Association for Information Systems
AIS Electronic Library (AISeL)

PACIS 2013 Proceedings

Pacific Asia Conference on Information Systems (PACIS)

6-18-2013

Analyzing the Impact of Brokered Services on the Cloud Computing Market

Richard D. Shang
Agency for Science, Technology and Research (A*STAR), richard-shang@ihpc.a-star.edu.sg

Jianhui Huang
Singapore Management University, jhhuang.2009@phdis.smu.edu.sg

Yinping Yang
Agency for Science, Technology and Research (A*STAR), yangyp@ihpc.a-star.edu.sg

Robert J. Kauffman
Singapore Management University, rkauffman@smu.edu.sg

Follow this and additional works at: http://aisel.aisnet.org/pacis2013

ANALYZING THE IMPACT OF BROKERED SERVICES ON THE CLOUD COMPUTING MARKET

Richard D. Shang, Institute of High Performance Computing, Agency for Science, Technology and Research (A*STAR), Singapore, richard-shang@ihpc.a-star.edu.sg

Jianhui Huang, School of Information Systems, Singapore Management University, Singapore, jhhuang.2009@phdis.smu.edu.sg

Yinping Yang, Institute of High Performance Computing, Agency for Science, Technology and Research (A*STAR), Singapore, yangyp@ihpc.a-star.edu.sg

Robert J. Kauffman, School of Information Systems, Singapore Management University, Singapore, rkauffman@smu.edu.sg

Abstract

This research offers a theoretical model of brokered services and provides an analysis of their impact on the cloud computing market with risk preference-based stratification of client segments. The model structures the decision problem that clients face when they choose among spot, reserved and brokered services. Although all the three types of services do not indemnify the cloud services client against other kinds of service outages, due to changes in market demand, service interruptions occur most frequently in the spot market, and are lower when brokered services are offered, and no risk of interruption is involved in reserved services. Based on our analysis, we show that the profitability and sustainability of a cloud service broker depends on its usage of reserved resources and its capability to mitigate the risk of interruptions. We further enrich our explanation through the consideration of the distribution of clients' risk preferences and the service vendor's pricing decisions for reserved and spot services.

Keywords: Cloud Computing, Cloud Service Brokerage, Economic Analysis, IT Services, Risk, Reserved Services, Spot Services.

1   INTRODUCTION

Cloud service brokerages have emerged as digital intermediaries in the information technology (IT) services market, creating value for cloud computing clients and vendors alike. According to Gartner (2012), which forecasts that annual IT spending on cloud service brokerages will reach $100 billion by 2014, cloud service brokerages derive value from the cloud resource management and brokerage services in different ways. One way is through *customization of services*, so adopting these services will be less risky for clients. Cloud computing services can be tweaked to meet the needs of clients in ways that are value-enhancing for the clients and are, at the same time, economical and profitable for the vendor. Another way for a cloud service broker to provide value is through the *aggregation or bundling of services*. Cloud computing services have become less expensive due to the development of connections between multiple kinds of cloud services, and new technological and managerial approaches to leverage their applications in the market. Aggregation also beneficially decreases the search costs that potential clients experience, since there is a possibility of one-stop shopping for cloud computing services. Furthermore, there is likely a broker brand effect, which will diminish the uncertainty discount that clients ascribe to the services of vendors they do not know. In this sense, the cloud service broker can act as an infomediary and guarantor of the quality of pre-qualified vendors' service.

Earlier research characterizes the main functions of service brokerages as intermediating transactions, and bundling services and products (Bakos et al. 2005). In the context of clouding computing, we are motivated by the market potential of cloud service brokerages and seek to provide theoretical insights into their impact by addressing the following question: How do brokered cloud computing services affect the cloud computing market? We use *risk preference*–the amount of risk the client wishes to take–as a key variable to analyze a client's decision-making process related to her service choice from among reserved, spot, and brokered cloud services.

In the section that follows, we describe existing services in current cloud computing market. We then develop a cost model that incorporates risk measures from financial economics, and derives critical values for a client's risk preference related to threshold values for client segments in the market. We also show the impact of brokered cloud services related to cloud service market shares and the gains that service vendors can make. We further demonstrate how the broker's price premium is affected by its use of reserved resources, its capacity to manage risk, the distribution of client risk preference, and the vendor's decision about pricing its services. We also will discuss how a cloud service broker's position in the market can be sustained.

