of the demand distribution. When the demand distribution is known, the monopolist's pricing problem in each period is independent, and a pricing policy that maximizes his one-period expected payoff also optimizes his overall payoff if repeated in the infinite-period game. Thus, we investigate the monopolist's optimcal price schedule in a single period. We then relax this assumption and investigate the monopolist's dynamic pricing problem when the demand distribution is initially unknown but can be learned over time (Section 2.2).

2.1. Pricing with Inelastic demand

Consider customers' decision: A customer seeks to maximize her surplus:

$$CS^*(\theta) = \max_{q \geqslant 0} U(q; q_\theta, v_\theta) - p\left(q\right). \tag{2.2}$$

Since a customer does not treat a consumption level greater than q_θ as inferior to a consumption level of q_θ, the solution to her optimization problem is simply to buy at quantity $\widehat{q}$, if $\widehat{q} \geqslant q_\theta$ and $p(\widehat{q}) \leqslant v_\theta$ and $\nexists q \geqslant q_\theta$ such that $p(q) < p(\widehat{q})$. That is if a customer decides to buy the service, she would buy at the quantity level that covers her efficient consumption level at the lowest price. One can show that a sufficient condition for a customer to buy at her own efficient consumption level if she decides to purchase ($\widehat{q} = q_\theta$) is that $p(q)$ is nondecreasing in q.

The monopolist's decision is to determine a price schedule $p(q)$ which maximizes his profit. Notice that customers with the same efficient consumption level may not derive the same utility from using the cloud service. For instance, the same computing capacity may generate different values for a financial institution than for a retailer. Therefore, given any type distribution $F(q_\theta, v_\theta)$, one can define $g(v;q)$ as the conditional probability that a consumer has $v_\theta = v$ given her $q_\theta = q$. Formally, this transformation from the standard formulation to ours is simply:

$$g(v;q) \underset{def}{=} \frac{f(q,v)}{\displaystyle\int_0^{\overline{v}} f(q,x)dx}.$$

And define

$$G(v;q) = \int_0^v g(x;q)dx$$

as the fraction of customers with efficient consumption level q whose willingness to pay for that level of consumption is at most v.

Given any nondecreasing pricing function, $p(q)$, one can show that the monopolist's objective function can be written, using our new formulation, as

$$\max_{p} \int_{q_L}^{q_H} p\left(q\right)\left(1 - G\left(p\left(q\right);q\right)\right)\left(\int_{0}^{\overline{v}} f(q,x)dx\right)dq. \tag{2.3}$$

Now define $\overline{p}(q)$ as the revenue-maximizing price if the monopolist trades with *just the set of customers* whose efficient level of consumption is q, or mathematically

$$\overline{p}(q) = \arg\max_{p} p\left(1 - G(p;q)\right). \tag{2.4}$$

One can show that if $g\left(p;q\right)$ is everywhere positive, and for any $q_1 < q_2$ and $q_1, q_2 \in [q_L, q_H]$, we have $G\left(v;q_2\right)$ hazard rate dominates (HRD) $G\left(v;q_1\right)$, then the solution to the pointwise optimization problem for each q, as in (2.4), constitutes a nondecreasing function of q, and hence, is the solution to the global optimization problem specified in (2.3). The following Theorem formalizes this result. All proofs not included in the paper are available upon request.

Theorem 1. *For each $q_1 < q_2$, where $q_1, q_2 \in [q_L, q_H]$, if $G(v;q_2)$ hazard rate dominates $G(v;q_1)$, or mathematically*

$$\frac{g(v;q_2)}{1-G(v;q_2)} \leqslant \frac{g(v;q_1)}{1-G(v;q_1)}$$

then there is no pricing function that yields a profit level higher than $\overline{p}\left(q\right)$.

This theorem characterizes the solution to a complicated pricing problem with discontinuous and inelastic individual demand functions, with weak restrictions on the distribution of customer types, and no single-crossing restrictions on the utility functions. Our results imply that under a fairly weak ordering restriction on g, the seller's pricing problem is completely separable: A cloud vendor can segment the market based on customers' anticipated consumption levels, and the optimal price schedule consists of revenue-maximizing prices for each segment of the market. This finding describes a solution that can be readily implemented by cloud vendors.

