An experiment in SLA decision-making

Ulrik Franke[1], Markus Buschle[2] and Magnus Österlind[2]

[1] FOI – Swedish Defence Research Agency
SE-16490 Stockholm, Sweden
ulrik.franke@foi.se
[2] Industrial Information and Control Systems
KTH – Royal Institute of Technology
Osquldas v. 12, SE-10044 Stockholm, Sweden
{markusb, magnuso}@ics.kth.se

Abstract. Decision-making with regard to availability service level agreements (SLAs) is investigated. An experimental economics approach was used to elicit the preferences for different SLA alternatives from the subjects ($N = 16$), all professionally working with IT management. A previously published scenario on downtime costs in the retail business was used in the experimental setup. Subjects made 18 pairwise choices under uncertainty. After the experiment, they were paid based on one of their choices, randomly selected. The subjects rarely behaved as expected utility maximizers in the experiment. This raises questions about company SLA management in real situations, and calls for further research.

Keywords: Service level agreements, availability, SLA management, decision-making, experiment

1 Introduction

Today, IT is increasingly being provisioned as a service. Distributed systems technology provides the basis of the "cloud", where enterprises can buy advanced IT services "off the shelf", gaining flexibility and scalability. However, the economic implications are just as important to investigate as the technology [1].

A key non-functional property of IT services bought and sold is *availability*. Annual costs of unplanned downtime were in the billion dollar range already 15 years ago [2], and have hardly improved since. Stock prices fall when business operations are disrupted by IT incidents [3, 4], and reliability costs rank as an important IT frustration for executives [5]. However, to maintain high availability today, IT executives need proper service level agreements (SLAs). Such contracts link business operations to the IT services bought off the shelf.

How to write proper SLAs is interesting both to academia and practitioners. Management by contract [6] can be said to be at the heart of this research area, along with the primacy of the business perspective [7, 8] and the fact that negotiations have to take place between parties with asymmetric information [9]. Gartner [10] and ITIL [11] offer practical advice on availability SLA writing.

The research question of this paper is: Do practitioners deviate from expected utility when procuring availability SLAs, and if so, how? Previous work identifies many potential deviations, e.g. bounded rationality [12, 13] and overconfidence [14]. This study extends previous theoretical work [15] with an empirical investigation. Our results show that practitioners do not necessarily behave as expected utility maximizers – indeed, they do so quite rarely in our experiment.

1.1 Outline

The remainder of the paper is structured as follows: Section 2 covers related work. Section 3 presents the availability investment model used in the experiment. Data collection methods are detailed in Section 4, followed by results in Section 5. Section 6 relates the outcome to previous findings and discusses the results. Finally, Section 7 offers some concluding remarks.

2 Related work

Optimal SLA management is a growing field. [16] offers models for optimal service-window scheduling to minimize business impact, but does not address *unplanned* outages. [17] derives optimal SLA strategies, but does not focus on availability. [18] considers SLA specifications, but without quantitative risk analysis. [19] investigates the service procurer's optimization problem, but does not empirically study human decision-making. Neither does the game theoretic approach of [20]. Technically oriented work such as frameworks for bridging SLA templates [21] or intelligent SLA negotiation agents [22] are important for well-designed SLAs, but does not further our understanding of human decision-making.

Turning to decision-making research, [23] presents a game theoretic framework for SLA negotiation. A bargaining process is envisioned, where an equilibrium between client and service provider is found by counter-offers. This is different from our study, where the client is offered a *take-it-or-leave-it* contract.

[24] presents a study on decision-making for duplex gambles where 34 undergraduate statistics students played hypothetical gambles. The study shows that in the loosing form of gambles (like those in our study, where the decision maker cannot gain money from the gamble) a majority of respondents (78%) maximize the expected value of the gamble, being highly consistent. A similar study with 42 undergraduate psychology students is presented in [25], with results again showing that most respondents are maximizing the expected value.