2   CLOUD COMPUTING SERVICES

There are two types of cloud services offered by existing cloud vendors, and they are resources for cloud service brokers to leverage with. Cloud computing services with reserved resources are pre-committed resources for clients by the vendor (we will refer to this type of resources as *reserved resources*). Reserved resources, priced based on either a subscription or pay-per-use, perform in a way that once a client launches a job, the vendor will set aside capacity for the job until the client terminates it.

The resource acquisition and allocation of cloud computing services with resources priced by the spot market (we will refer to this type of resources as *spot resources*) differ from reserved resources. Spot resources have no commitment on the part of the vendor to guarantee access at a given time, other than via the client's willingness to pay the spot market price for services. The acquisition of resources for spot services varies according to a client's bid price and the changing relationship between supply and demand. Clients submit bids representing the maximum unit prices they are willing to pay for a predefined type of spot service. As soon as the service price in the spot market goes above the client's bid price, the vendor will terminate the service for that client. Clients, as a result, will receive service allocations that are affected by the interplay between supply and demand.

Computing jobs running on spot resources may occasionally be interrupted due to price spikes in the

market. Most of the time, though, spot service prices are below one quarter of all the jobs which use reserved resources. Significant cost savings from using spot resources are attractive for clients who need cloud computing services for compute-intensive jobs. This has led to the design and application of a variety of resource allocation approaches, involving predictive analytics, machine learning, and other models that support value-conscious use of limited server resources (Das et al. 2011; Mazzuco and Dumas 2011). Empowered by those resource management techniques, cloud service brokers are able to provide less costly and more reliable services to clients (Elliott 2012).

3 MODEL

To start with a sufficient simplicity that has a reasonable representation of the real world market, we model a cloud computing market with a monopoly vendor, a service broker, and a pool of clients.

3.1 The Vendor and Broker of Cloud Computing Services

The vendor offers two types of cloud computing services: reserved and spot services. The reserved services are offered to clients with a guarantee of uninterrupted job execution at a fixed price. The spot services are offered to clients in an auction-like manner that may have unintentional termination under certain conditions. When a service interruption occurs to a job that was initiated based on the spot services, clients may need to use other form of computing resources. This can be their in-house computing resources or some other services to which they have a subscription.

The cloud service broker offers cloud computing services based on spot and reserved resources made available by the monopoly vendor. It mitigates the risk of service interruption by dynamically scheduling clients' jobs. The clients still bear some job interruption risk, though at a lower level than spot services. The ratio of reserved versus spot resources used by the broker is a and $1 - a$ respectively (see Table 1 for the details of our modeling notation). The cost of the broker's service is a combination of the reserved resource and spot resource costs with similar proportions. The proportion of reserved resources in the broker's resource management model together with technology advance is what drives the interruption risk downwards.

Variable	Description
JobCost, *InterruptCost*	Vendor cost for completing a job; unit cost of service interruption.
TotalCost	Total cost of executing a job, based on the job and interruption costs.
RiskPref, (Δ)	Client's risk (difference in degree of risk) preference.
SpotJobCost, (*ReservedJobCost*)	Fixed cost for using spot (reserved) resources to execute a job.
%ReservedResources	Proportion of reserved cloud computing resources used by the service broker.
BrokerJobCost	Cost for using the service broker to execute a job, depending on a.
#Interrupts, *Mean#Interrupts*, *SD#Interrupts*	Number (mean, standard deviation) of interruptions for the spot services market.
MeanInterruptRisk (.), *SDInterruptRisk* (.)	Risk-mapping for mean (standard deviation) of interruptions in broker's services.
F(.)	Cumulative distribution function for risk preference of client population.
Revenue, *PricePremium*, *MktShare*	Stakeholder's revenue, price premium, and market share of services consumed.
A	The broker's percentage of use of reserved cloud computing resources.
Δ	Shift in demand from one cloud computing service to another different service.

