

Price Elasticity of the Enterprise Computing Resource Market

Qingye Jiang, *Senior Member, IEEE*, Young Choon Lee, and Albert Y. Zomaya, *Fellow, IEEE*

Abstract—Pricing is an important property of a product in that it significantly affects customer's purchasing behavior. In this paper, we study the supply and demand relationship of the global enterprise computing resource market, including the server sales business, server rental business, and public clouds. We find out that consumers in different market segments are not equally sensitive to prices changes, leading to significantly different purchasing behavior. The price elasticity is extremely inelastic for server sales and server rental business, and is modestly elastic for public clouds. In the server sales business, we reveal that (a) the three biggest vendors - IBM, HP and DELL - dominate and have the market power over pricing in their respective fine-grained market segment, and (b) price reduction is not an effective way to win market share when the price elasticity is inelastic.

Index Terms—cloud computing, enterprise computing resource market, microeconomics, price elasticity of demand

1 INTRODUCTION

Pricing strategies are crucial business decisions in the life cycle of a product, especially when there are multiple competitors in the market. In a perfect competition market, there are multiple suppliers offering similar products, while no participant is powerful enough to control the price of the product. To compete for market share, suppliers often use price reduction in promotional campaigns. In order to maintain competitiveness in the market, competitors are forced to lower their prices to match, resulting in a series of pricing wars. Many of such “matches” in price reduction are the outcome of business intuition rather than rigorous reasoning, and often at the cost of losing revenue and profit. It is not uncommon to see enterprises go out of business as the result of pricing wars.

The enterprise computing resource market seems to be a perfect competition market with multiple players. Traditionally, enterprise can buy servers from hardware vendors, or rent servers from hosting service providers. In recent years, Infrastructure-as-a-Service (IaaS) is gradually making its way into the enterprise computing resource market. Amazon Web Services (AWS) are widely recognized as examples of public cloud services. Inspired by the dramatic growth of AWS, there have quickly emerged many followers in the market, including Microsoft Windows Azure and Google Compute Engine. New service providers often employ price reduction as a strategy to win customers from competitors. Existing service providers are forced to carry out further price reductions to maintain competitiveness in the market. In this study, we are interested in the following three questions:

- Is the enterprise computing resource market a perfect competition market?
- In the enterprise computing resource market, how sensitive consumers are to price changes?

- Is price reduction an effective way to win in the enterprise computing resource market?

To answer these questions, we study the supply and demand relationship of the global enterprise computing resource market between 2006 and 2013. The data source includes Gartner's quarterly reports on worldwide server shipments, as well as the quarterly / annual reports from Amazon and Rackspace. (The references to all of the actual data sources being used in this paper are listed in <http://www.qyjohn.net/?p=3507>.) We quantitatively measure the consumer's sensitivity to price changes using the concept of price elasticity of demand [1]. We analyze the impact of different pricing strategies on the performance of the business, using the server sales business of IBM, HP and DELL as case studies. This research provides new insights into the microeconomics characteristics of the enterprise computing resource market, and can help participants in various market segments design better pricing strategies.

2 BACKGROUND AND RELATED WORK

2.1 The Enterprise Computing Resource Market

Enterprises have three major options to acquire computing resources. They can either buy new servers from hardware vendors, lease servers from hosting service providers, or use public clouds. The different ways to acquire computing resources represent different segments of the enterprise computing resource market, namely server sales business, server rental business, and public clouds (Table 1).

