

# IT service outage cost: Case study and implications for cyber insurance

Ulrik Franke  
RISE Research Institutes of Sweden  
P.O. Box 1263, SE-164 29 Kista, Sweden  
`ulrik.franke@ri.se`

Suggested running title:  
IT service outage cost

# IT service outage cost: Case study and implications for cyber insurance

Ulrik Franke

RISE Research Institutes of Sweden  
P.O. Box 1263, SE-164 29 Kista, Sweden  
`ulrik.franke@ri.se`

## Abstract

Today, almost all enterprises are highly dependent on IT services. Thus, high availability IT services and the cost of downtime have received a lot of attention in recent years. One increasingly used tool for cyber risk management and transfer is cyber insurance, which typically offers some form of business interruption coverage. However, cost structures of IT service outages are still poorly understood, as costs are often just reported as lump sums. This article contributes a multiple case study of IT service outage cost in three sectors in Sweden: transport companies ( $N = 11$ ), food companies ( $N = 9$ ), and government agencies ( $N = 19$ ). The contribution is three-fold: (i) the measurement instrument itself, (ii) the insights into different cost structures gained, and (iii) the implications of different cost structures on availability investment strategies. Main findings include that whereas some enterprises incur only fixed outage cost, some incur (almost) only lost productivity, and some incur almost only lost revenue. In the public sector, lost revenue is often negligible. The results are further contextualized by a discussion of cyber insurance implications.

**Keywords**— IT service availability, outage cost, cyber insurance, business interruption, availability investment, fixed and variable cost

## 1 Introduction

In modern society, individuals, private companies and public agencies all depend on functioning and available IT services in their daily lives and operations. For business, outages can have dramatic consequences, including lost revenue, lost productivity and damage to reputation (Lerner et al., 2016). Though the latter category of damage can be difficult to quantify, it has been demonstrated that IT failures result in abnormal drops in stock prices, by 2% on average, in the days following the event (Bharadwaj et al., 2009).

To avoid such consequences, enterprises pose demanding requirements on IT service continuity. The consultancy Gartner reports of “aggressive” targets such as recovery times of at most four hours and availability levels of at least 99.5% (Morency, 2014). However, in many cases even these targets are not sufficient, and it has been argued that enterprises typically require four ‘nines’ IT service availability, i.e., 99.99%, or more (Durkee, 2010). For a service running 24 hours a day, 365 days a year, this corresponds to slightly below a single hour of allowed annual downtime.

However, despite technical efforts, outages can never be avoided altogether. This is one of the rationales for writing cyber insurance policies that cover *business interruption* and related incident management costs, both first party coverage of internal IT service failures and third party coverage of external

failures (OECD 2017, pp. 33–38; Biener et al. 2015). Oftentimes, academic and practitioner discussions of cyber insurance instead focus on *data breach*, possibly due to the availability of the Privacy Rights Clearinghouse curated data set on breaches, which has been used in several papers (e.g. Carfora et al. 2019; Tonn et al. 2019; Eling and Loperfido 2017; Edwards et al. 2015). However, business interruption and data breach are both important in their own right. Neither should overshadow the other. Depending on the circumstances, the one or the other could be the greatest risk. For example, Lloyd’s recently explored a scenario similar to the NotPetya and WannaCry cases, where malware rapidly spreads and causes IT service outages across the globe (Daffron et al., 2019). In all the variations of this scenario, business interruption constitutes the greatest insured loss, several times those of incident response or liabilities (Daffron et al., 2019, Table 9, p. 41).

This is the background to the overall research question of this article: *How can the cost of IT service outages be measured, investigated, and insured?* Of course, IT service outage cost has been studied before. Hourly cost of downtime have regularly been estimated to be in the range of hundreds of thousands or even millions of US dollars (Rapoza, 2014; IBM Global Services, 1998; Ponemon, 2016), at least for large companies in certain lines of industry. However, such reports often stop short at lump numbers. This article, by contrast, aims to contribute insight into how outage *cost structures* differ and to demonstrate the relevance of these differences both to the enterprises and their insurers. Briefly stated, some cost components are insurable, while others are not, and different cost structures imply different managerial strategies. Thus, knowledge about cost structure can be crucial. The vehicle of investigation is a multiple case-study, where a questionnaire instrument to measure outage cost has been applied at 39 Swedish enterprises from various sectors in society. The article makes three main contributions: (i) the measurement instrument itself, (ii) empirical evidence on how IT service outage costs differ between enterprises and sectors, and (iii) the implications of different cost structures on availability investment strategies, as modeled mathematically.

IT service outage cost at Swedish enterprises may seem like a provincial concern, but the results are interesting to a wider audience for two reasons. First, Sweden is often near the top when countries are ranked by digital maturity. For example, Sweden was ranked 2nd in the EU Digital Economy & Society Index 2019 (European Commission, 2019) and 3rd in the World Economic Forum’s Networked Readiness Index 2016 (World Economic Forum, 2016). Therefore, insights from Sweden, as a forerunner, might offer valuable insights also for other countries. Second, Sweden has a rapidly developing cyber insurance market (Insurance Sweden, 2019).

More precisely, the following research questions are addressed:

1. How should a measurement instrument suitable for investigating IT service outage cost look? This has implications for the forms used in cyber insurance underwriting (Woods et al., 2017), as well as claims adjustment.
2. How do IT service outage cost structures vary across enterprises and sectors? This has implications for how insurers and re-insurers should manage accumulation risk, arguably the greatest concern about cyber insurance (Hofmann et al., 2018), in their portfolios.
3. How does IT service outage cost structure affect availability investment strategies? This has implications for how insurers can proactively nudge their customers towards more security, a question at the heart of recent policy interest in cyber insurance (ENISA, 2016; OECD, 2017; World Economic Forum, 2018).

The remainder of the paper is structured as follows: Section 2 puts the contribution in context by discussing related work. Section 3 explains the data collection method used in the study. The results are first described in Section 4, then discussed and contextualized using a mathematical investment allocation problem in Section 5. Validity and reliability are addressed in Section 6, which also spells out implications for cyber insurance, before Section 7 concludes the paper, and offers some directions for future work.