3 The decision-making problem

SLAs govern many non-functional requirements, but our focus is on *availability*. The *average availability* can be computed as the Mean Time To Failure (MTTF) divided with the total time of operation, i.e. the sum of MTTF and the Mean Time To Repair/Restore (MTTR) [26]:

$$A = \frac{\text{MTTF}}{\text{MTTF} + \text{MTTR}} \tag{1}$$

Availability is a good experimental topic for many reasons: Requirements are easy to understand, there is a tangible economic impact, and it is often at the heart of SLAs. In the experiment (cf. Section 4), participants were subjected to a decision-making problem re-used from [15], where more details can be found.

3.1 A simple investment model

Availability investments have diminishing returns. Each additional hour of up-time comes at a higher cost. This is modeled by Eq. (2)

$$A = f(A_0, c) = 1 - (1 - A_0)e^{-\alpha c} \tag{2}$$

where $A \in [0, 1]$ is the availability resulting from an investment $c \geq 0$ made at an initial availability level $A_0 \in [0, 1]$, where $\alpha \in (0, 1)$ determines the shape of the function. Though simplified, it reflects some important real world characteristics.

An estimated average cost of 1 hour of downtime is the following [27]:

$$\begin{array}{l} \quad \text{Empl. costs/hour} \cdot \text{\% Empl's affected by outage} \\ \underline{+ \text{ Avg. Rev./hour} \cdot \text{\% Rev. affected by outage} \qquad\quad} \\ = \text{Estimated average cost of 1 hour of downtime} \end{array} \tag{3}$$

If this cost is multiplied with the number of hours per operating year (e.g. 365 days $\cdot$ 24 hours for 24/7 systems) a maximum potential loss L is found. With availability A, the annual loss is $(1 - A)L$, e.g. $A = 95\%$ entails a loss of $0.05L$. In this simplified model, hourly cost is independent of outage duration.

By adding downtime costs and investment costs a net cost function is found:

$$\text{Net cost} = (1 - f(A_0, c))L + c \tag{4}$$

This net cost function has a level of investment c^* that minimizes the cost:

$$c^* = \frac{\ln(\alpha \cdot L \cdot (1 - A_0))}{\alpha} \tag{5}$$

3.2 The variance of outage costs

A better model does away with averages and lets the outage cost depend on the time of occurrence, giving each hour a separate random cost variable L_i. The expected total cost becomes a sum over the set Out of hours when outages occur:

$$\text{Net cost} = (1 - f(A_0, c)) \sum_{i \in Out} E[L_i] + c \tag{6}$$

In the stochastic model, net cost variance becomes important. As shown in [15], the variance depends a lot on whether the outage hours are consecutive or non-consecutive, assuming that the covariance of consecutive hours is larger than that of non-consecutive. In practice, this is often the case: two consecutive outage hours in a retail business before Christmas probably have a greater covariance than one hour from before Christmas and one hour from a February Monday morning. Thus, the *number of outages* becomes important for the variance of downtime costs. In our model, this is modeled by a homogeneous Poisson process (HPP). The probability that a failure occurs n times in the time interval $[0, t]$ is

$$P(N(t) = n) = \frac{(\lambda t)^n}{n!} e^{-\lambda t} \text{ for } n \in \mathbb{N} \tag{7}$$

$N(t)$ belongs to the Poisson distribution: $N(t) \in Po(\lambda t)$. λt is the expected number of outages in a year: the product of λ, the intensity of the HPP [occurrences/time] and t, the length of the time interval.

3.3 An actual dataset of revenue data

The final component of the model is a dataset based on [28], a report from the Swedish Retail Institute, with statistics on the revenue distribution in the Swedish retail sector. Hourly and monthly data is given in Tables 1 and 2. Based on these statistics, a dataset of 13 hours times 365 days was generated and normalized, reflecting relative hourly revenues over the operating year.

Table 1. Hourly retail sector revenue distributions (normalized) for normal and pay weeks [28].