Table 1.   Modeling notation

3.2 The Clients of Cloud Computing Services

The clients in this market use spot resources to save money when the spot market price is low, but

they face the possibility of a service interruption during the job execution period. We assume a client population of unit mass, with computing jobs that involve parallel execution. We assume that clients will incur a constant additional cost when a service interruption occurs.

3.3 The Mean-Risk Measure and the Total Cost Function

We first consider a base case. We assume clients have identical computing jobs and their objective is to maximize the net gains from executing them. The net gains will be the same for clients irrespective of which of the three alternatives that they use, however, they will experience different total costs.

To address the effects of client uncertainty when job interruptions are possible, we employ a *mean-risk measure* to assess the interruption risk and the extent to which risk preference drives a client's decision (Mookerjee 2000; Wang et al. 2008). It is a straightforward approach to communicate to investment management clients the risk associated with the price volatility of financial investments. Researchers have adopted this measure in studies of IT investment and risk management in IS research (Paleologo 2004; Kauffman and Sougstad 2008). Such measures incorporate the mean and variance of the risk factor that affects the outcome of interest. Variance acts as a proxy for risk that captures the extent of the volatility and predisposition to change of a value-relevant factor in the firm's market or operating environment. In our setting, the mean-risk measure quantifies the effects of risky factors on cloud computing service choice. This enables comparative statics analysis for the various factors.

Clients may perceive the financial impacts of service interruptions to be different, leading them to make different choices among the services (spot, reserved, brokered). Clients who perceive there to be too high a frequency of service interruption for spot services may shift to brokered services, if the broker's services have fewer interruptions at an appropriate price. We capture the client's risk attitude toward uncertainty about financial losses with the variable *RiskPref*.

Based on the mean-risk measure, the total cost of job execution is:

$$TotalCost = JobCost + InterruptCost \cdot [Mean\#Interrupts + (RiskPref \cdot SD\#Interrupts)]$$

The second term in the right-hand side of the equation is the mean-risk measure, which captures the impact of the variance associated with service interruption on the total job cost.

3.4 Modeling Assumptions

We set the following assumptions for this model.

- **Assumption 1 (Risk-Neutral Clients).** *A risk-neutral client, represented by RiskPref = 0, is indifferent about variations in performance metrics such as total cost or utility.*

This rationale of this assumption is consistent with prior research (Mookerjee and Mannino 2000). Risk-neutral clients tend to find spot services to be more attractive and economical compared to reserved services. This suggests that risk-neutral clients will also view the total cost of reserved services to be high compared to the total cost of spot services, i.e., $TotalCost_{Reserved} > TotalCost_{Spot}$. Here, $TotalCost_{Reserved} = JobCost_{Reserved}$, and $TotalCost_{Spot} = JobCost_{Spot} + InterruptCost \cdot Mean\#Interrupts_{spot}$. We represent this information with the following constraint: $JobCost_{Reserved} - JobCost_{Spot} > InterruptCost \cdot Mean\#Interrupts_{spot}$.

- **Assumption 2 (Mean and Standard Deviation of Service Interruption Risk in Broker's Services).** *MeanInterruptRisk(0) = SDInterruptRisk(0) = 1, and MeanInterruptRisk(1) = SDInterruptRisk(1) = 0.*

This assumption reflects the structure that, when the broker uses only spot resources to build its own service without stable computing resources as backup, it will bear the same level of interruption risk as if it were operated as a spot service vendor. When the broker's service is based only on reserved resources though, there is no risk of service interruption.

4 ANALYSIS

We next derive the critical values for the client's risk preference in relation to her choice to use spot, reserved or brokered services. To accomplish this, we solve the set of *TotalCost* equations for different client service choices when the difference between the total costs for each of the services is zero. Then, based on risk preference-based stratification of the client segments, we analyze how brokered services will affect the vendor's total revenue and the condition for brokered services to be sustainable.