While the ordering condition of Theorem 1 is quite natural, the following extension to this theorem shows that even when this ordering condition does not hold, the optimal pricing policy comprises of prices solely from the range of $\overline{p}(q)$. We need a few additional assumptions. Let $I\left(p\left(q\right)\right) = p\left(q\right)\left(1 - G\left(p\left(q\right); q\right)\right)$.

Assumption 3. $I\left(p\left(q\right); q\right)$ is strictly quasi-concave.

Assumption 4.

$$\begin{aligned} \int_x^y \frac{\partial I}{\partial p}\left(p\left(q\right); q\right)\left(\int_{\{\theta: q_\theta = q\}} dF(q_\theta, v_\theta)\right) dq &\geqslant 0 \longrightarrow \\ \int_x^y \frac{\partial^2 I}{\partial p^2}\left(p\left(q\right); q\right)\left(\int_{\{\theta: q_\theta = q\}} dF(q_\theta, v_\theta)\right) dq &< 0,\ x < y. \end{aligned}$$

The following theorem generalizes the results in Theorem 1.

Theorem 2. *If there exists an optimal response function $p^*(q)$, then there exists a set of subintervals*

$$\left\{\left[x^i, y^i\right] \subseteq \left[q_L, q_H\right] | x^i > y^{i-1}\right\}_{i \in I},$$

possibly empty, such that for all i

(i) $\overline{p}(x^i) = \overline{p}(y^i)$

Figure 2.2: Optimal Pricing without the HRD Condition

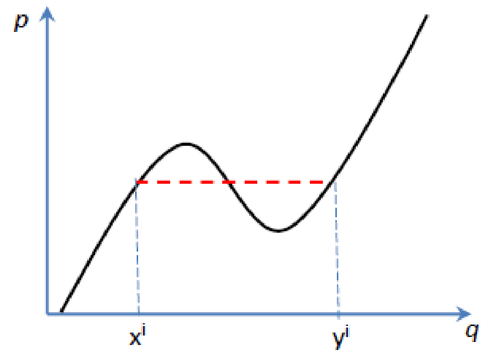

(ii) $\int_x^y \frac{\partial I(\overline{p}(q);q)}{\partial p} \left(\int_{\{\theta:q_\theta=q\}} dF(q_\theta, v_\theta) \right) dq \leqslant 0$, *where* $x < y$, $y \notin \cup [x^i, y^i]_{i \in I}$, *and equality holds only when* $x = x^i$ *and* $y = y^i$.

$$\text{(iii) } p^*(q) = \begin{cases} \overline{p}(y^i) & \text{if } q \in [x^i, y^i] \text{ for some } i \\ \overline{p}(q) & \text{otherwise,} \end{cases}$$

Given assumption 2 and 3, and that the collection of subintervals satisfying (i) and (ii) exist, the price schedule defined in (iii) is optimal if assumption 4 also holds.

Intuitively, if the HRD condition does not hold, price menu $\overline{p}(q)$ is no longer nondecreasing in q. An example is shown in the following Figure 2.2. Given assumption 2 and 3, one can define a set of intervals that satisfy (i) and (ii), such that the optimal pricing is $\overline{p}(q)$ outside these intervals, and is equal to $\overline{p}(y^i)$ for any quantity $q \in [x^i, y^i]$ for some $i \in I$.

2.2. Dynamic pricing with demand discovery

This section considers the dynamic game in which the monopolist does not known the demand system $F(q, v)$ initially, but can learn over time. To highlight the impact of learning, we compare the optimal pricing policy of a myopic monopolist, whose objective is to maximize his current period payoff given his demand information, with that of an experimenting monopolist, whose objective is to maximize his overall payoff given the same information.