Server Sales Business – The major suppliers in the server sales business are hardware vendors, including IBM, HP, and DELL. Consumers can order servers directly from the vendors, or through a third-party distributor. The cost to acquire new servers is denoted as capital expense (CapEx), and is usually paid to the

TABLE 1: Different Segments in the Enterprise Computing Resource Market

Market segment	Server Sales Business	Server Rental Business	Public Clouds
Major Vendors	IBM, HP, DELL, Fujitsu, Oracle	Rackspace, OVH, SoftLayer	Amazon Web Services, Google Compute Engine, Microsoft Windows Azure
Acquisition Time	weeks	days	minutes
Expense Type	CapEx	OpEx	OpEx
Unit Cost	2,000–50,000 USD one-time payment	100–2000 USD monthly charge	0.02–6.82 USD hourly charge
Ownership	consumer	supplier	supplier
Operations	consumer	supplier	supplier
Life Cycle	3 to 5 years	1 to 36 months	hourly
Contracts	optional service contract with service level agreement	long term contract with service level agreement	service level agreement
Market Size (#servers)	10 million annually or 40 million in total	2 million in total	1 million in total
Market Share	93%	5%	2%

supplier in a single payment before or after hardware delivery. The acquisition time usually falls within weeks for low-end servers up to months for high-end servers. When the servers are delivered to the consumer, their ownership is also transferred to the consumer. The life cycle of servers is between three and five years.

Server Rental Business – The major suppliers in the server rental business are hosting service providers, including Rackspace, OVH, and SoftLayer. Consumers usually establish a server rental contract directly with the service provider. According to the contract, the consumer pays a recurring monthly fee to maintain access to designated computing resource through network. Such cost is denoted as operations expense (OpEx). The time needed to acquire access to computing resource after signing the contract usually falls within days. The server rental contract does not change the ownership of computing resource, but some service providers choose to give the ownership to the consumer after the successful execution of a long contract (e.g., 3 years).

Public Clouds – Major cloud providers are Amazon Web Services, Google Compute Engine, and Microsoft Windows Azure. Consumers can directly request for, and gain instant access (within minutes) to, computing resource through web portals or application programming interfaces (API). In a manner similar to public utilities such as gas and electricity, consumers only pay for the amount of computing resource they use for the duration of their usage. The ownership of the physical computing resource always belongs to the service providers. Consumers’ responsibility is essentially confined to the pay-as-you-go expense.

2.2 Related Work

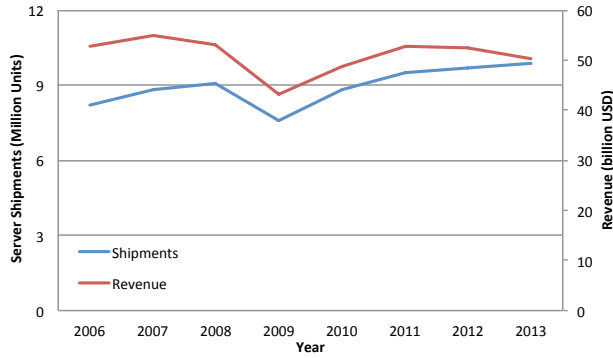
There exists an abundance of literature on the pricing strategy in the computing resource market. Most of them focus on maximizing revenue / profit for a particular vendor or service provider within its own customer base. These researches usually assume a certain resource consumption pattern from the customers,

without considering the impact from its competitors and the overall business environment. Only very few researchers studied the supply-demand relationship in the computing resource market, as presented below.

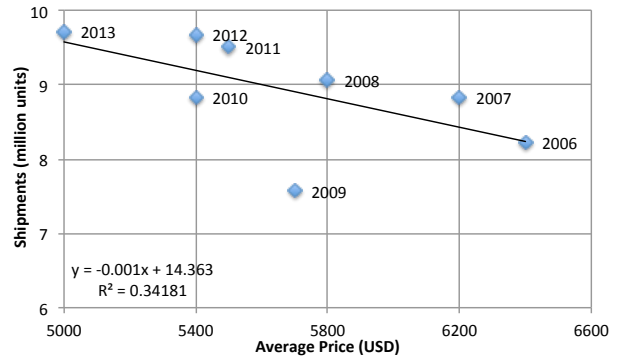
Chow [2] studies the technological change and the demand for computers with the price and quantity of computers in the United States between 1955 and 1965. Tam et al. [3] extend Chow’s work with annual computer hardware spending data in the United States between 1955 and 1984. Their findings indicate that the price elasticity of demand for computer hardware is inelastic, but the degree of elasticity changes over time. The data sources being used in [2] and [3] are extremely old, while technology advancements have dramatically changed the horizon of the computing resource market in general. On the technology side, expensive mainframes are being replaced by commodity x86_64 hardware. On the business side, server rental and public clouds are gaining increasing level of popularity among enterprises. Such changes in the demand for computing resource, as well as the relationship between supply and demand, are largely under-explored.

