

# Pricing of insurance policies against cloud storage price rises

Loretta Mastroeni  
Università di Roma Tre  
Dip. di Economia  
Rome, Italy

mastroen@uniroma3.it

Maurizio Naldi  
Università di Roma Tor Vergata  
Dip. di Ingegneria Civile e Ingegneria Informatica  
Rome, Italy

naldi@disp.uniroma2.it

## ABSTRACT

When a company migrates to cloud storage, the way back is neither fast nor cheap. The company is then locked up in the storage contract and exposed to upward market prices, which reduce the company's profit and may even bring it below zero. We propose a protection means based on an insurance contract, by which the cloud purchaser is indemnified when the current storage price exceeds a pre-defined threshold. By applying the financial options theory, we provide a formula for the insurance price (the premium). By using historical data on market prices for disks, we apply the formula in realistic scenarios. We show that the premium grows nearly quadratically with the duration of the coverage period as long as this is below one year, but grows more slowly, though faster than linearly, over longer coverage periods.

## Categories and Subject Descriptors

H.3 [Information Storage and Retrieval]: Information Storage; K.6 [Management of Computing and Information Systems]: Project and People Management

## Keywords

Cloud storage, Pricing, Insurance

## 1. INTRODUCTION

Cloud storage is a major component of the virtualization process. In cloud storage, a customer stores its data on the facilities of a cloud provider and can then access them through the Internet. The cloud replaces the customer's data infrastructure and allows the cloud purchaser to pursue an outsourcing strategy: it is a major example of the IaaS (Infrastructure as a Service) paradigm.

Cloud storage service is paid for through a periodic fee related to the amount of data stored. Several papers have been devoted to the economical analysis of the migration from an owned infrastructure to the cloud [4, 7, 6]. The profitability of migration relies heavily on the price required by the cloud provider: too high a price erodes the profit margins and may make infrastructure ownership the best business proposition. Since the migration decision must be evaluated in the long run [6], any decision has to be based on forecasts for future prices and storage needs. Given the

long evaluation window, there is a significant probability of taking the wrong decision. In [6], the Value-at-Risk metric was employed to analyse the extent of such risk. The cloud purchaser has to deal with such a risk and look for mitigating strategies.

In [1] a technology-based strategy was proposed to reduce the problems due to price hikes, through the application of RAID-like techniques (*Redundant Array of Independent Disks*). That strategy is not free, since it involves managing multiple storage contracts and the extra storage costs to achieve redundancy.

Here we propose an insurance-based strategy, whereby the cloud purchaser underwrites an insurance policy to hedge against price increases. This may be employed as either an alternative or an accompanying strategy for any technology-based solution. Actually, since a downward trend is expected for cloud storage prices, our insurance approach gets protection against any upward deviation from the downward trend. We envisage an insurance contract covering a long-term migration strategy. Since a major issue is the cost associated to the strategy, we derive a pricing formula for the insurance contract and evaluate the resulting price for several realistic scenarios. We show that the insurance price (the *premium*) grows nearly quadratically with the length of the coverage period as long as this is below one year, but grows more slowly, though faster than linearly, over longer coverage periods.

The paper is organized as follows. In Section 2, we review the current cloud storage market and describe a model for future price fluctuations. The insurance policy and the pricing formula are described in Section 3. In Section 4, we set some sample scenarios and evaluate the resulting premium.

## 2. PRICE FLUCTUATIONS

In order to get established as an operational paradigm, cloud storage must be economically profitable for both the cloud provider and the cloud purchaser. Costs for the cloud purchaser are determined by the prices offered by the cloud provider, which in turn are lower bounded by its costs. Such costs pose different constraints on the two parties. Cloud providers must set prices at least so high as to recover their costs: costs represent a lower bound for prices. Instead, cloud purchasers seek prices lower than their costs (otherwise they would stay with their data center): costs represent an upper bound for prices. Whatever the role of the party, we expect costs to vary over time, under the influence of a number of factors (e.g., cost of disks, operational expenses, level of competition). Likewise, we expect prices to reflect

**Table 1: Unit prices of cloud providers**

Provider	Unit price [€]	
	Consumers	Business
Average	0.0955	0.185
Standard deviation	0.0663	0.145

those changes and vary as well. We can model the fluctuations of price through a stochastic model. In this section, starting from data derived from the market of disks and storage facilities, we describe the stochastic model we employ for future cloud storage prices.

All the most important providers charge a fixed price for each time period (either a month or a year), with a maximum amount of storage space (increased in brackets). We have examined the price plans offered by major cloud providers, which offer separate plans for business customers and consumers, one of the main differences being that consumers can access their data from just one computer. Since the unit price decreases as the amount of stored data increases, we have taken the least cost for each provider, i.e., that corresponding to the maximum amount of data. The results of this analysis are summarized in Table 1, the prices being expressed in euros per GB and per month (prices in US dollars have been converted into euros through a conversion rate of 1.3). Prices show a significant dispersion around their average value, though we must recognize that the price itself is just one of the characteristics of the service bundle: different prices may reflect quite different service features.