	Mo.	Tu.	We.	Th.	Fr.	Sa.	Su.	Sum
Normal week (Pay week)								
09.00-12.00	1 (2)	2 (2)	2 (2)	3 (3)	4 (4)	4 (3)	2 (2)	18 (18)
12.00-16.00	3 (3)	3 (3)	4 (3)	5 (5)	8 (8)	9 (9)	6 (6)	37 (37)
16.00-22.00	5 (5)	5 (5)	7 (6)	10 (11)	10 (11)	4 (5)	3 (4)	44 (47)
Total	**10 (10)**	**10 (10)**	**12 (11)**	**18 (19)**	**22 (23)**	**17 (17)**	**11 (12)**	**99 (100)**

Table 2. Monthly retail sector revenue distribution (normalized) over a year [28].

Jan.	Feb.	Mar.	Apr.	May.	Jun.	Jul.	Aug.	Sep.	Oct.	Nov.	Dec.
7	7	8	8	8	8	8	8	8	8	8	11

As seen in the tables, there is a lot of variance: A payment system outage during a single high revenue hour might cost as much as a dozen low revenue hour outages, if no transactions can be made with the payment system down. In the experiment, downtime costs are calculated by Eq. (6), substituting expectations with hourly costs from the normalized dataset described in Tables 1 and 2. To summarize, the subjects thus face two important features of availability SLA decision-making: diminishing marginal returns on investment, and variance of outage costs. While the data might not be representative of all industries (cf. [15] for a further discussion), its variance offers an interesting case. The problem is easy to understand, but the stochastic model makes it hard to solve.

4 Data collection method

To empirically investigate the preferences of IT professionals with regard to availability SLAs, an experimental economics approach was used with 16 research subjects. All of the subjects work in the intersection of business and IT, some with a focus on availability. The subjects participated in an evening course on Enterprise Architecture for practitioners in the field of enterprise IT. Based on their background, the subjects are more likely to be on the procuring than the providing side of an SLA, though this was not explicitly investigated.

First, the subjects were introduced to the problem presented in Section 3. Diminishing returns on investments were introduced to the subjects using textbook diagrams [29, 11] and a table [30]. The importance of variance was illustrated with the following thought-provoking wording on a PowerPoint slide: "99.9% availability 24-7 means almost 9 hours of annual downtime. Case 1: A single 9 hour outage. Case 2: 100 separate 5 minute outages. Which one do you prefer?" Then, Tables 1 and 2 – background facts in the experiment – were shown and remained on display throughout the session. It was explicitly pointed out that the decision-making problem is a simplified one, not aiming to capture the entire complexity of real systems and their availability, but rather to investigate the behavior of IT decision-makers under uncertainty.

The subjects were asked 18 questions, with no pre-test. Each question represented a binary choice between two SLA scenarios. Each scenario had an SLA price, a resulting minimum availability (percent) (by Eq. (2)) and a number of expected outages (by Eq. (7)). The subjects were asked to procure the payment service for a retail store with revenue streams/downtime costs according to Tables 1 and 2 and Eq. (6). The translated questionnaire can be found in Fig. 1.

Each subject received an initial endowment of 300 Swedish kronor (SEK) for each of the 18 questions. The subjects received the information that one percent unavailability would correspond to 47.45 hours of annual downtime with *average* cost of 100 SEK. As a motivation for making wise decisions, following the data collection an answer from each subject was selected at random, its outcome simulated according to Section 3 and the resulting amount was paid out. The subjects did not have calculators. All data was fully anonymized before analysis.

The 18 questions were grouped into 3 categories. In the first category the questions were phrased as follows:

1. Do you prefer to pay 6 SEK for 99% availability with 1 expected outage or 12 SEK for 99 % availability with 2 expected outages?
2. Do you prefer to pay 12 SEK for 99% availability with 2 expected outages or 18 SEK for 99 % availability with 3 expected outages?

i.e. both alternatives always offered 99% availability, but the first alternative was cheaper with fewer outages. This pattern was followed until:

Consider the 18 questions below. Every question is a choice between two alternatives.

For every question your initial capital is 300 SEK that you should invest in order to get the most out of. One question will be selected randomly. simulated and paid out!