For simplicity, let us consider a case in which there are N efficient quantity levels $\{q_1, ..., q_N\}$, where $0 < q_1 < ... < q_N$. The willingness-to-pay of customers with the same efficient consumption level follows a uniform distribution. That is $g(v; q_i) = 1/v_{q_i}$, and $G(v; q_i) = v/v_{q_i}$, where v_{q_i} denotes this market segment's consumers' highest willingness-to-pay for using quantity q_i of services, $i = 1, ..., N$. Let v_q denote the vector $(v_{q_1}, ...v_{q_N})$. It is conceivable that the monopolist knows the measure of customers with the same efficient consumption level, or

$$w_i = \int_{\{\theta : q = q_i\}} f(q_i, x)\, dx,$$

although the value of v_q is unknown to him. This may be a reasonable assumption because it is relatively easy for a vendor to obtain, for instance, data on the distribution of clients' firm size, and hence, estimate the fraction of customers requiring a certain number of user licenses. Nonetheless, clients' willingness-to-pay for the service is much harder to obtain.

The monopolist does not know v_q, but knows $v_q \in \Phi = \{v^1, ..., v^M\}$. At the beginning of each period t, the monopolist holds a belief about the value of v_q, represented by the prior μ_t, where $\mu_t \in P(\Phi)$, and $P(\Phi)$ denotes the set of probability measures on Φ. He announces a price schedule $p_t(q) \in X$, where X is a compact subset of R^N, and observes a market outcome. There are $N + 1$ possible market outcomes: the customer purchases the service at one of the quantity levels or does not purchase the service. Denote the market outcome by $y_t = (y_{t,0}, y_{t,1}, ..., y_{t,N})$, where for $i = 1, ..., N$, $y_{t,i} = 1$ if the customer buys the service at quantity q_i, and zero otherwise, and

$$y_{t,0} = 1 - \sum_{i=1}^{N} y_{t,i}.$$

Let Y be the set of all possible market outcomes. Given any v_q, denote by $h(y|p, v_q)$ the conditional probability of a market outcome y given a nondecreasing price schedule $p(q)$. That is

$$h(y|p, v_q) = \prod_{i=1}^{N} (w_i (1 - G(p; q_i)))^{y_i} \times \left(\sum_{i=1}^{N} w_i G(p; q_i) \right)^{1-\sum_{i=1}^{N} y_i}.$$

Denote by $r(p_t, y_t)$ the seller's payoff in period t given price schedule $p_t(q)$ and outcome y_t.

After observing the market outcome, the monopolist updates his belief on the value of v_q following Bayes' rule, since the observation y_t contains information on the value of v_q. In particular, his revised belief, $\mu_{t+1}(v_q)$, depends on the conditional probability of the parameter given his prior μ_t, the price schedule $p_t(q)$ and the market outcome y_t. That is for each $v^m \in \Phi$,

$$\mu_{t+1}(v^m; p_t, y_t) = \begin{cases} \frac{\mu_t(v^m) \times w_i \left(1 - p_i / v^m_{q_i}\right)}{\sum_{j=1}^{M} w_i \left(1 - p_i / v^j_{q_i}\right) \mu_t(v^j)} & \text{if } y_{t,i} = 1 \\ \frac{\mu_t(v^m) \times \left(\sum_{i=1}^{N} w_i \times p_i / v^m_{q_i}\right)}{\sum_{j=1}^{M} \left(\sum_{i=1}^{N} w_i \times p_i / v^j_{q_i}\right) \mu_t(v^j)} & \text{if } y_{t,0} = 1 \end{cases}.$$

The monopolist's problem is to determine a price schedule $p_t(q)$ in period t that maximizes his expected dicounted sum of profits given the history of the game. Note, the history of the game impact the monopolist's continuation profit only through providing information on the demand distribution, and this information is summarized by his current belief on the demand distribution. One can show that the monopolist's problem can be formulated into a dynamic programing problem in which the state variable in each period is the monopolist's current belief on the probability distribution of v_q, i.e., $\mu_t(v_q)$. The monopolist chooses his price schedule, $p_t(q)$, to maximize the expected discounted sum of profits. In doing so, he recognizes that $p_t(q)$ influences both his current period profit and through learning, his future payoff. In particular,

$$V(\mu_t) = \max_{p_t} \pi_t(p_t, \mu_t) + \delta K(p_t, \mu_t), \tag{2.5}$$

where

$$\pi_t(p_t, \mu_t) = \sum_{y \in Y} \left[r(p_t, y_t) \cdot \left(\sum_{j=1}^{M} h\left(y_t | p_t, v^j\right) \mu_t\left(v^j\right) \right) \right]$$

represents the monopolist's current period expected payoff, and

$$K(p_t, \mu_t) = \sum_{y_t \in Y} \left[V\left(\mu_{t+1}(v; p_t, y_t)\right) \cdot \left(\sum_{j=1}^{M} h\left(y_t | p_t, v^j\right) \mu_t\left(v^j\right) \right) \right]$$

represents his expected future payoff.