## 2 Related work

Given the importance of available IT services, it is not surprising that there is much work on availability modeling, including cost modeling. Recent examples from the technical literature include important topics such as cloud infrastructure design (Sousa et al., 2017), IT service design (Bosse et al., 2016), and disaster recovery strategy evaluation (Andrade et al., 2017; Nguyen et al., 2016). However, these contributions typically take the cost of outages as exogenously given starting points, readily available to be plugged into mathematical frameworks to find solutions to availability engineering problems. The same is true of the abundant recent literature on optimization of IT service availability, addressing topics such as optimal placement of virtual machines (Jammal et al., 2017), optimal reconfiguration decisions (Logeswaran et al., 2017), and optimal backup (Li et al., 2017) in clouds in order to minimize downtime. These all differ from the objective of this paper, which is to study the cost, and its structure, itself.

Unsurprisingly, there are also treatments of outage cost in the cyber insurance literature. Franke (2017) observes that business interruption coverage for attacks is included in all insurance policies on offer in Sweden, whereas non-malicious interruptions are treated differently by different insurers. Similarly, an OECD survey found that standard business interruption coverage was included in almost all policies investigated (OECD, 2017, p. 62). In an investigation of US and UK cyber insurance proposal forms, Woods et al. (2017) found that out of 24 forms studied, 15 evaluate the business continuity plan of the prospective client. Based on mandatory regulatory filings of US insurers, Romanosky et al. (2019) offer an overview of cyber insurance policy contents. Findings indicate that business interruption is priced by insurers as an additional percentage of a base-rate premium. With respect to service provider failure BI, Romanosky et al. also document that many proposal forms ask insureds to list which services have been out-sourced to which providers, a practice that is also in line with previous findings from the Swedish market Franke (2017). Based on the same mandatory US filings, Woods et al. (2019) study the pricing information in the filings of 26 US insurers more in-depth. Out of the 26 insurers studied, 15 offered business interruption coverage, at an average premium of 0.36 times the corresponding premium for cyber liability (4 offered *contingent* business interruption coverage at 0.06 times the cyber liability premium). While all these studies offer insights on outage cost, such evidence is indirect—observed, as it were, through the lens of insurers. In this paper, observations are instead made through the organizations bearing the cost.

There are also a few related contributions in the empirical risk management and insurance literature. In this strand of research, data from financial sector operational risk databases or similar is often used. Examples include Goldstein et al. (2011) who analyze 25 years of IT operational failure events in US financial service firms, Rachev et al. (2006) who test alternatives to the Basel II rules on operational risk data, Franke et al. (2014) who do statistical analysis of IT service outage durations in a large Nordic bank, Ibrahimovic and Franke (2016) who use operational risk data from a Bosnian bank to propose improved risk management under Basel II, and Biener et al. (2015) who test the insurability of cyber risk using financial sector operational risk data. Biener et al. in particular connect this empirical investigation to cyber insurance, but also illustrate a shortcoming of the existing literature: the losses from operational risks are not analyzed in the context of, e.g., outage durations for business interruptions, number of records lost for data breaches, or similar contextual numbers. This makes it hard to connect knowledge about, e.g., outage duration with knowledge about cost entailed. This paper aims to contribute to this area, shedding more light on the structure of outage cost, beyond lump numbers.

Though, to the knowledge of the author, there are no peer reviewed empirical studies of IT service outage cost in Sweden, a limited study by the Swedish Social Insurance Inspectorate should be mentioned. Using a methodology similar to ours (cf. Section 3) the cost of lost productivity due to system failures and outages were estimated to 19 million SEK at the Swedish Social Insurance Agency and 1.7 million SEK at the Swedish Pensions Agency in 2012 (Sohlberg and Jansson, 2012). (10 SEK is roughly one euro or one US dollar.)

To conclude, while there is abundant literature on (i) IT service outages, (ii) the treatment of business interruption in cyber insurance policies, and (iii) operational risk and its cost, literature connecting these three fields by addressing cost structures is scarce. This paper makes a novel contribution precisely in this area, thus also responding to calls for more empirical research on cyber insurance in general Eling

and Schnell (2016) and cyber cost measurement in particular (Falco et al., 2019, Section 3).

### 3 Methodology

The empirical part of this investigation is a multiple case study, conducted using an electronic questionnaire to investigate the magnitude of IT service outage cost incurred by the respondents in one year. In all, 41 responses were received.

The study used *purposive sampling*, aiming to cover respondents from several contrasting sectors. Respondents were sought primarily from sectors identified as particularly important by the Swedish Civil Contingencies Agency, MSB. In the end, three sectors were represented in sufficient numbers to be included: transport companies ( $N = 11$ ), food companies ( $N = 9$ ), and government agencies ( $N = 19$ ). Two responses from the financial industry were also received, but are too few to constitute a proper case study, and are excluded from further analysis.

To put these sectors in context, there are some 6 000 Swedish road transport companies, employing some 150 000 people, with an annual revenue of some 160 GSEK.<sup>1</sup> In 2017, the Swedish food industry comprised some 4 600 companies, employing some 49 000 people, with an annual revenue of some 195 GSEK.<sup>2</sup> In 2017, there were some 340 Swedish government agencies employing some 220 000 people, with an annual revenue (2016) of 1 344 GSEK (Statskontoret, 2017).

Since the sampling is nonrandom, this investigation is not a statistical survey, but rather a multiple case study. The implications of this, in particular with regard to generalization, are further discussed in Section 6.

#### 3.1 Questionnaire construction: core questions

The questionnaire aims to estimate respondent cost of IT service outages. However, just asking for this cost *simpliciter* entails a substantial risk that different respondents will interpret the question in different ways. Therefore, it was decided to let respondents answer several simpler component-wise questions before these answers were combined into the grand total sought for. Respondents were also given the opportunity to interactively refine their grand totals by revising the component-wise answers until the result was judged reasonable. All outage costs are covered, regardless of cause. Thus, the questionnaire instrument is agnostic on distinctions such as whether, for instance, (i) outages are caused by internal or external IT service failures, or (ii) caused by attacks or non-malicious mistakes.