One percent unavailability corresponds to 47.45 hours downtime. 47.45 hours downtime cost 100 SEK *on average*.

1. ☐	Pay 6 SEK for 99% availability with 1 expected outage	☐ Pay 12 SEK for 99% availability with 2 expected outages	7. ☐	Pay 0 SEK for 99% availability with 2 expected outages	☐ Pay 15 SEK for 99.53% availability with 2 expected outages	13. ☐	Pay 0 SEK for 99% availability with 20 expected outages	☐ Pay 15 SEK for 99.53% availability with 20 expected outages
2. ☐	Pay 12 SEK for 99% availability with 2 expected outages	☐ Pay 18 SEK for 99% availability with 3 expected outages	8. ☐	Pay 15 SEK for 99.53% availability with 2 expected outages	☐ Pay 30 SEK for 99.78% availability with 2 expected outages	14. ☐	Pay 15 SEK for 99.53% availability with 20 expected outages	☐ Pay 30 SEK for 99.78% availability with 20 expected outages
3. ☐	Pay 18 SEK for 99% availability with 3 expected outages	☐ Pay 30 SEK for 99% availability with 5 expected outages	9. ☐	Pay 30 SEK for 99.78% availability with 2 expected outages	☐ Pay 45 SEK for 99.89% availability with 2 expected outages	15. ☐	Pay 30 SEK for 99.78% availability with 20 expected outages	☐ Pay 45 SEK for 99.89% availability with 20 expected outages
4. ☐	Pay 30 SEK for 99% availability with 5 expected outages	☐ Pay 48 SEK for 99% availability with 8 expected outages	10. ☐	Pay 45 SEK for 99.89% availability with 2 expected outages	☐ Pay 60 SEK for 99.95% availability with 2 expected outages	16. ☐	Pay 45 SEK for 99.89% availability with 20 expected outages	☐ Pay 60 SEK for 99.95% availability with 20 expected outages
5. ☐	Pay 48 SEK for 99% availability with 8 expected outages	☐ Pay 60 SEK for 99% availability with 10 expected outages	11. ☐	Pay 60 SEK for 99.95% availability with 2 expected outages	☐ Pay 75 SEK for 99.98% availability with 2 expected outages	17. ☐	Pay 60 SEK for 99.95% availability with 20 expected outages	☐ Pay 75 SEK for 99.98% availability with 20 expected outages
6. ☐	Pay 60 SEK for 99% availability with 10 expected outages	☐ Pay 120 SEK for 99% availability with 20 expected outages	12. ☐	Pay 75 SEK for 99.98% availability with 2 expected outages	☐ Pay 90 SEK for 99.99% availability with 2 expected outages	18. ☐	Pay 75 SEK for 99.98% availability with 20 expected outages	☐ Pay 90 SEK for 99.99% availability with 20 expected outages

Fig. 1. The questionnaire used (translation).

6. Do you prefer to pay 60 SEK for 99% availability with 10 expected outages or 120 SEK for 99 % availability with 20 expected outages?

In the second category the questions were phrased as follows:

7. Do you prefer to pay 0 SEK for 99% availability with 2 expected outages or 15 SEK for 99.53 % availability with 2 expected outages?

In this case the number of outages was always 2, but the first alternative was cheaper with a lower availability. This pattern was again followed until:

12. Do you prefer to pay 75 SEK for 99.98% availability with 2 expected outages or 90 SEK for 99.99 % availability with 2 expected outages?

In the third and final category the questions were phrased as follows:

13. Do you prefer to pay 0 SEK for 99% availability with 20 expected outages or 15 SEK for 99.53 % availability with 20 expected outages?

The number of outages was 20, but the first alternative was cheaper with a lower availability. The availability numbers were the ones of the second category.

The subjects were allowed as much time as they needed in order to complete the questionnaire. The authors were available to answer questions related to the subjects' understanding of the questions.

5 Results

A visual guide to the different behaviors described below is offered in Fig. 2.