On the other hand, a myopic monopolist which does not consider the future benefits of learning would simply choose $p_t(q)$ such that

$$V(\mu_t) = \max_{p_t} \pi_t(p_t, \mu_t).$$

Denote the solution to this problem by p^O. Therefore, the value of information generated by a price schedule $p(q)$ can be measured by the improvement in expected future profits because of the information obtained from the current period. Specifically, *Definition.* Price schedule p_t is *more informative* than price schedule p_t' if

$$\sum_{y_t \in Y} \left[\Pr(y_t; p_t, \mu_t) \cdot \left(\sum_{j=1}^{M} \mu_{t+1}\left(v^j | p_t, y_t\right) u\left(p, v^j\right) \right) \right]$$
$$\geqslant \sum_{y_t \in Y} \left[\Pr(y_t; p_t', \mu_t) \cdot \left(\sum_{j=1}^{M} \mu_{t+1}\left(v^j | p_t', y_t\right) u\left(p, v^j\right) \right) \right]$$

for any real valued payoff function $u(p, v)$, and set of prices A.

We need to define another ordering condition before moving to our findings:

Definition. Price schedule $p = (p(q_1), ..., p(q_N))$ is *sufficient* for price schedule $p' = (p'(q_1), ..., p'(q_N))$ if there exist a $(N+1) \times (N+1)$ nonnegative matrix $X = (x_{ij})_{i,j=0,\dots,N}$, such that x_{ij} are independent of v_q, and $\sum_i x_{ij} = 1$, for all j, and

$$h(y_m | p', v) = \sum_j x_{mj} h(y_j | p, v).$$

The optimization problem defined in (2.5) is well known to be complex, and often does not lend itself to analytical solutions. Even numerical estimation is less likely to be solved in polynomial time than NP-complete problems (Papadimitriou and Tsitsiklis 1987).One objective of this paper is to provide cloud vendors with some guidelines on how price experimentation can be done in such a market. The following two theorems show that a monopolist is able to obtain more (valuable) information about the demand distribution by simply proportionally increasing his price schedule.

Assumption 5. $\min_j \left\{v_i^j\right\} > \max_j \left\{v_i^j\right\} \times \frac{1}{2}$.

Theorem 3. *Let* $p = (p(q_1), ..., p(q_N))$, $p' = (p'(q_1), ..., p'(q_N))$, *and* $p(q_i), p'(q_i) \leqslant \min_j \left\{v_{q_i}^j\right\}$.

If there exists $\rho \in (0,1)$ *such that* $p^{'} = \rho \cdot p$*, then* p *is sufficient for* p'.

We apply Blackwell's (1953) well-known result that sufficiency implies informativeness.

Theorem 4. *If* p *is sufficient for* p'*, then* p *is more informative than* p'*, or* $K(p,\mu) \geqslant K(p',\mu)$.

The above results hold independent of the value of $\{v^m | m = 1, ..., M\}$ and the monopolist's current belief μ_t, as far as the price schedules fall in $B = \left\{ p | 0 \leqslant p(q_i) \leqslant \min_j \left\{ v^j_{q_i} \right\} \right\}$. Moreover, we show below that, in general, other price schedules that are not parallel to price schedule p (or, equivalently, are not proportional to p) are not guaranteed to improve the monopolist's expected future payoff through demand information gained from the current period. This finding is formalized in the following theorem.