Stavins [4] studies the price elasticity of demand for the personal computer market with annual personal computer sales data in the United States between 1976 and 1988. Goolsbee [5] compares the price sensitivity on computer purchases from online stores versus in retail stores. Both [4] and [5] focus on personal computers instead of enterprise computing resources.

Rogers et al. [6] use Game Theory to analyze the effect of price reduction on the public cloud market. The authors concluded that there is no Nash equilibrium in this game, resulting in a constant stalemate. Cloud service providers practice price reductions to gain media attention and up-sell value-added services, at the cost of reduced margin from commoditized services such as storage and compute. The problem with this approach is that unless the profit from the value-added services exceeds the lost from down pricing, lowering price always results in reduced profit and can threaten the survival of a business.



(a) The volume of server shipments and revenues on a global scale, based on the data in Gartner's quarterly reports on worldwide server shipments.



(b) The demand curve of the global server sales business, calculated from the data in Gartner's quarterly reports on worldwide server shipments.

Fig. 1: Server Sales Business

3 PRICE ELASTICITY OF DEMAND

In this section, we first introduce the microeconomics concept of price elasticity of demand; describe its application to the enterprise computing resource market. Then, we present our analysis on price elasticity of the three market segments.

Price elasticity of demand is defined by the percentage change in demand over the percentage change in price. More formally,

$$E = \frac{\left(\frac{dQ}{Q}\right)}{\left(\frac{dP}{P}\right)}, \quad (1)$$

where E represents the price elasticity of demand, P represents the price and Q represents the quantity sold, while dP and dQ represent the percentage changes (in absolute values) in price and demand. When E is greater than 1, the percentage change in demand is greater than the percentage change in price, which is classified as elastic demand. When E is smaller than 1, the percentage change in demand is smaller than the percentage change in price, which is classified as inelastic demand. When E equals to 1, the percentage change in demand equals to the percentage change in price, which is classified as unitary demand. When demand is inelastic, a rise in price leads to a rise in revenue. When demand is elastic, a fall in price leads to a rise in revenue.

In general, there are two factors contributing to the price elasticity of demand of a product. The first factor is the availability of close substitutes. If a buyer has greater choice among close substitutes in consumption, the price elasticity of demand is greater. The second factor is the proportion of a buyer's budget that is devoted to a product. The larger the proportion of budget, the more responsive is the quantity demanded to price changes, and the price elasticity of demand is greater.

To identify the price elasticity of demand in this study, we carry out linear regression with history

price and demand data. With such a technique, we can express the relationship between price and demand as

$$Q = b_0 + b_1P. \quad (2)$$

The slope of a linear curve b_1 can be expressed as

$$b_1 = \frac{dQ}{dP}. \quad (3)$$

Therefore


$$E = -b_1 \frac{P}{Q}. \quad (4)$$

In the enterprise computing resource market, vendors and service providers offer computation units (such as servers and EC2 instances) with different capacities. In this study, we carry out price elasticity analysis from a statistics point of view. That is, we ignore the difference in the actual computation units, but focus on the total amount and the average price of the computation units.

3.1 Price Elasticity of Server Sales Business

Figure 1 shows the volume of server shipments and revenues, as well as the demand curve for the server sales business. The significant drop in both server shipments and revenue in 2009 was the result of the global economic recession. Using Equation 4, the price elasticity of demand for the worldwide server market is calculated as 0.51 in 2013, i.e., extremely inelastic.