As a starting point, Patterson’s “simple way” to estimate the average cost of 1 hour of downtime was used (Patterson, 2002):

$$\begin{aligned} & \text{Employee costs/hour} \cdot \% \text{ Employees affected by outage} \\ + & \text{Average revenue/hour} \cdot \% \text{ Revenue affected by outage} \\ = & \text{Estimated average cost of 1 hour of downtime} \end{aligned} \tag{1}$$

In Eq. (1), the average employee costs per hour are multiplied with the fraction of employees affected by the outage, representing loss of productivity, i.e., unproductive working hours that are still paid for. For example, with employee costs per hour at 300 000 SEK (e.g., with 1 000 employees who cost 300 SEK an hour) and half of the employees are affected, the first line of Eq. (1) is 150 000 SEK. Second, the average revenue per hour is multiplied with the fraction of revenue affected by the outage, representing loss of revenue. For example, with an average hourly revenue of 1 000 000 SEK (e.g., with 100 stores selling goods for 10 000 SEK an hour) and 30 % of the point-of-sales systems being down so that sales cannot take place, the second line of Eq. (1) is 300 000 SEK. Given these figures, the estimated average cost of 1 hour of downtime is then the sum of lines one and two, or 450 000 SEK in the example. This is essentially the method used to study outage cost at Swedish government agencies (Sohlberg and Jansson, 2012), also pedagogically explained in a recent Gartner report (Vecchio, 2016).

<sup>1</sup><https://www.akeri.se/en/node/161>, accessed March 25, 2020.

<sup>2</sup><https://www.livsmedelsforetagen.se/in-english/>, accessed March 25, 2020.

However, Eq. (1) only considers the *variable* cost associated with an outage, i.e., the cost per hour. There could also be a *fixed* cost associated with an outage, i.e., a cost which is constant whether the outage lasts a second, a minute, or an hour. To see the importance of fixed costs, consider the difference between 30 interruptions of 2 minutes and a single interruption of 1 hour. In the *absence* of fixed cost, the difference between the two cases—continuous and aggregated downtime—disappears.

From a cyber insurance perspective, it is worth noting that such fixed first-party response cost is often covered in policies. In an investigation of cyber insurance policies, Meland et al. cite data showing that response costs (which are mostly fixed costs) are on average three times as great as loss of business income (which are mostly variable costs) (Meland et al., 2017).

Based on these considerations, the following core cost questions were included in the questionnaire, along with examples to make them easier to answer (numbering as in the final questionnaire, translation from Swedish, slightly abbreviated, italics in original):

#### Unplanned IT service downtime in 2016

5. Number of outages during the year:

*Example: 7*

6. Number of hours of unplanned IT service outages during the year:

*Example: The 7 outages sum to a grand total of 235 hours of downtime*

#### Cost to restore service per IT service outage (fixed cost)

7. Cost to restore service per IT service outage (average):

*Example: Overtime, crisis management, extra work, and consultancy bills sum to 10 000 SEK per outage*

#### Average cost of one hour of IT service downtime (variable cost)

#### Lost productivity

8. Hourly cost of an employee including payroll tax, excluding overhead costs (average):

*Example: 250 SEK*

9. Number (equivalents) of employees affected by an IT service outage (average):

*Example: 5 employees who can do nothing gives the answer 5. 10 employees who can work with half productivity gives the equivalent 5. 5 employees who can do nothing and 2 employees who leave their regular tasks gives the answer 7.*

#### Lost revenue

10. Revenue per hour (average):

*Example: 100 stores selling goods for 10 000 SEK an hour gives 1 000 000 SEK*

11. Fraction of revenue affected by an IT service outage (average):

*Example: If 100 stores are all of the same size and a single store has a payment system outage, that gives the answer 1%*

Once these questions were answered by the respondent, the electronic questionnaire automatically calculated the estimated total cost of IT service outages in 2016, and displayed the derivation explicitly, for example as follows:

12.  $7 \text{ outages} \cdot 10\,000 \text{ SEK per outage} + 235 \text{ hours} \cdot (250 \text{ SEK per employee and hour} \cdot 5 \text{ employees} + 1\,000\,000 \text{ SEK per hour} \cdot 1\% \text{ of the revenue}) = 2\,713\,750 \text{ SEK}$

The respondent was then given the choice either to confirm that the estimate was reasonable, or go back and revise the answers to the cost component questions. Once the estimate was thus established, its representativeness was investigated by an explicit question:

13. Is the 2016 cost representative of previous years and assessments of future years? If not, in what way does the 2016 cost differ?

- The 2016 cost was *greater* than previous years and assessments of future years
- The 2016 cost was *less* than previous years and assessments of future years

Respondents who answered that the 2016 cost was *not* representative could also give additional free text comments on representativeness. Concluding the core questions, all respondents were also asked about additional impact:

14. Additional impact on the enterprise not reflected in the estimated figure

*Example: The reputation of the company among customers and stock owners has been damaged, patients have had to wait for care, citizens have had to wait for service, employees have been stressed etc.*

(The examples in question 14 were tailored to respondents so that, e.g., private companies were given the example of customers and stock owners, whereas government agencies were given the example of citizens waiting for service.)

### 3.2 Questionnaire construction: additional questions

The questionnaire also contained some demographics on size and IT dependency (questions 1–4) *before* the core questions, and some additional questions (questions 15–19) *after* the core questions. When asked about cyber insurance (questions 15–17), it turned out that no respondent had such a policy, which is not surprising given the small uptake on the Swedish market at the time (Franke, 2017), though this has changed dramatically since then (Insurance Sweden, 2019). Some questions on Service Level Agreements (SLA) for IT service availability governance (questions 18–19) are separately treated by Olsson and Franke (2019). The full questionnaire, translated into English, is available as supplemental material.

### 3.3 Questionnaire validation

The questionnaire was iteratively developed with external feedback. As the original model (Patterson, 2002) fits commercial operations best, feedback on the suitability for the public sector was important. Discussions were therefore held first with experts at the Swedish Social Insurance Inspectorate and the Swedish National Audit Office, then with two Swedish municipalities.