The significant dispersion may be due to the infancy of the storage service, where cloud providers have quite different backgrounds and operational strategies. However, they are just present prices and tell us very little about the future trends of prices. Hence, they do not provide enough information to evaluate the profitability of the cloud approach, especially for the typically long range adopted to evaluate the migration to cloud storage as in [6].

For that purpose, we need a forecast of future storage prices. Unfortunately, the time series of cloud storage prices are not long and rich enough to draw conclusions. More data are instead available on market prices for disks. In fact, we can use the data coming from an extensive survey conducted on SATA (*Serial Advanced Technology Attachment*) disk prices every week over more than 5 years. That survey has shown that the price  $P(t)$  of disks follows a decaying trend, which is approximately exponential [7]:

$$P(t) \simeq P_0 e^{-\beta t}, \quad (1)$$

where  $P_0$  is the price at the beginning of the observation period,  $t$  is the time expressed in years, and  $\beta = 0.438$ .

In the survey reported in [7], Equation (1) was taken as a deterministic one. Actually, that equation was obtained by regression, since real prices fluctuate around the values predicted by that exponential curve. In [5] we have abandoned the deterministic model, assuming instead that the future price is described by a Geometric Brownian Motion (GBM) random process [2], where the price at time  $t$  is

$$P(t) = P_0 e^{-(\beta + \sigma^2/2)t + \sigma W(t)}, \quad (2)$$

with  $W(t)$  being a Wiener process, and  $\sigma^2$  the variance of

the process. The expected value of the stochastic process described by Equation (2) is exactly that given by Equation (1) Since cloud providers have to purchase storage facilities and operate them, we expect their costs (and likewise their prices) to follow a process similar to that described by Equation (2). We adopt therefore the Geometric Brownian Motion as a suitable model for future cloud storage prices.

### 3. INSURANCE COVERAGE AND PRICING

In Section 2, we have set a model to describe price fluctuations. Though we expect prices to go downward, they may be higher than expected. A purchaser of storage space has to protect itself against upward fluctuations. It can do so by buying an insurance policy. In this section, we formulate the insurance policy as a multiperiod call option, and derive the fair premium for such an insurance policy.

As shown in Section 2, prices of cloud storage can go above and below their expected (negative exponential) trend. That trend is what the cloud storage purchaser expects to pay in the future. If the actual price is lower than what is predicted by the exponential curve (1), the convenience margins are higher than expected. Instead, in the case of upward fluctuations of prices, the profit margins shrink and may get negative: upward fluctuations represent a risk for the cloud purchaser. We can define the risk  $R(t)$  incurred at time  $t$  for each unit of storage (e.g., a GB) as the difference between the actual unit price  $S(t)$  and that predicted by Equation (1)

$$R(t) = (S(t) - \mathbb{E}[P(t)])|_+. \quad (3)$$

The cloud purchaser seeks protection against that risk. A suitable insurance policy may consist in the payment to the insured (the cloud purchaser) of a sum equal to  $R(t)$  whenever the actual price exceeds the expected one, against a payment by the insured of an initial premium (the price of the insurance contract) when the insurance contract is underwritten. By means of this policy, the cloud purchaser is completely indemnified against any price increase above the expected value  $\mathbb{E}[P(t)]$ . The main problem is now the computation of the fair value  $\mathcal{V}$  of the premium that the cloud purchaser has to pay to buy the insurance policy (at the time conventionally set equal to zero). We have also to define when the comparison represented by Equation (3) is performed to assess if the insured has the right to receive the compensation. A natural solution is to use the same time period employed in pricing plans (e.g., every month). If we consider the timing at which the purchase contract is renewed (e.g., each month), we have a sequence  $\{t_1, t_2, \dots, t_n\}$  of times when the cloud purchaser renews its contract with the cloud provider and compares the actual price with the expected one. If the actual price is larger than the expected one, the insured (the cloud purchaser) obtains the difference from the insurer. The length  $n$  of the sequence is determined by the insurance contract. The sequence of cash flows for the cloud purchaser is

$$C_i = \begin{cases} -\mathcal{V} & \text{if } i = 0 \\ R(t_i) & \text{if } i = 1, 2, \dots, n \end{cases} \quad (4)$$

Each individual cash flow falling at time  $t_i$  is exactly what we would obtain with a call option. Call options are financial contracts whose reward for their holder is determined by the difference between the actual price of an underlying asset and a predefined exercise price [3]. In practice, if

the actual price is larger than the exercise price, the option holder gets the difference. In our case, the underlying asset is the cloud storage service, the exercise price is the expected price  $\mathbb{E}[P(t)] = P_0 e^{-\beta t}$ , and the time at which the option can be exercised is each  $t_i$ . Since the cloud service contract renewal is done at a number of times  $t_i$ , the insurance contract is equivalent to a set of call options, one for each time  $t_i$  (a multiperiod call option). As in the case of insurance contracts, a call option guarantees a right for its holder, and is therefore to be paid for.