4.1 The Cloud Computing Services Market without Brokered Services

The solution for the value of *RiskPref* consistent with values of *TotalCost* for spot and reserved services is:

$$RiskPref_{Spot,Reserved} = \frac{JobCost_{Reserved} - JobCost_{Spot}}{SD\#Interrupts_{Spot} \cdot InterruptCost} - \frac{Mean\#Interrupts_{Spot}}{SD\#Interrupts_{Spot}}$$

Clients with a level of risk preference lower than $RiskPref_{Spot, Reserved}$ will choose to use spot services, while the others will select reserved services.

With this critical value for risk preference, we can derive the market shares for each of the service types and the vendor's revenue from the market. The market share for reserved services is given by $MktShare_{Reserved} = 1 - F(RiskPref_{Spot, Reserved})$, and for spot services it is $MktShare_{Spot} = F(RiskPref_{Spot, Reserved})$. Hence, the existing service vendor's total revenue will be $Revenue_{Vendor} = JobCost_{Reserved} \cdot (1 - F(RiskPref_{Spot, Reserved})) + JobCost_{Spot} \cdot F(RiskPref_{Spot, Reserved})$.

4.2 The Cloud Computing Services Market with Brokered Services

Using a similar analysis approach, we next derive a decision rule for the client to choose one of the three services. We solve for other critical values of *RiskPref* by comparing *TotalCost* for spot versus brokered services, and reserved versus brokered services. We already have derived a similar value for spot versus reserved services, which will be useful for this analysis.

First, we calculate the total cost for clients when they use the brokered services. The cloud service broker has room to make profit with the capability to provide services with a very small probability of service interruption, while utilizing more spot resources and less reserved resources that can keep their average cost close to the average cost of spot services. The broker can charge a price that is higher than the spot price but is lower than the reserved service price to gain a price premium. In our model, we call this the broker's *PricePremium*. Thus, we have

$$Mean\#Interrupts_{Broker} = MeanInterruptRisk(a) \cdot Mean\#Interrupts_{Spot} \tag{1}$$

$$SD\#Interrupts_{Broker} = SDInterruptRisk(a) \cdot SD\#Interrupts_{Spot} \tag{2}$$

$$JobCost_{Broker} = (1 + PricePremium) \cdot (a \cdot JobCost_{Reserved} + (1 - a) \cdot JobCost_{Spot}) \tag{3}$$

With the $JobCost_{Broker}$, $Mean\#Interrupts_{Broker}$ and $SD\#Interrupts_{Broker}$ variables, we can compute the client's risk preference for brokered versus reserved services, and for spot versus brokered services, as follows:

$$RiskPref_{Broker, Reserved} = \frac{JobCost_{Reserved} - JobCost_{Broker}}{SD\#Interrupts_{Broker} \cdot InterruptCost} - \frac{Mean\#Interrupts_{Broker}}{SD\#Interrupts_{Broker}}$$

$$RiskPref_{Spot,Broker} = \frac{JobCost_{Broker} - JobCost_{Spot}}{\left(SD\#Interrupts_{Spot} - SD\#Interrupts_{Broker}\right) \cdot InterruptCost} - \frac{Mean\#Interrupts_{Spot} - Mean\#Interrupts_{Broker}}{SD\#Interrupts_{Spot} - SD\#Interrupts_{Broker}}$$

Existing cloud service vendor's revenue. With the results we have established, we can construct the equation for the service vendor's revenue when the broker enters the market, $Revenue_{Vendor}'$, as fol-

lows:

$$Revenue_{Vendor}' = JobCost_{Reserved}$$

$$\cdot (1 - F(RiskPref_{Broker, Reserved}) + a \cdot (F(RiskPref_{Broker, Reserved}) - F(RiskPref_{Spot, Broker})))$$

$$+ JobCost_{Spot} \cdot (F_{risk}(RiskPref_{Spot, Broker}) + (1 - a) \cdot (F(RiskPref_{Broker, Reserved}) - F(RiskPref_{Spot, Broker})))$$

We next consider the change in revenue that occurs for the vendor when the broker enters the market, as follows:

$$\triangle Revenue = JobCost_{Reserved} \cdot (F(RiskPref_{Spot, Reserved})$$

$$- (1 - a) \cdot F(RiskPref_{Broker, Reserved}) - a \cdot F(RiskPref_{Spot, Broker}))$$

$$+ JobCost_{Spot} \cdot (a \cdot F(RiskPref_{Spot, Broker})$$

$$+ (1 - a) \cdot F(RiskPref_{Broker, Reserved}) - F(RiskPref_{Spot, Reserved}))$$

Based on this analysis approach, not all clients will be willing to choose brokered services. This is because the broker's price premium, consistent with the client's value, may exceed the reduction in *TotalCost* due to the decrease in perceived risk for the brokered services.