### 3.4 Respondents and questionnaire distribution

Respondents were approached through three avenues of contact. First, the author visited 5 public-private fora for coordination between government agencies and private companies, who distributed the questionnaire to their members. Second, trade associations were contacted, 4 of which distributed the survey to their members (or a subset thereof). Third, 35 government agencies not represented in the public-private fora were approached.

E-mails with links to the questionnaire were distributed in February and March 2017 to respondents, who were typically given some 1-2 weeks to answer. As deadlines approached, reminders were sent, and opportunities to submit late answers were given.

For reasons of respondent confidentiality, in particular with regard to the trade association members whom the author only accessed indirectly through the respective association, the exact total number of invited respondents is not known.

## 4 Results

In the following, the results from the questionnaire are briefly described, first broken down into the three sector case studies, then with some general remarks on the bigger picture. As this is a multiple case study, the section focuses on qualitative descriptions of the data obtained and its interpretation, rather

than on quantitative data. Throughout, the k ( $10^3$ ), M ( $10^6$ ) and G ( $10^9$ ) prefixes are used on cost figures. For convenience, a brief summary of typical results is also given in Table 1.

#### 4.1 Case 1: Transportation

The first case study concerns transport companies, 11 of which completed the questionnaire. A short characteristic of their responses follows.

**Employees:** Typically below 100; a few hundred in a single case.

**Revenues:** In the order of tens or hundreds of MSEK.

**IT dependency:** Typically high (8-10), a few medium (5).

**IT service outage cost 2016:** Typically ranging from tens to hundreds of kSEK; a few MSEK in a single case, zero cost in a single case.

**Number of IT service outages 2016:** Ranging from single digits to a dozen, or a score in a single case.

**Hours of unplanned IT service outages 2016:** Ranging from single hours to tens of hours; a few hundred hours in a single case.

**Representativeness:** All except for the company with zero losses report that the 2016 cost figures are representative of previous years and assessments of future years.

**Additional impacts:** Overtime and stress among employees, poor route planning for trucks, inability to fulfill obligations to customers, having to switch to manual procedures, general anxiety about IT service outages, customers who cannot order services and thus may be tempted to use competitors instead.

It is worth pointing out that many of the additional impacts listed by respondents should actually be included in the cost estimate; in particular, this applies to overtime and customers switching to competitors.

#### 4.2 Case 2: Food

The second case study concerns food companies, 9 of which completed the questionnaire.

**Employees:** Typically one or two hundreds; a few below a hundred and a single case with a thousand.

**Revenues:** In the order of hundreds or thousands of MSEK, a single case below a hundred MSEK.

**IT dependency:** High (8-10).

**IT service outage cost 2016:** Typically ranging from tens to hundreds of kSEK; zero cost in a single case.

**Number of IT service outages 2016:** Typically single digits, ten in a single case.

**Hours of unplanned IT service outages 2016:** Ranging from single hours to tens of hours.

**Representativeness:** Three companies report higher cost in 2016, the rest report that the 2016 cost figures are representative of previous years and assessments of future years.

**Additional impacts:** Stress among employees who have to switch to manual procedures, annoyed customers who get decreased service levels, longer outages would dramatically affect customer experience (but the respondent has avoided this so far).

It is noteworthy that the single company with zero IT service outage cost in 2016 reports that this is representative.

### 4.3 Case 3: Government

The third case study concerns government agencies, 19 of which completed the questionnaire.

**Employees:** Typically hundreds or thousands, below 100 in a single case.

**Revenues:** Typically hundreds or thousands of MSEK; a few dozen GSEK in a single big case, a few dozen MSEK in a single small case.

**IT dependency:** High (7-10).

**IT service outage cost 2016:** Ranging from tens to hundreds of kSEK to single digit and tens of MSEK; zero cost in a single case.

**Number of IT service outages 2016:** Ranging from single digits to a few dozen, to a few hundred in a few cases.

**Hours of unplanned IT service outages 2016:** Ranging from single hours to tens of hours; a few hundred hours in a few cases and a thousand hours in a single case.

**Representativeness:** Four agencies report higher cost in 2016, the rest report that the 2016 cost figures are representative of previous years and assessments of future years.

**Additional impacts:** In some cases citizens and external partners have had to wait for information, but in other cases short outages are reported as not having affected citizens or anyone else outside the organization. Stress among employees, annoyed users, some customers choosing other alternatives.

Table 1: A summary of *typical* results from the three sector case studies. For more details, see the text. Costs, number and durations of outages are cumulative and annual, reflecting the sum total of the entire year. The  $\sim$  sign is used to denote ‘order of’, i.e.,  $\sim 1$  denotes single digits,  $\sim 10$  is roughly from 10 to 90, and  $\sim 10^3$  is roughly from 1 000 to 9 000.

	Transportation	Food	Government
Employees	$\sim 10$ ; a few $\sim 100$	$\sim 10$ to $\sim 100$	$\sim 100$ to $\sim 1000$
Revenues (SEK)	$\sim 10^7$ to $\sim 10^8$	$\sim 10^8$ to $\sim 10^9$	$\sim 10^8$ to $\sim 10^9$
Outage cost (SEK)	$\sim 10^4$ to $\sim 10^5$	$\sim 10^4$ to $\sim 10^5$	$\sim 10^4$ through $\sim 10^5$ and $\sim 10^6$ to $\sim 10^7$
Outages	$\sim 1$ to $\sim 10$	$\sim 1$	$\sim 1$ to $\sim 10$ ; $\sim 100$ in a few cases
Cumulative outages (h)	$\sim 1$ to $\sim 10$	$\sim 1$ to $\sim 10$	$\sim 1$ to $\sim 10$ ; $\sim 100$ in a few cases

### 4.4 IT service outage costs relative to revenues

It is also interesting to consider cost magnitudes relative to revenues, as a crude way of normalizing the outage cost with respect to the size of the enterprise. Fig. 1 displays these relative outage costs as box plots categorized by the three case studies, giving an overview of the data.