For the majority of enterprise consumers, when they need computing resource the first option that comes to mind is purchasing some server hardware. Other forms of computing resource are not close substitutes for various reasons—renting server hardware from hosting service providers might pose security problems, while public clouds have not been widely adopted. From a budget perspective, server hardware purchases, as fixed-asset investments, are usually decided at the organization level, and are only a small proportion of an organization's annual budget. The lack of close substitutes, and the small proportion of



(a) Total servers and revenues, based on the data in Rackspace's annual reports to SEC.

(b) Annual server growth and average server price, calculated from the data in Rackspace's annual reports to SEC.

(c) The demand curve for Rackspace's server rental business, calculated from the data in Rackspace's annual reports to SEC.

Fig. 2: Server Rental Business: Rackspace

budget, combined contribute to the inelastic demand in server hardware. In other words, organizations invest on server hardware based on their business plans, regardless of price changes in the market.

A fundamental characteristic of inelastic demand is that the percentage change in demand is less than the percentage change in price. In this case, price reduction might produce more demand, but will also result in revenue lost. In a perfectly monopoly industry, the vendor would always do the opposite, i.e., raising the price and harvesting more revenue. The problem is, the competition in the server hardware market is extremely intensive. Leading server vendors are cutting prices to win market share. Regardless of the steady growth in worldwide server shipments, the overall revenue from server sales seems to be decreasing, as shown in Figure 1a. This results in decreasing margin from server sales, and makes it difficult for server vendors to survive.

3.2 Price Elasticity of Server Rental Business


Figure 2a shows Rackspace's number of servers and revenue, as well as the revenue break down for dedicated hosting (server rental business) and public cloud between 2006 and 2013. Figure 2b shows the estimated number of servers and the estimated average revenue per server for Rackspace's dedicated hosting business. It should be noted that the number of worldwide server shipments discussed in Section 3.1 represents the newly added computing resource every year. Furthermore, the life cycle of server hardware is usually more than one year, and in most cases falls within the range of three to five years. The number of servers for Rackspace's dedicated hosting business represents the accumulated computing resource over the past years. To facilitate a fair comparison, we calculate the annual server increments for Rackspace's dedicated hosting business, and use that number to compare with the server sales business. Figure 2c is a supply and demand plot with the average revenue per server on the horizontal axis and the annual server increments on the vertical axis. The calculated price elasticity of demand for Rackspace's dedicated hosting business is 0.70 in 2013, which is modestly inelastic.

Enterprise consumers choose server rental over buying server hardware for various reasons. For small and medium businesses, the cost and effort to build and maintain a small data center with the appropriate networking capacity is overwhelming. There are many server hosting service providers available, including international and local ones. An enterprise usually needs to go through a comparison and selection process before deciding which one to work with. Once a decision has been made, a long-term contract with the selected service provider will be setup. The length of such contract is usually measured by months or even years, depending on the consumer's business plan. Such a long-term contract effectively eliminates the possibility of migrating to alternative options when there is a price change in the market. Thus, enterprise consumers invest on server rental based on business plans, which are usually not sensitive to price changes. This is why the price elasticity of demand for the server rental business is inelastic.

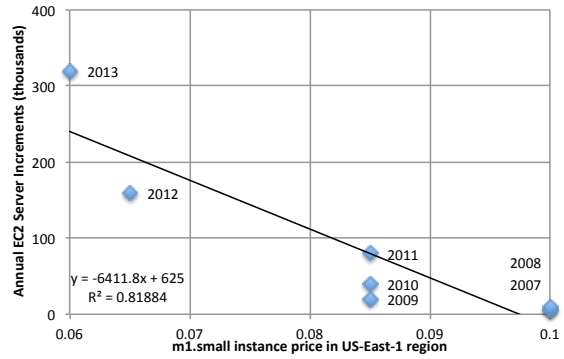
3.3 Price Elasticity of Public Clouds

For the public cloud business, we use the estimated number of servers for AWS EC2 derived by a number of cloud computing professions in the industry as the foundation to carry out our analysis. AWS is currently the largest public cloud service provider. It is commonly recognized that the size of AWS is far larger than all the competitors combined. Figure 3a shows the estimated number of servers for AWS EC2 on a global scale. Figure 3b is a supply and demand plot showing the estimated amount of annual EC2 server increments and the price of m1.small instance in US-East-1 region between 2008 and 2013, with a linear regression trend line. The calculated price elasticity of demand is 1.20 in 2013, which is modestly elastic.