The broker's profit and its constraints. In the market in which spot and reserved services are offered, clients with risk preference higher than $RiskPref_{Spot, Reserved}$ will purchase the reserved service and the rest will use the spot service. When the broker enters the cloud market, a portion of the clients who utilize spot and reserved services will shift to the brokered services. Clients with risk preference higher than $RiskPref_{Spot, Broker}$ and lower than $RiskPref_{Broker, Reserved}$ will adopt the brokered service. The shift in demand from reserved to brokered services is represented by $\delta_{Reserved, Broker} = RiskPref_{Broker, Reserved} - RiskPref_{Spot, Reserved}$. And the shift in demand from spot services to brokered service is represented by $\delta_{Spot, Broker} = RiskPref_{Spot, Reserved} - RiskPref_{Spot, Broker}$.

For brokered services to be feasible, two conditions must be satisfied to guarantee that the market shares for each of the three types of services are positive, that is, $\delta_{Reserved, Broker} > 0$ and $\delta_{Spot, Broker} > 0$. Although we cannot see it directly by visually assessing the mathematical forms of the critical value functions for *RiskPref*, it can be shown that the former condition will dominate the latter one. This is true in the sense that, if the first condition is satisfied then the second one will automatically be satisfied. This leads to the following condition for the brokered services:

$$(JobCost_{Reserved} - JobCost_{Broker}) \cdot SD\#Interrupts_{Broker} - (JobCost_{Reserved} - JobCost_{Spot})$$

$$\cdot SD\#Interrupts_{Broker} - Mean\#Interrupts_{Broker} \cdot InterruptCost \cdot SD\#Interrupts_{Spot}$$

$$+ Mean\#Interrupts_{Spot} \cdot InterruptCost \cdot SD\#Interrupts_{Broker} > 0 \qquad (4)$$

4.3 The Brokered Services Feasibility Condition

By combining Equations 1, 2, 3 and 4, we obtain the following proposition:

- **Proposition (Brokered Services Feasibility Condition).** *A feasibility condition for the brokered services is*:

$$JobCost_{Reserved} \cdot (1 - a - a \cdot PricePremium - SDInterruptRisk(a)) + JobCost_{Spot}$$

$$\cdot (a + a \cdot PricePremium + SDInterruptRisk(a) - 1 - PricePremium) + InterruptCost \quad (5)$$

$$\cdot Mean\#Interrupts_{Spot} \cdot (SDInterruptRisk(a) - MeanInterruptRisk(a)) > 0$$

This condition will be satisfied regardless of the scale of $JobCost_{Reserved}$ and $JobCost_{Spot}$ when $SDInterruptRisk(a) > MeanInterruptRisk(a)$ and *PricePremium* = 0. Because the left-hand side of Equation 5 is decreasing in the broker's *PricePremium*, we can derive an upper bound for the *PricePremium* when the Brokered Services Feasibility Condition holds, and is denoted by *UB[PricePremium]*:

$$[(JobCost_{Reserved} - JobCost_{Spot}) \cdot (1 - a - SDInterruptRisk(a)) + InterruptCost \cdot Mean\#Interrupts_{Spot}$$

$$\cdot (SDInterruptRisk(a) - MeanInterruptRisk(a))] / [JobCost_{Reserved} \cdot a + SpotMktJobCost \cdot (1 - a)]$$