As seen, the vast majority of annual IT service outage costs represent a few basis points (i.e., hundredths of a percent) of revenues. All typical values are below fifty basis points (0.5%); only a few outliers

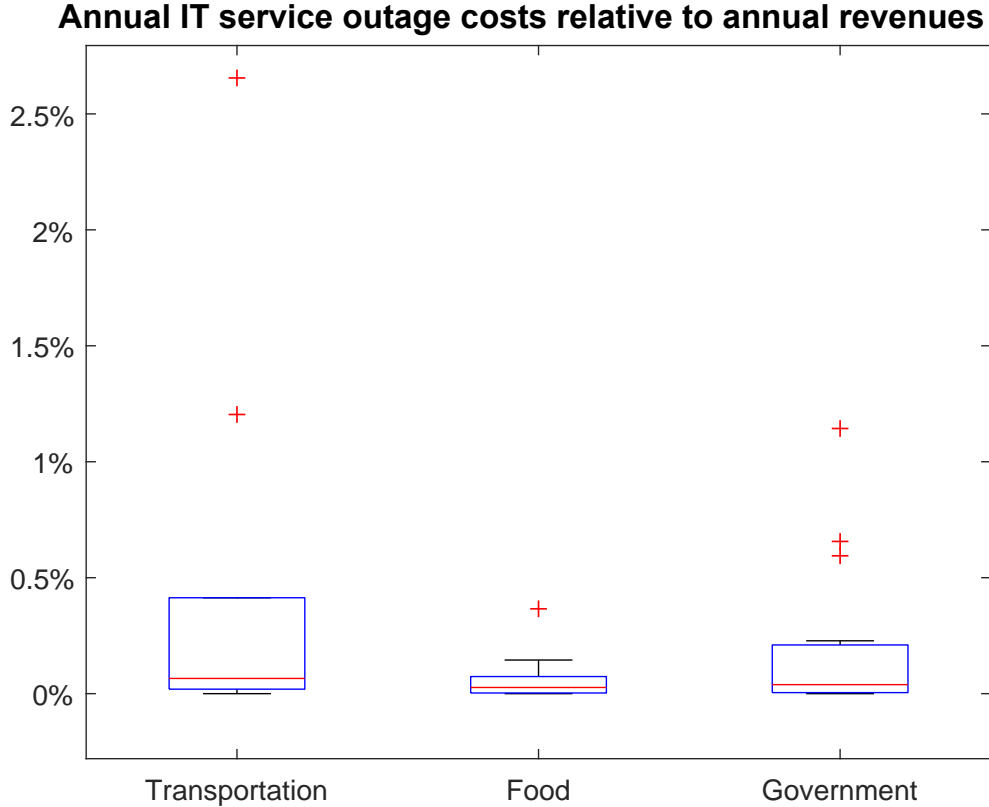


Figure 1: Annual IT service outage costs relative to annual revenues. For transportation,  $N = 9$ , with two respondents removed due to spurious answers (on revenue and lost productivity, respectively). For food,  $N = 9$ , with all respondents included. For government  $N = 18$ , with one respondent removed due to a spurious answer (on revenue).

reach one or even two percent. This is as expected: if IT service outage cost typically reached several percent of revenue, it would be in the same order of magnitude as overall profit margins, and enterprises would regularly go bankrupt in the wake of outages. This does not happen *regularly*. However, the outliers in the box plot do suggest that it *could* happen, which is an important observation for an insurer underwriting these risks.

Furthermore, only two of the six outlier-costs in the plot have been classified by the respondents as not being representative. In fact, these are the two smallest ones. The two greatest costs from the transportation case and the two greatest costs from the government case have, somewhat surprisingly, been classified as representative.

#### 4.5 Loss distribution over cost components

It is also interesting to study the relative importance of the cost components. Specifically, the questionnaire described in Section 3.1 defines three components: (i) the *fixed cost* to restore service per IT service outage (insurable as first-party incident response cost), (ii) the variable cost from *lost productivity* (typically *not* insurable), and (iii) the variable cost from *lost revenue* (insurable as first-party business interruption cost). For each respondent, total annual IT service outage cost can be broken down into contributions from these three.

Fig. 2 displays this loss component-wise distribution in a ternary (simplex) plot. Each position within the triangle represents a possible distribution over the three cost components, as can be read