5.1 Category 1

Expected behavior The *expected* reward is the same in all alternatives: 1% expected unavailability means an expected loss of 100 SEK. Thus, a decision-maker that maximizes expected utility would always chose the cheapest alternative, i.e. never be willing to pay to spread unavailability over a greater number of outages. Such maximization of expected utility would be consistent with the findings of [24] and [25]. However, because of the large variance in the outage costs, a more risk-averse decision-maker would be willing to pay to reach a certain number of outages, determined by her level of risk aversion. Once that number is reached, she would not be willing to pay more for an even greater number of outages. Thus, there would be a unique turning-point, below which a risk-averse decision-maker would pay for more outages, and above which she would not pay for more outages. The expected utility maximizer and the risk-averse agent are the two types of decision-makers discussed in [15].

Observed behavior 7 participants (44%) maximized the utility by always choosing the cheapest alternative. 5 participants (31%) behaved as risk-averse decision-makers and exhibited turning-points. One participant had a turning-point at 3 outages, two at five outages, one at eight outages and one at ten outages. 4 participants (25%) exhibited *non-monotonic* preferences in the sense that they, at some point, were not willing to pay to go from n to $n+m$ outages, but were willing to pay to go from $n+m$ to $n+k$ outages, where $k > m$.

5.2 Categories 2 and 3

Expected behavior The *expected* reward changes with the alternatives: Each basis point (i.e. one hundredth of a percentage point) of expected unavailability has an expected cost of 1 SEK. Thus, a decision-maker that maximizes expected utility would always pay for increased availability at a rate of more than 1 basis point per SEK, and never pay for increased availability at a rate of less than 1 basis point per SEK. In the given case, the utility-maximizer would pay 30 SEK to reach 99.78%, but not 45 SEK to reach 99.89%. However, a moderately risk-averse decision-maker might forgo this principle in the category 2 questions (where two expected outages make for large variance), but not in the category 3 questions (where twenty expected outages make for small variance).

Observed behavior 1 participant (6%) behaved as a consistent utility maximizer, with a turning-point at 30 SEK in both cases. 1 participant (6%) behaved as a risk-averse utility maximizer, with a turning-point at 45 SEK in category 2 and 30 SEK in category 3. 3 participants (19%) behaved as flawed but consistent utility maximizers, with equal but non 30 SEK turning points in both cases. 4 participants (25%) exhibited extreme behavior (not illustrated in Fig. 2) in always choosing to pay for more availability (2 participants) or never choosing to pay for more availability (2 participants). 1 participant (6%) exhibited non-monotonic preferences in both categories 2 and 3. 1 participant (6%) behaved as a utility maximizer (30 SEK turning-point) in category 2, but was extreme in category 3 by always choosing to pay for more availability. 1 participant (6%)

behaved as a flawed utility maximizer (15 SEK turning-point) in category 2, but exhibited non-monotonic preferences in category 3. 2 participants (11%) behaved as risk-averse utility maximizers in category 2 (turning-points at 45 SEK), but exhibited non-monotonic preferences in category 3. 2 participants (11%) behaved as risk-averse utility maximizers in category 2 (turning-points at 45 and 75 SEK), but were extreme in category 3 (one always choosing to pay for more availability, one never choosing to pay for more availability).

The results are summarized in Table 3. The payments, following random selection and simulations, ranged from a maximum of 261 SEK to a minimum of 123 SEK, with a median of 236 SEK and an mean of 216 SEK.

Table 3. A summary of the results. EUM = expected utility maximizer (i.e. no risk aversion of risk seeking), FUM = flawed utility maximizer (i.e. a non-optimal turning point), RUM = risk averse utility maximizer (i.e. paying more than a strict expected utility maximizer to decrease variance), Non-mon = non-monotonic preferences (i.e. multiple turning points), Extreme = always choosing to pay for more availability or never choosing to pay for more availability (not illustrated in Fig. 2).