For a single call option, when the underlying asset price follows a GBM process, the fair price is given by the well known Black-Scholes formula. For the coverage against a price increase at the time  $t_i$ , the general expression of the fair price  $C(t_i)$  of the option (the insurance contract) is

$$V(t_i) = S(0)\mathbb{G}(d_1) - K(t_i)e^{-rt_i}\mathbb{G}(d_2), \quad (5)$$

with

$$d_1 = \frac{\ln \frac{S(0)}{K(t_i)} + \left(r + \frac{\sigma^2}{2}\right)t_i}{\sigma\sqrt{t_i}} \quad d_2 = \frac{\ln \frac{S(0)}{K(t_i)} + \left(r - \frac{\sigma^2}{2}\right)t_i}{\sigma\sqrt{t_i}}, \quad (6)$$

where  $\mathbb{G}(\cdot)$  is the standard normal distribution function,  $K(\cdot)$  is the exercise price,  $r$  is the risk-free interest rate, and  $\sigma$  is the standard deviation of the value of the underlying asset (i.e., the cloud storage unit price).

In our case, since the present price of the service is  $S(0) = P_0$  and the exercise price at time  $t_i$  is  $K(t_i) = P_0 e^{-\beta t_i}$ , Equation (5) becomes

$$V(t_i) = P_0 \left[ \mathbb{G}(d_1) - e^{-(\beta+r)t_i} \mathbb{G}(d_2) \right] \quad (7)$$

with

$$d_1 = \frac{r + \beta + \sigma^2/2}{\sigma} \sqrt{t_i} \quad d_2 = \frac{r + \beta - \sigma^2/2}{\sigma} \sqrt{t_i}. \quad (8)$$

Finally, the overall price for the insurance contract is simply the sum of the prices pertaining to the single call options

$$\mathcal{V} = \sum_{i=1}^n V(t_i) = P_0 \sum_{i=1}^n \left[ \mathbb{G}(d_1) - e^{-(\beta+r)t_i} \mathbb{G}(d_2) \right] \quad (9)$$

#### 4. A SAMPLE APPLICATION

In Section 3, we have derived the pricing formula for our insurance policy. The resulting price is a function of the current unit price, the risk-free interest rate, the rate of storage price decrease, and the duration of the insurance scheme. In this section, we set suitable values for these quantities and compute the insurance price for several realistic scenarios.

We start by dealing with the standard deviation of cloud storage prices. In Section 2, we have analysed the current market prices for cloud storage for consumers and business customers. Here we employ the standard deviation of prices, as reported in Table 1, for  $\sigma$  in Equations (8).

As to the risk-free interest rate, we should consider a government bond of the highest credit merit. For that purpose, we consider US Treasury bills and bonds, with a duration equal to that of the insurance policy. In the following, we consider two rates: 0.2% for 1 year and 0.99% for 5 years.

Since the granularity that cloud providers adopt for their pricing plans is monthly (see Section 2), we assume that the time step in Equation (9) is likewise monthly. Our insurance policy includes therefore a single payment (the premium) at time 0 and a sequence of monthly claims.


**Table 2: Insurance prices**

Category	Period [years]	Insurance price
		Present monthly price
Consumer	1	2.469
Consumer	5	36.504
Business	1	2.481
Business	5	36.516

We also note that the premium in Equation (9) is proportional to the current price of cloud storage. In the following, we refer to the normalized premium  $\mathcal{V}/P_0$ . It is however to be noted that the normalizing quantity  $P_0$  is the present monthly price, while the premium refers to the overall duration of the insurance policy (1 or 5 years in our case). Both quantities apply, however, to the same unit of storage space.

We can now compute the premium for four different cases, obtained by combining the two customer categories and two contract durations. In Table 2, we report the resulting insurance premium. We can examine the impact of several quantities intervening in the price formula (9). Though the two categories of customers differ for the standard deviation of prices (with the standard deviation of prices for business customers being more than twice that for consumers), the differences in the premium to be paid look negligible, quite less than 1%, since the differences in the volatility of prices are overshadowed by the overall downward trend. The duration (which also brings along a variation in the risk-free interest rate) has instead a relevant role: the premium for an insurance policy covering 5 years is nearly 15 times as large as that for a single year.