## Loss distribution over cost components

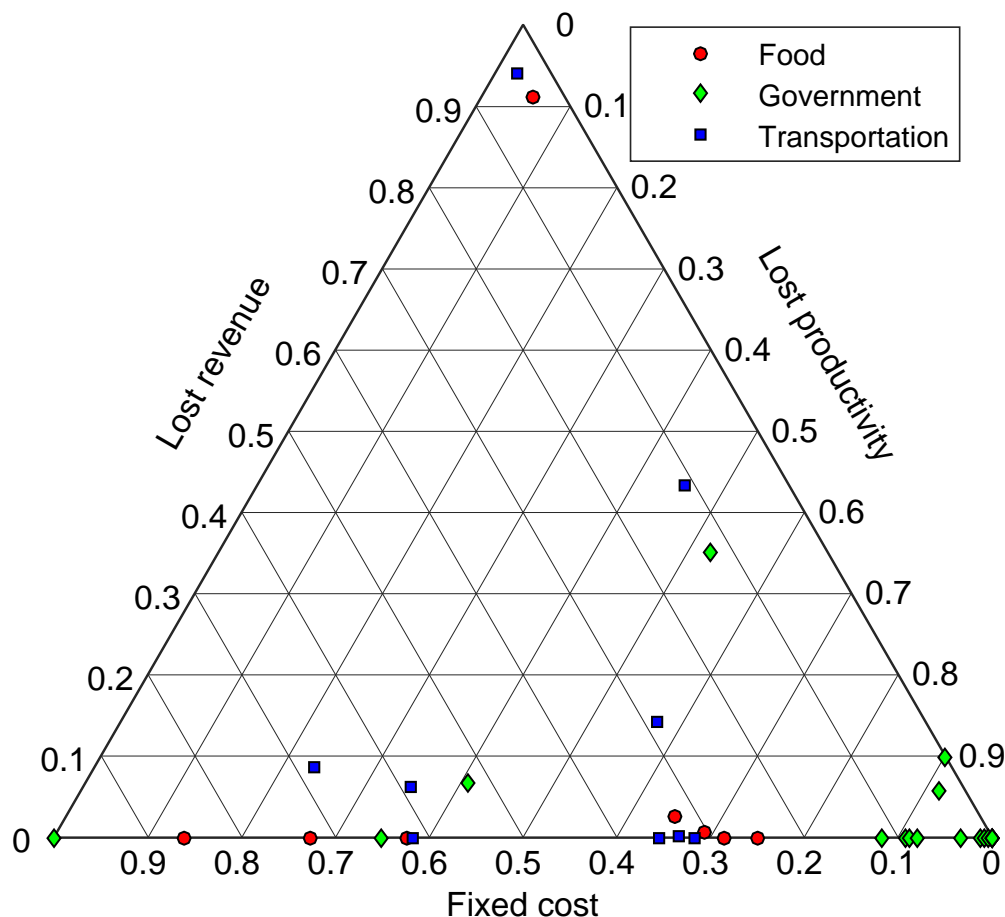


Figure 2: Loss distribution over the three cost components. For transportation,  $N = 9$ , with two respondents removed (one for a spurious answer on lost productivity, one for zero outage cost). For food,  $N = 8$ , with one respondent removed (for zero outage cost). For government,  $N = 18$ , with one respondent removed (for zero outage cost).

off the axes. For example, the transportation company found to the right in the vertical middle of the triangle represents a distribution of 11% fixed cost (potentially insurable), 46% lost productivity (not insurable), and 43% lost revenue (potentially insurable). Similarly, the food company found almost at the top of the triangle represents a distribution of 3% fixed cost (potentially insurable), 5% lost productivity (not insurable), and 91% lost revenue (potentially insurable). As a final example, the government agency at the leftmost corner at the bottom of the triangle represents the extreme distribution of 100% fixed cost (potentially insurable), and no variable cost for lost productivity or revenue.

Some tendencies can be observed. First, in a cluster of government agencies in the lower right corner, almost all outage cost come from lost productivity, with only small contributions from fixed cost or lost income. Indeed, for 13 agencies (hard to distinguish in the plot) lost productivity represents more than 90% of total IT service outage cost. These losses are virtually uninsurable on the present market. Second, there is a large number of enterprises where the revenue lost in IT service outages is comparatively small,

represented by the large number of markers in the lower part of the triangle. Indeed, for 31 of the 35 enterprises in the plot revenue lost represents less than 20% of total losses. Third, the few enterprises with a very high proportion of revenue lost (top of the triangle) are both private companies.

## 4.6 Losses and waiting periods

Cyber insurance coverage of lost revenue from business interruption typically applies a *waiting period*, i.e., a kind of temporal deductible. Compensation for lost revenue is only paid for outage durations that exceed the waiting period, though compensation for incident response cost can be paid for shorter outages as well. Typical waiting periods can be in the order of 24, 36, 48, or even 72 hours, though some 6 or 8 hours can be negotiated (Franke, 2017). Looking at the mean outage durations inferred in the case studies, they all fall short of an hypothetical 24 hour waiting period. Instead assuming a 6 hour waiting period, 8 (out of 36 respondents with outages in 2016) have longer inferred mean outage durations, and would thus in principle be eligible for compensation. However, of these 8, only 4 actually had non-zero lost revenue.

## 4.7 Passing cost on in supply chains

One interesting observation was made by a representative of a cloud service provider, who declined to respond based on the argument that their outage cost is typically passed along to their customers. This is relevant from an insurance perspective, as the ability of cloud service providers to pass outage cost along has been identified as a driver for their customers buying insurance coverage instead (Franke, 2017).

# 5 Implications for investment allocation

The results reported in the previous section are descriptive. However, they also have interesting implications of a more prescriptive nature. To investigate this further, we develop a mathematical investment allocation problem to discuss and contextualize the results. How, conceptually, should an enterprise invest its scarce resources to minimize its net cost of downtime, and what should its insurer incentivize?

## 5.1 An IT service availability investment allocation problem

Steady state availability  $A$  is typically defined as the Mean Time To Failure (MTTF) divided by the total time of operation, which consists of the MTTF plus the Mean Time To Repair/Restore (MTTR):

$$A = \frac{\text{MTTF}}{\text{MTTF} + \text{MTTR}} \quad (2)$$

Adhering to the model introduced in (Franke, 2014) (where more details can be found), assume that an enterprise can invest resources aimed either at increasing MTTF (so that outages occur less frequently) or decrease MTTR (so that outages are recovered more quickly). We assume simple Cobb-Douglas production functions:

$$\text{MTTF} = k_K K^{\beta_K} \quad (3)$$

$$\text{MTTR} = k_L L^{-\beta_L} \quad (4)$$

The  $\beta$  parameters are known as *output elasticities*, and for production with decreasing returns to scales, as is reasonable here,  $\sum \beta_i < 1$ . Note the minus sign in Eq. (4)—as shorter MTTR is better, it decreases with greater investment in  $L$ . The  $k$  parameters are scaling factors, the importance of which will be discussed shortly. The Cobb-Douglas production function (Cobb and Douglas, 1928) is a simple model widely used in economics, which has also been used in studies of software reliability (Kapur et al., 2012).

The intuition behind this model is that Capital  $K$  can buy better hardware and software, increasing MTTF, while Labor  $L$  can be used to monitor a system and take action to recover it upon failure, decreasing MTTR. As also stressed in (Franke, 2014), Labor and Capital should be taken with a grain of salt (a

failover system is Capital that decreases MTTR). The model does, however, distinguish two conceptually different availability investments: production factors that affect MTTF and MTTR, respectively.

Following Section 3.1, the total annual cost of IT service outages derived can be summarized as follows:

$$c_{\text{tot}} = n_{\text{out}}(c_{\text{fix}} + t \cdot c_{\text{var}}) \quad (5)$$

In Eq. (5),  $c_{\text{tot}}$  is the total annual cost of IT service outages,  $n_{\text{out}}$  is the number of outages (question 5),  $c_{\text{fix}}$  is the fixed restoration cost per outage (question 7),  $t$  is the mean duration, in hours, of an outage, and  $c_{\text{var}}$  is the variable cost per hour, composed by lost productivity (questions 8–9) and lost revenue (questions 10–11).