Participant	Category 1	Category 2	Category 3
1	EUM	FUM	Non-mon
2	Non-mon	EUM	Extreme
3	Non-mon	Non-mon	Non-mon
4	EUM	Extreme	Extreme
5	Non-mon	Extreme	Extreme
6	RUM	Extreme	Extreme
7	EUM	Extreme	Extreme
8	EUM	FUM	FUM
9	Non-mon	FUM	FUM
10	EUM	FUM	FUM
11	RUM	RUM	EUM
12	EUM	EUM	EUM
13	RUM	RUM	Extreme
14	RUM	RUM	Extreme
15	EUM	RUM	Non-mon
16	RUM	RUM	Non-mon

6 Analysis

The experimental evidence is somewhat surprising, as few participants behave as (risk-averse) expected utility maximizers, whereas many exhibit non-monotonic or extreme preferences. Very few individuals were consistent with the ideal (risk-averse) utility maximizers hypothesized before the experiment; in essence only participants 11 and 12. The behavior of the respondents is different from those presented in [24, 25], especially the high amount of inconsistent respondents.

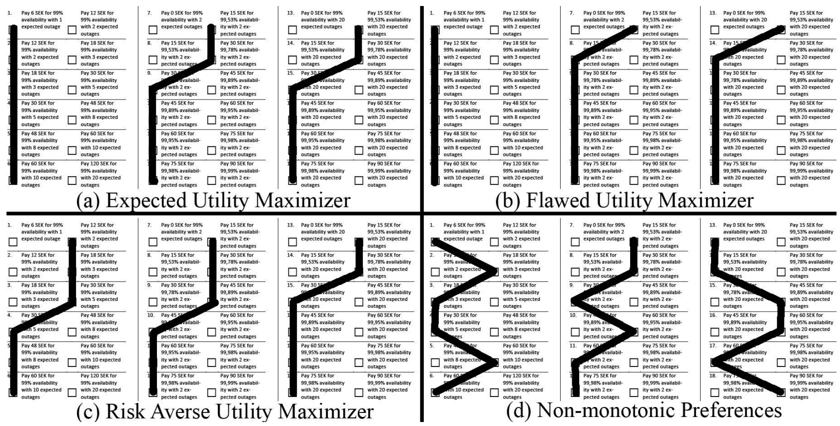

Fig. 2. Categories of behaviors found. This figure provides a visual guide to illustrate the different behaviors – the text of the questionnaire is more legible in Fig. 1.

As noted above, previous work has identified many deviations expected utility maximization. However, all the subjects were professionals, managing enterprise IT in their line of work. Therefore, their deviations from expected utility in decisions relating to SLAs are interesting. It might be the case that the incentives were simply too low to properly motivate the subjects (unfortunately, it is prohibitively costly to use realistically large incentives), especially as the subjects did not have calculators (though for all 18 questions, a mere three calculations suffice to find the appropriate turning-point in each category). Indeed, companies might always do their math properly in real situations with higher stakes. However, at least some practitioners self-report that their companies are immature in SLA writing [31], and thus might not be much better utility-maximizers than the individual decision-makers in the experiment. Furthermore, knowledge gaps can exist even between knowledgeable procurers and service providers, affecting SLA quality [32]. In light of the large deviations from expected utility, it would be interesting to redo the experiment on a larger student sample to see whether professional experience matters.

It is worth elaborating on two reasons why the expected utility is appropriate for the enterprise IT service SLA setting. First, expected (monetary) utility is appropriate because of the corporate context. Second, SLA decision-making does not aim to replicate decisions or decision-making principles of any actual individuals. It aims to do what is in the best interest of the enterprise. To this end, it is often a distributed process, in the sense that someone investigates business-side requirements, someone maps dependencies between IT services, someone does ROI calculations, and someone negotiates with service-providers, before someone (nominally) finally makes the decision and signs the contract.

This distributed nature of decision-making is both a strength and a weakness of the experimental setup. The strength is that all of the participants were relevant, in the sense that even though they professionally belong to different parts of the decision-making chain; they all have a role in it. The weakness is that decisions are rarely taken by a single individual. Still, a more complex (collaborative) experimental setup might have unnecessarily clouded the results.