To be sustainable, while cloud service brokers benefit clients with lower-risk and lower-cost services, brokered services should also benefit the existing cloud services vendor by increasing the vendor's revenue. Otherwise, the vendor will stop offering its resources to the broker. This requires another condition to be true. Because the broker's service cost will incur to the clients who originally used reserved or spot services, to benefit the vendor, the broker's service cost should be larger than what the vendor incurred to this group of clients prior to the entry of the broker. The differences in the critical values of *RiskPref* act as proxies for the broker's market share. Based on additional mathematical analysis involving the critical values for *RiskPref* for the different services, we can state the conditions under which brokered services will add value for the vendor:

$$a > \frac{F(RiskPref_{Broker,Reserved}) - F(RiskPref_{Spot,Reserved})}{F(RiskPref_{Broker,Reserved}) - F(RiskPref_{Spot,Broker})}$$

With these constraints that ensure brokered services are feasible and sustainable, we now can formalize the broker's decision problem as:

$$Max_{(a, PricePremium)} \; PricePremium \cdot [\, a \cdot JobCost_{Reserved} + (1 - a) \cdot JobCost_{Spot}]$$

$$\cdot [F(RiskPref_{Broker, Reserved}) - F(RiskPref_{Spot, Broker})]$$

$$\text{subject to} \; a > \frac{F(RiskPref_{Broker,Reserved}) - F(RiskPref_{Spot,Reserved})}{F(RiskPref_{Broker,Reserved}) - F(RiskPref_{Spot,Broker})}$$

$$PricePremium < UB[PricePremium]$$

5   NUMERICAL ILLUSTRATION

We provide an illustration based on the parameters in Table 2 to show the constraints to the brokered services due to the feasibility condition. The parameters satisfy the modeling assumptions introduced in Section 3.4, and the risk-mapping functions for the mean and standard deviation of broker's job interruptions are the same.

$JobCost_{Reserved}$	4
$JobCost_{Spot}$	1
$InterruptCost$	1
$Mean\#Interrupts_{Spot}$	1
$SD\#Interrupts_{Spot}$	0.5

Table 2.   Parameter settings for illustrative numerical analysis of brokered services

We assume a uniform distribution for clients' risk references, and thus constrain a is simplified as to be greater than $\frac{1 - SDInterruptRisk(a)}{SDInterruptRisk(a)}$. We also constraint the broker's price premium to the upper bound $UB[PricePremium]$. We randomly generate 100 combinations of values of percentage of reserved resource usage and risk mapping function that satisfy constraints, and use the values to calculate the broker's profit and the vendor's gain from brokered services. The results suggest that while reducing reserved resource usage and increasing its capability to mitigate interruption risk improve the broker's profitability, such an approach decreases the vendor's gain from the brokered service. Figure 1 illustrates the conflict of interest between the broker and the vendor. We rank the combinations according to the broker's profit, and the figure shows that the combinations that generate the maximum profit for the broker result in the lowest gain for the vendor. When we change the parameter setting of $JobCost_{Reserved}$ to 5 from 4, we can observe a shift-up of vendor's gain associated with a shift-down of broker's profit. This finding suggests that though new technological approaches may further enhance the broker's capability of resource utilization and risk mitigation, to ensure brokered services sustainable, the broker needs to balance the benefits of and the vendor's potential strategic responses to implementing the new techniques.
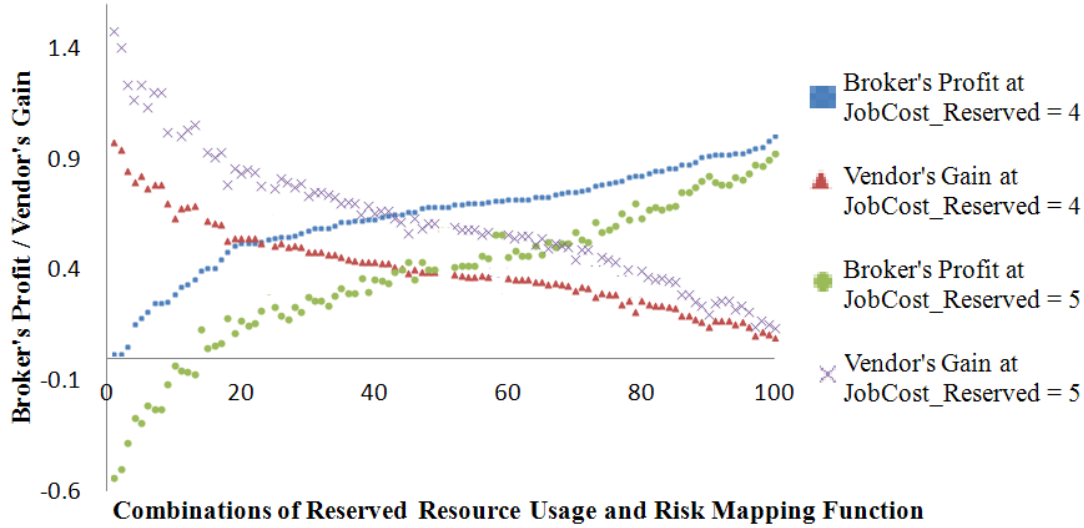
Figure 1. *Broker's profits and vendor's gain at different combinations of percentage of reserved resources usage and the risk-mapping function*