Eqs. (3)–(5) now allow us to formulate a net cost as the sum of investments  $K$  and  $L$  to avoid outages and the outage cost incurred when outages nevertheless happen:

$$\text{Net cost} = K + L + \frac{t_{\text{op}}}{\text{MTTF}} (c_{\text{fix}} + c_{\text{var}} \cdot \text{MTTR}) = K + L + \frac{t_{\text{op}}}{k_K K^{\beta_K}} (c_{\text{fix}} + c_{\text{var}} \cdot k_L L^{-\beta_L}) \quad (6)$$

$t_{\text{op}}$  is the annual operating time, e.g.,  $365 \cdot 24 = 8760$  hours per year. Dividing it by MTTF gives the number of annual outages, each of which entails (i) a fixed cost independent of duration and (ii) a variable cost proportional to MTTR. (More precisely,  $t_{\text{op}}$  should be divided by the sum of MTTF and MTTR, but we know that  $\text{MTTF} \gg \text{MTTR}$ .)

An enterprise wants to minimize Eq. (6), and doing so, we can now analyze how availability investments should be allocated between  $K$  and  $L$ , i.e., between increasing MTTF and decreasing MTTR. Setting a fixed budget  $M = K + L$  and using the method of Lagrange multipliers, we find that an optimal allocation  $\{K^*, L^*\}$  should satisfy the following first order condition (details in appendix):

$$\frac{K^*}{L^*} = \frac{\beta_K}{\beta_L} \left( \frac{c_{\text{fix}}}{c_{\text{var}}} \frac{L^{*\beta_L}}{k_L} + 1 \right) \quad (7)$$

Since Eq. (7) does not give  $K^*$  and  $L^*$  on a closed form, it can be difficult to draw conclusions from it. However, two special cases can readily be seen:

$c_{\text{fix}} \ll c_{\text{var}}$  In the case of *negligible fixed cost*, the first term within the parentheses of Eq. (7) is close to zero, and  $\frac{K^*}{L^*} = \frac{\beta_K}{\beta_L}$ . Thus,  $K$  and  $L$  are allocated by the ratio of their output elasticities.

$c_{\text{fix}} \gg c_{\text{var}}$  In the case of *negligible variable cost*, it can be seen directly from Eq. (6) that any investment in  $L$  is wasted, so  $K^* = M$  and  $L^* = 0$ .

For cost structures between these two extreme cases, the relationship between  $K^*$  and  $L^*$  is non-linear and also dependent on the value of  $k_L$  (but not on  $k_K$ ).

## 5.2 Empirical cost structures and their implications

As the outage cost structure thus has important implications for the allocation of availability investments, it is interesting to revisit the cost structures found empirically in the case studies. Fig. 3 displays these cost structures on the  $x$ -axis with the implied optimal allocation  $K^*$  as a proportion of the total budget on the  $y$ -axis. ( $L^*$ , of course, is the remaining part of the budget). The graphs are based on numerical solutions of Eq. (7). It should be stressed that these are theoretical numbers—investment allocations were not investigated in the questionnaire.

To illustrate the dependence of Eq. (7) on  $k_L$ , three different curves are plotted for different hypothetical  $k_L$  values, expressed in terms of the total budget  $M = K + L$ . For legibility, the respondents are explicitly shown only on the lowest curve, but as their positions on the  $x$ -axis are fixed, their corresponding locations on the other curves are readily found by just shifting them upwards along the  $y$ -axis to the next curve.

Looking at Fig. 3, we first note that the entire range of cost structures is present in the data, including the two limiting cases discussed above. Several respondents have no, or almost no, fixed cost, leading to the optimal allocation where  $\frac{K^*}{L^*} = \frac{\beta_K}{\beta_L}$ . This level is marked separately on the  $y$ -axis, to the lower

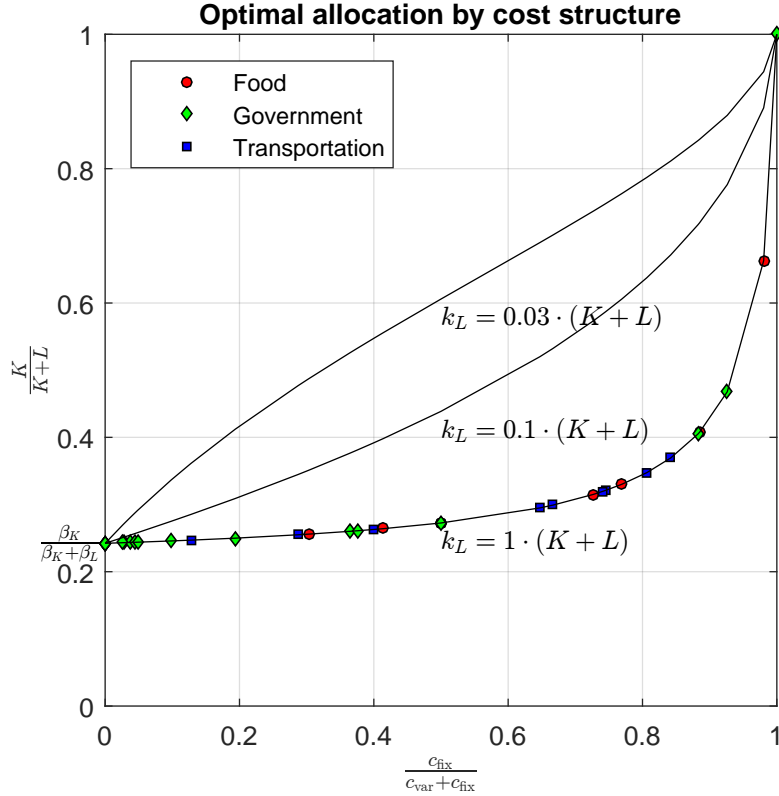


Figure 3: Optimal allocation of a fixed budget  $M = K + L$  ( $y$ -axis) as functions of the cost structure ( $x$ -axis). The respondents plotted are the same as in Fig. 2.

left.<sup>3</sup> On the other side, on the upper right, one respondent has no variable cost, leading to the optimal allocation  $K^* = M$  and  $L^* = 0$ . A few others also have relatively large fixed cost.