The small number of participants ($N = 16$) clearly deserves a remark, as it limits the reliability. Follow-up experiments, with a larger number of participants, would obviously be desirable. However, it should be noted that this weakness of reliability is related to a strength of validity: all of the participants were actual IT management professionals, lending the result a greater credibility. Validity is further increased by the realistic data-set (re-used from [15]).

An improvement of the questionnaire would be to include baseline questions on binary choices between a sure thing (e.g. 100 SEK) and a lottery (e.g. lottery 1: 200 SEK with 40% probability, 0 SEK with 60% probability or lottery 2: 200 SEK with 60% probability, 0 SEK with 40% probability). This would clarify each subject's tendency to maximize expected utility or to avoid risk. A question related to professional experience would also have been interesting.

7 Conclusions

This paper presents an investigation of availability SLA decision-making with subjects from the IT management profession. The scenario required the subjects to make pairwise choices between alternatives, under uncertainty. Subjects were incentivized by a payment based on one of their choices, randomly selected.

The results indicate that decision-makers rarely maximize expected utility. Some previous work indicated that they would, whereas there are also many deviations identified in the literature. The implications for company SLA management in real situations require more research. The sample size ($N = 16$) is small and reliability thus moderate, whereas validity is high due to the background of the participants and the realistic data-set (re-used from [15]) used in the payment simulations.

In addition to re-doing our experiment with a larger number of participants, an interesting direction for future work is to investigate whether decision-support systems of various kinds could help improve SLA decision-making. Another interesting approach for future experiments would be to have research subjects act both as IT service providers and procurers, playing out a negotiation scenario. It would also be interesting to investigate the impact of varying years of experiences; how do experienced professionals compare with their less experienced colleagues, or with inexperienced students?

Acknowledgment

The authors thank Shahid Hussain, who acted as the discussant of this paper at a research seminar at KTH. Furthermore, the three anonymous reviewers offered a number of comments that have substantially improved the quality of the paper.

Bibliography

[1] Marston, S., Li, Z., Bandyopadhyay, S., Zhang, J., Ghalsasi, A.: Cloud computing – the business perspective. Decision Support Systems **51**(1) (2011) 176–189

[2] IBM Global Services: Improving systems availability. Technical report, IBM Global Services (1998)

[3] Bharadwaj, A., Keil, M., Mähring, M.: Effects of information technology failures on the market value of firms. The Journal of Strategic Information Systems **18**(2) (2009) 66–79

[4] Baek, S.I., Lee, S.J., Lim, G.G.: Exploring impacts of IT service failure on firm's market value. In: Networked Computing and Advanced Information Management, 2008. NCM '08. Fourth International Conference on. Volume 2. (September 2008) 450 –454

[5] Harris, K.: Gartner CEO and Senior Business Executive Survey, 2010: Perceptions of IT and Tactical Fixes. Technical report, Gartner, Inc. (March 2010)

[6] Salle, M., Bartolini, C.: Management by contract. In: Network Operations and Management Symposium, 2004. NOMS 2004. IEEE/IFIP. Volume 1., IEEE (2004) 787–800

[7] Sauvé, J., Marques, F., Moura, A., Sampaio, M., Jornada, J., Radziuk, E.: SLA design from a business perspective. In: Ambient Networks. Springer (2005) 72–83

[8] Buco, M.J., Chang, R.N., Luan, L.Z., Ward, C., Wolf, J.L., Yu, P.S.: Utility computing SLA management based upon business objectives. IBM Systems Journal **43**(1) (2004) 159–178

[9] Hasija, S., Pinker, E.J., Shumsky, R.A.: Call center outsourcing contracts under information asymmetry. Management Science **54**(4) (2008) 793–807

[10] Williams, D.: The Challenges and Approaches of Establishing IT Infrastructure Monitoring SLAs in IT Operations. Technical report, Gartner, Inc. (September 2010)