Theorem 5. *Let* $p = (p(q_1), ..., p(q_N))$, $p' = (p'(q_1), ..., p'(q_N))$*, and* $0 < p(q_i), p'(q_i) \leqslant \min_j \left\{ v^j_{q_i} \right\}$*. Define the following vectors,*

$$\overrightarrow{v}^1 = \begin{pmatrix} \frac{1}{v^1_{q_1}} \\ ... \\ \frac{1}{v^M_{q_1}} \end{pmatrix}, \overrightarrow{v}^2 = \begin{pmatrix} \frac{1}{v^1_{q_2}} \\ ... \\ \frac{1}{v^M_{q_2}} \end{pmatrix}, ..., \overrightarrow{v}^N = \begin{pmatrix} \frac{1}{v^1_{q_N}} \\ ... \\ \frac{1}{v^M_{q_N}} \end{pmatrix}, \text{ and } \overrightarrow{v}^0 = \begin{pmatrix} 1 \\ ... \\ 1 \end{pmatrix}.$$

If $\overrightarrow{v}^0, \overrightarrow{v}^1, ..., \overrightarrow{v}^N$ *are linearly independent, then if* p *is sufficient for* p'*, there exists* $\rho \in (0,1)$ *such that* $p' = \rho p$.

The linear independence condition in this theorem can be easily satisfied. For instance, when there are two efficient quantity levels, or $N = 2$, this condition requires that, in a two-dimensional space, $\left(1/v^m_{q_1}, 1/v^m_{q_2}\right)_{m=1,...,M}$ do not *all* fall on the same straight line.

One implication of the above three theorems is that we can partition the set of all nondecreasing price schedules in B, denoted by S, into subsets defined as follows: given any $\overline{p} \in S$, let $\Gamma(\overline{p}) = \{p | p = \rho\overline{p}$, where $\rho > 0$ and $p \in S\}$. Clearly $S = \underset{p \in S}{\cup} \Gamma(p)$, and if $\widehat{p} \in \Gamma(\overline{p})$, then $\Gamma(\widehat{p}) = \Gamma(\overline{p})$. Within each partition, we have

$$\text{If } p' = \rho_1\overline{p},\ p = \rho_2\overline{p} \text{ and } \rho_1 < \rho_2, \text{ then } p \text{ is sufficient for } p', \text{ and } K(p,\mu) \geqslant K(p',\mu).$$

Moreover, if $\widehat{p} \notin \Gamma(\overline{p})$, then $\overline{p}$ is not sufficient for $\widehat{p}$, and the reverse is not true either. Since sufficiency is equivalent to informativeness by notion of Blackwell (1953), this result implies that a price schedule not in the subset $\Gamma\left(p^O\right)$ may not produce more valuable information about the demand system than the myopic price menu p^O. A forward-looking monopolist, therefore, may improve his expected future payoff by proportionally increasing his entire myopic price menu p^O. This result provides a simple guideline to cloud vendors who wish to learn about the demand system that they are facing.

3. Concluding Remarks

This paper studies a cloud service vendor's pricing problem when customers' demand features: (1) the presence of local demand inelasticity; (2) the need to use pricing policy to discover the demand system. Our analysis indicates that these characteristics necessitate a reformulation of the traditional nonlinear pricing problem. We present a first step towards this reformulation with the objective to provide practical guidance to cloud vendors facing such pricing challenges. We show that under a very weak ordering condition of customer types, the vendor's pricing problem is completely separable. A vendor's optimal pricing policy consists of revenue-maximizing prices for each market segement. Moreover, a vendor is able to collect more information about the demand system simply by proportionally increasing his entire myopic price menu. More theoretical and practical implications and future directions are discussed in the complete version of the paper.

References

Blackwell, D., 1953. Equivalent comparisons of experiments. *The Annals of Mathematical Statistics*, 24(2), 265-272.

Bakos, Y. and Brynjolfsson, E. 1999. Bundling information goods: Pricing, profits, and efficiency. *Management Science*, 45(12), 1613-1630.

Carr, N., 2009. The big switch: Rewiring the world, from Edison to Google. WW Norton.

Papadimitriou, C. H., and Tsitsiklis, J. N., 1987. The complexity of Markov decision processes. *Mathematics of operations research*, 12(3), 441-450.

Sundararajan, A., 2004. Nonlinear pricing of information goods. *Management Science*, **50**, 1660-1673.