Many consumers consider EC2 as an alternative to traditional server rental services and virtual private server (VPS) services, while ignoring those advanced features such as CloudWatch, AutoScaling, and Elastic Load Balancing. Furthermore, spending on EC2, as an operations expense, is usually decided at the project level, and is usually a big portion of a project's budget. The availability of close substitutes, and the



(a) The estimated number of servers for the AWS EC2 service, estimated and adjusted based on Amazon's quarterly earning report in the second quarter of 2015.



(b) The demand curve for AWS EC2 in the us-east-1 region, calculated based on EC2 pricing history and the estimated number of EC2 servers.

Fig. 3: Public Clouds - AWS EC2

big proportion of budget, combined contribute to the elastic demand for EC2. In other words, enterprise consumers spend on EC2 based on actual business needs rather than business plan, and are sensitive to price changes because of budget limitations.

In general, inelastic demand is closely related to planned buying, while elastic demand is closely related to unplanned buying. With the fixed-asset model, when enterprise consumers need computing resource, the budget for the needed computing resource must first be secured before making an order. Then it takes weeks, if not months, for the needed computing resource to be delivered and deployed. With the utility model, consumers have instant access to computing resource with no capital expense, zero lead-time, at affordable prices. In other words, the utility model makes it possible for customers to react to price changes in a timely manner.

4 CASE STUDIES

In this section, we analyze the impact of different pricing strategies on the performance of the business, using IBM, HP and DELL as case studies; they represent over 75% of the market share. Other vendors including Sun (later acquired by Oracle), Fujitsu, Cisco, Lenovo, etc. combined only represent 25% of the market share.

Figure 4a shows the annual server shipments and the average prices for IBM, HP and DELL. The significant gaps in the average prices indicate that these vendors are not competing in the same fine-grained market segment. IBM dominates the high-end server market, HP occupies the mid-range server market, while DELL dominates the low-end server market.

For IBM, the average price does not change much between 2006 and 2010. The slight decrease in its server shipments is an indicator that the market preference is gradually shifting from high-end servers to mid-range and low-end servers. Between 2010 and 2012, there is an 11% increase in IBM's average price,

with a corresponding 8% decrease in server shipments. The price elasticity of demand for IBM servers is 0.73, which is modestly inelastic. For HP, there is an 18% percent decrease in average price between 2006 and 2013, with a corresponding 13% increase in server shipments. The price elasticity of demand for HP servers is 0.72, which is very similar to IBM. For DELL, there is a 28% increase in average price, with a corresponding 15% increase in server shipments. The simultaneous increase in price and shipments indicates that the price elasticity of demand for DELL servers is extremely inelastic - the demand increases regardless of price increases.

The in-elasticity in demand, plus the significant gaps in the average prices, suggest that enterprise consumers do not perceive IBM, HP and DELL servers as close substitutes for each other. These vendors dominate in their own fine-grained market segment, and have the market power over pricing. IBM responded to the inelastic demand with up-pricing. However, the market preference was shifting from high-end servers to mid-range and low-end servers, resulting in the decrease in revenue. HP responded to the inelastic demand with down-pricing, resulting in increase in shipments but decrease in revenue. DELL responded to the inelastic demand with up-pricing. Coupled with the shift in market preference, DELL achieves a simultaneous increase in server shipments and revenue. The situation for both HP and DELL can be directly derived from Equation (1) - when the demand is elastic, a fall in price leads to a fall in revenue, a rise in price leads to a rise in revenue. Therefore, the HP and DELL cases can be considered as classical examples of how the principle of price elasticity works in microeconomics.