[11] Taylor, S., Lloyd, V., Rudd, C.: Service Design (ITIL). The Stationery Office, TSO (2007)

[12] Camerer, C.: Bounded rationality in individual decision making. Experimental Economics **1**(2) (1998) 163–183

[13] Goo, J., Kishore, R., Rao, H.R., Nam, K.: The role of service level agreements in relational management of information technology outsourcing: An empirical study. MIS Quarterly **33**(1) (2009) 119–145

[14] Vetter, J., Benlian, A., Hess, T.: Overconfidence in IT investment decisions: why knowledge can be a boon and bane at the same time. In: Proc. Thirty Second International Conference on Information Systems, ICIS. (2011)

[15] Franke, U.: Optimal IT Service Availability: Shorter Outages, or Fewer? Network and Service Management, IEEE Transactions on **9**(1) (2012) 22–33

[16] Setzer, T., Bhattacharya, K., Ludwig, H.: Change scheduling based on business impact analysis of change-related risk. Network and Service Management, IEEE Transactions on **7**(1) (March 2010) 58 –71
[17] Courcoubetis, C., Siris, V.: Managing and pricing service level agreements for differentiated services. In: Quality of Service, 1999. IWQoS '99. 1999 Seventh International Workshop on. (1999) 165 –173
[18] Trienekens, J.J., Bouman, J.J., van der Zwan, M.: Specification of service level agreements: Problems, principles and practices. Software Quality Journal **12** (2004) 43–57
[19] Kieninger, A., Schmitz, B., Straeten, D., Satzger, G.: Incorporating business impact into service offers–a procedure to select cost-optimal service contracts. In: AMCIS 2012 Proceedings. (2012) Paper 19.
[20] Demirkan, H., Goul, M., Soper, D.S.: Service level agreement negotiation: A theory-based exploratory study as a starting point for identifying negotiation support system requirements. Hawaii International Conference on System Sciences **1** (2005) 37b
[21] Brandic, I., Music, D., Leitner, P., Dustdar, S.: VieSLAF framework: Enabling adaptive and versatile SLA-management. In: Grid Economics and Business Models. Springer (2009) 60–73
[22] Silaghi, G.C., Şerban, L.D., Litan, C.M.: A framework for building intelligent SLA negotiation strategies under time constraints. In: Economics of Grids, Clouds, Systems, and Services. Springer (2010) 48–61
[23] Figueroa, C., Figueroa, N., Jofre, A., Sahai, A., Chen, Y., Iyer, S.: A game theoretic framework for SLA negotiation. Technical report, HP Laboratories (2008)
[24] Davenport, W., Middleton, M.: Expectation theories of decision making for duplex gambles. Acta Psychologica **37**(3) (1973) 155 – 172
[25] Lindman, H., Lyons, J.: Stimulus complexity and choice inconsistency among gambles. Organizational Behavior and Human Performance **21**(2) (1978) 146 – 159
[26] Høyland, A., Rausand, M.: System reliability theory: models and statistical methods. Wiley New York (1994)
[27] Patterson, D.: A simple way to estimate the cost of downtime. In: Proc. 16th Systems Administration Conf.— LISA. (2002) 185–8
[28] Bergström, F., Arnberg, J.: När handlar vi? – Om konsumtionsmönster i den traditionella detaljhandeln. Technical report, Handelns Utredningsinstitut (March 2005) In Swedish.
[29] Marcus, E., Stern, H.: Blueprints for high availability, second edition. John Wiley & Sons, Inc., Indianapolis, IN, USA (2003)
[30] Malik, B., Scott, D.: How to Calculate the Cost of Continuously Available IT Services. Technical report, Gartner, Inc. (November 2010)
[31] Franke, U., Johnson, P., König, J.: An architecture framework for enterprise IT service availability analysis. Software & Systems Modeling (2013) 1–29
[32] Köppel, A., Böning, D., Abeck, S.: How to support the negotiation of service level agreements (SLAs) for your client/server application. In: Proc. SCI'99 and ISAS'99. (1999)