6 DISCUSSION

This research-in-progress provides analytical insights into quantifying the feasibility conditions of brokered cloud services. The condition is constrained by the proportion of reserved services resources the broker uses, and the capability the broker possesses to mitigate the interruption risk the clients face. Therefore, it is a useful tool for a cloud service broker to assess its business feasibility based on the risk mitigation capability, vendor's pricing, and client related parameters.

The cloud brokers have a room to promote the use of cloud services by customizing them for cost and risk-conscious clients. This is similar to what financial institutions do that provide portfolio management services. They attract investors who want to achieve some balance between investing with some risk, while achieving an attractive level of return. The brokers, however, must do nothing to harm the value that the vendors already create in the market, while doing new things the vendors are not capable of to increase revenue and welfare. Meanwhile, the clients should also be better off, and are willing to pay for and engage the brokered services. This ties in with how client segments and market shares arise in the presence of different approaches to the provision of cloud services.

There is space for us to extend the model. One approach is to analyze the interplay of the pricing strategies of the broker and the vendor. Our numerical illustration already showed the conflict of interest between the broker and the vendor, and the impact of the vendor's price change. We can extend the model to include the vendor's strategic price changes of reserved and spot resources, and the broker's change of price premium. Another approach is to look at market saturation. Current model assumes a saturated market, and it results in the strict constraints for brokered services to be sustainable because the broker's gain of market share will be at the provider's loss. If there are previously unsatisfied demands, both brokers and vendors will benefit from market shares arise in the presence of brokered cloud services.

References

Bakos, Y., Lucas, H.C., Jr., Oh, W., Simon, G., Viswanathan, S. and Weber, B.W. (2005). The impact of e-commerce on competition in the retail brokerage industry. *Information Systems Research*, 16 (4), 352-371.

Das, S., Du, A., Gopal, R. and Ramesh, R. (2011). Risk management and optimal pricing in online storage grids. *Information Systems Research*, 22 (4), 756-773.

Elliott, S. (2012). PiCloud and Princeton Consultants win the first Amazon EC2 spotathon. Amazon Web Services Blog, December 3, 2012. Available at aws.typepad.com/aws/2012/12/picloud-and-princeton-consultants-win-the-first-amazon-ec2-spotathon.html.

Gartner. (2012). Predicts 2013: Cloud Services Brokerage, Gartner Research Report, ID Number: G00239877, December 18, 2012.

Kauffman, R. and Sougstad, R. (2008). Risk management of contract portfolios in IT services: The profit-at-risk approach. *Journal of Management Information Systems*, 25(1), 17-48.

Mazzucco, M. and Dumas, M. (2011). Achieving performance and availability guarantees with spot instances. In *Proceedings of the 2011 IEEE International Conferences on High Performance Computing and Communications (HPCC 2011)*, p. 296-303, IEEE Computer Society, Banff, Alberta, Canada.

Mookerjee, V.S. and Mannino, M.V. (2000). Mean-risk trade-offs in inductive expert systems. *Information Systems Research*, 11 (2), 137-158.

Paleologo, G. (2004). Price-at-risk: A methodology for pricing utility computing services. *IBM Systems Journal*, 43(1), 20-31.

Wang, J.G., Chaudhury, A. and Rao, H.R. (2008). A value-at-risk approach to information security investment. *Information Systems Research*, 19 (1), 106-120.