Figure 4c shows the market share for IBM, HP and DELL, as defined by the percentage of server shipments. As shown in this figure, price reduction is not an effective way to win market share.




Fig. 4: Case Studies - HP, IBM and DELL. All the data presented were calculated from the data in Gartner's quarterly reports on worldwide server shipments.

5 CONCLUSIONS

In this paper, we study the price elasticity of demand of the global enterprise computing resource market, including the server sales business, server rental business, and public clouds. We reveal that from a microeconomics point of view public clouds are fundamentally different from traditional server sales business and server rental business. The difference in business model leads a different purchasing behavior, which is reflected by the difference in price elasticity of demand. The price elasticity of demand is extremely inelastic for both server sales business and server rental business. Since consumers are not sensitive to price changes, a rise in price leads to a rise in revenue, while a fall in price results in a fall in revenue. The price elasticity of demand is modestly elastic for public clouds. Since consumers are sensitive to price changes, a fall in price results in a rise in revenue. As such, it is vital for vendors and service providers to carefully evaluate the price elasticity of the product or service they offer when developing pricing strategies.

In our case studies with IBM, HP and DELL, the significant gaps in their average unit prices indicate that the hardware vendors are not competing in the same fine-grained market segment. IBM dominates the high-end server market; HP occupies the mid-range server market; while DELL dominates the low-end server market. These vendors have the market power over pricing in their respective fine-grain market segment. As shown in the historical market share data, price reduction is not an effective way to win market share.


REFERENCES

- [1] A. Marshall, *Principles of Economics*, vol. III, ch. IV. Digireads.com Publishing, 2004.
- [2] G. C. Chow, "Technological change and the demand for computers," *The American Economic Review*, pp. 1117–1130, 1967.
- [3] K. Y. Tam and K. L. Hui, "Price elasticity and the growth of computer spending," *IEEE Transactions on Engineering Management*, vol. 46, no. 2, pp. 190–200, 1999.
- [4] J. Stavins, "Estimating demand elasticities in a differentiated product industry: the personal computer market," *Journal of Economics and Business*, vol. 49, no. 4, pp. 347–367, 1997.

- [5] A. Goolsbee, "Competition in the computer industry: Online versus retail," *The Journal of Industrial Economics*, vol. 49, no. 4, pp. 487–499, 2001.
- [6] O. Rogers and A. Sadowski, "Prices: to cut or not to cut? that is the question facing cloud providers." <http://tinyurl.com/ozjeozh>, 2014.



Qingye Jiang received the B. Eng. Degree in civil engineering from Tsinghua University, China, in 1999 and the M.S. degree in civil engineering from University of Illinois at Urbana-Champaign, USA, in 2001. He is currently pursuing the Master of Philosophy degree in the School of Information Technologies at The University of Sydney. He is a senior member of IEEE.



Young Choon Lee is currently an Australian Research Council DECRA Fellow and a lecturer in the Department of Computing, Macquarie University. His current research interests are in the areas of distributed systems and high performance computing.



Albert Y. Zomaya is the Chair Professor of High Performance Computing & Networking in the School of Information Technologies, The University of Sydney. He is also the Director of the Centre for Distributed and High Performance Computing which was established in late 2009. Professor Zomaya published more than 500 scientific papers and articles and is author, co-author or editor of more than 20 books. He served as the Editor in Chief of the IEEE Transactions on Computers (2011-2014). Also, Professor Zomaya serves as an associate editor for 22 leading journals.

Professor Zomaya is the recipient of the IEEE Technical Committee on Parallel Processing Outstanding Service Award (2011), the IEEE Technical Committee on Scalable Computing Medal for Excellence in Scalable Computing (2011), and the IEEE Computer Society Technical Achievement Award (2014). He is a Chartered Engineer, a Fellow of AAAS, IEEE, IET (UK). Professor Zomaya's research interests are in the areas of parallel and distributed computing and complex systems.