We can further analyse that issue by observing the normalized individual monthly premium  $V(t_i)$  to be paid at time 0 to cover a single monthly period  $i$  months ahead. In Figure 1, we report the resulting  $V(t_i)$  for a contract covering five years. Over the relatively short period of 1 year, the




**Figure 1: Normalized monthly premium for a single period over five years**

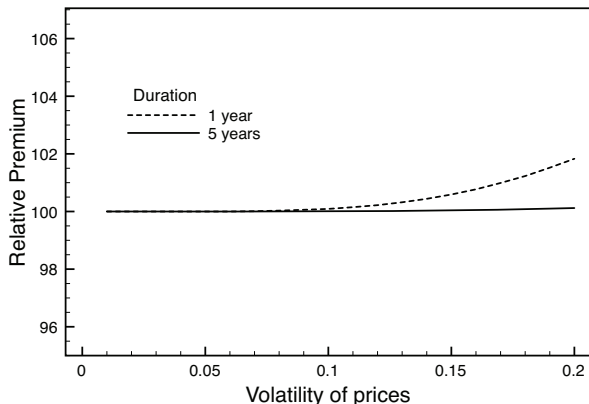
monthly premium grows roughly linearly. The premium for the first month is just 3.6% of the current price, but grows up to 35.6% for the premium to be paid for the twelfth month. Because of the quasi-linear growth of the monthly premium,

over short periods we expect the cumulative premium  $\mathcal{V}$  to grow as the sum of an arithmetic series, i.e., roughly as the square of the number of months. Instead, when we look at the trend over a longer policy duration, the growth of the monthly premium becomes quite less than linear (and correspondingly the growth of the cumulative premium will be quite less than quadratic). For the farthest month, the monthly premium reaches nearly 90% of the current storage price (against the asymptotic limit of 100%).

The differences in the single monthly premium build up when we consider the overall premium. In Figure 2, we show the actual relation between premium and duration for the consumer case, where we have applied the risk-free rates pertaining to the duration.



**Figure 2: Impact of contract duration on the overall premium**



**Figure 3: Impact of the volatility of prices on the premium**

Finally, we consider the effect of the volatility  $\sigma$  of cloud storage prices.

As to the former quantity, in Section 2 we obtained an estimate for the dispersion of prices, which impacts on the premium: an incorrect estimate of the volatility would distort the insurance price. In Figure 3, we draw the relation between the premium and the volatility of cloud storage prices for the two contract durations considered so far, using the premium for the minimum value of the standard deviation considered (0.01) as the baseline value. For both durations, the larger the uncertainty on the storage price (larger variance), the larger the premium. But the effect is negligible for the longer duration: over the whole range considered for the standard deviation of prices, the premium increases by 1.83% if the duration is 1 year, but just by 0.12% if the duration is 5 years.

## 5. CONCLUSIONS

Cloud storage prices are subject to changes, exposing the cloud purchaser to upward price fluctuations. But switching back from the cloud to in-house infrastructure is neither fast nor cheap, and neither is switching to another cloud provider. We have introduced an insurance scheme that protects the cloud purchaser against cloud prices exceeding the expected ones. We provide an insurance pricing formula, which delivers the premium for an insurance policy covering a sequence of monthly claims. The resulting premium has been evaluated for some realistic scenarios and grows first quadratically with the duration of the coverage period and then more slowly. The premium for a single month is initially a very small fraction of the current monthly cloud storage price, but can become even 90% of that when the covered month is 5 years from now. The impact of the dispersion of storage prices on the premium is quite small for short durations and practically negligible for longer ones.

## 6. REFERENCES

- [1] H. Abu-Libdeh, L. Princehouse, and H. Weatherspoon. Racs: a case for cloud storage diversity. In *Proceedings of the 1st ACM Symposium on Cloud Computing, SoCC 2010, Indianapolis, Indiana, USA*, pages 229–240, June 10–11, 2010.
- [2] N. H. Chan and H. Y. Wong. *Simulation techniques in financial risk management*. John Wiley, 2006.
- [3] J. Hull. *Options, Futures, and Other Derivatives (7th Edition)*. Prentice Hall, May 2008.
- [4] A. Khajeh-Hosseini, D. Greenwood, J. W. Smith, and I. Sommerville. The cloud adoption toolkit: supporting cloud adoption decisions in the enterprise. *Softw., Pract. Exper.*, 42(4):447–465, 2012.
- [5] L. Mastroeni and M. Naldi. Storage Buy-or-Lease Decisions in Cloud Computing under Price Uncertainty. In *7th EuroNF Conference on Next Generation Internet*, Kaiserslautern, 2011.
- [6] L. Mastroeni and M. Naldi. Long-range evaluation of risk in the migration to cloud storage. In *13th IEEE Conference on Commerce and Enterprise Computing, CEC 2011*, pages 260–266, September 5–7, 2011.
- [7] E. Walker, W. Brisken, and J. Romney. To lease or not to lease from storage clouds. *IEEE Computer*, 43(4):44–50, 2010.