It is evident from Fig. 3 that the magnitude of  $k_L$  is important in determining the optimal allocations for cost structures between the extremes. Recalling Eq. (4), a large  $k_L$  implies a large impact of  $L$  investments. This is seen on the lowest curve in Fig. 3, where  $L$  is important and the optimal  $L$  share remains almost constant over a large region of cost structures, before the  $K$  share begins to grow substantially. Conversely, a small  $k_L$  implies a small impact of  $L$  investments, so on the upper curves,  $L$  is less important and the optimal  $K$  share grows more rapidly as the cost structure changes.

Of course, the allocation problem described here is stylized. The exact figures needed to make these calculations in a real case are difficult to find in practice. Nevertheless, the problem does lead to interesting conceptual insights, especially when combined with the case study findings:

- First, while it is sometimes the case that all investments should go into increasing MTTF, it is *never* the case that all investments should go into decreasing MTTR. Regardless of cost structure, investments in MTTF should never go below the proportion implied by its output elasticity over that of MTTR investment.
- Second, ignoring fixed cost can skew investment decisions. This might seem like a trivial observation, but well-cited models such as Patterson (2002) *do* ignore fixed cost. Using Patterson's model,

<sup>3</sup>The axes are normalized, to emphasize the general nature of the plot. As in (Franke, 2014), the particular  $\beta$  values used are  $\beta_K = 0.212$  and  $\beta_L = 0.663$ , building on empirical work by Hitt and Brynjolfsson (1996). For a discussion of the applicability of these numbers, see (Franke, 2014, Section II) .

the optimal ratio of investments into increasing MTTF and decreasing MTTR is always the same, viz. the ratio of output elasticities. But as seen in Figs. 2 (per annum) and 3 (per outage), fixed cost cannot be ignored. It is an important insight of practical relevance that the greater the share of fixed cost, the greater the share of the budget that should go to increasing MTTF.

- A third insight relates to the *non-linearity of outage cost*. As put by Gartner (Vecchio, 2016): “While one hour out of 2,000 may represent one-twentieth of 1% of annual revenue, a five-day (40-hour) outage may have an impact substantially greater than 2% of revenue. That is, for some firms, a brief interruption may be of little consequence, while a substantial outage could jeopardize the firm’s financial viability.” In this respect, Eq. (1) is clearly a simplification, though often a necessary one in practice, e.g., to elicit questionnaire responses. Nevertheless, it is possible to modify Eq. (6) to derive some more insight into this problem. In (Franke, 2012), the snowball effect illustrated in the Gartner quotation is modeled by a cost term growing quadratically with outage duration. If Eq. (6) is modified by squaring MTTR, it is straightforward to derive the new first order optimality condition for  $K^*$  and  $L^*$ :

$$\frac{K^*}{L^*} = \frac{\beta_K}{2\beta_L} \left( \frac{c_{\text{fix}}}{c_{\text{var}}} \frac{L^{*2\beta_L}}{k_L^2} + 1 \right) \quad (7A)$$

Compared to Eq. (7), taking MTTR to the power of 2 in Eq. (6) has translated into multiplying  $\beta_L$  with 2 and taking  $k_L$  to the power of 2 in Eq. (7A). *Mutatis mutandis*, the same changes apply to any exponent  $n$  applied to MTTR. As seen in Eq. (7A), the optimal allocation in the case of negligible fixed cost has changed, so that  $\frac{K^*}{L^*} = \frac{\beta_K}{2\beta_L}$ . In fact, except for the extreme case of fixed cost only, where all resources are devoted to  $K$ , the new condition in Eq. (7A) shifts the optimal balance towards more  $L$  over the entire range of cost structures. Thus the presence of a snowball effect consistently implies devoting relatively more resources to decreasing MTTR, in order to limit a cost growing polynomially in time.

## 6 Discussion

This section discusses the results, starting by scrutinizing the individual cost estimates, then addressing implications for cyber insurance.

### 6.1 Validity and reliability of individual cost estimates

There are several reasons why the measurements of total annual IT service outage cost should be considered valid *for each individual enterprise* responding: (i) The cost is broken down into several components, each of which should be straightforward for the respondent to understand. (ii) The cost resulting from the combination of these component-wise answers was transparently shown to the respondent, who either confirmed that the answer was reasonable, or went back to revise the component-wise answers. This error-correcting mechanism ensures that the answers correspond to the *deliberated assessments* of the respondents, rather than merely their *prima facie* intuitions. (iii) Respondents were also asked to assess the representativeness of the estimated 2016 cost compared to previous years and assessments of future years, implicitly forcing them to make yet another check that the answer was reasonable. (iv) Respondents were given the opportunity to comment, in free text, on the total cost estimate, including any shortcomings.

Since Florêncio and Herley (2013) have pointed out some important limitations of cyber security surveys, it is a good idea to also discuss the results in light of their criticism. In short, a problem of many cyber security surveys is that aggregate numbers are derived from data which is very sensitive to errors. For example, consider determining the total cost of business interruption in a population. A straightforward method to do so is to ask respondents about their costs, finding the average, and then multiply this average with the total number of enterprises in the population. However, if any one respondent has overestimated the cost, this carries over into the aggregate number. Such errors are

indeed a concern, as some actors might have incentives to produce certain results. However, the results reported in Section 4 are *not* aggregated in such ways. Indeed, it is precisely concerns such as that voiced by Florêncio and Herley that motivate the free text format chosen in Sections 4.1–4.3, as well as the use of box plots in Figure 1. For example, the outlier response in the transportation sector seen in Figure 1 would be precisely the kind of extreme value that concerns Florêncio and Herley. While its accuracy is hard to judge, it is clear that it would significantly impact an average value. Therefore, no such average value is given. Instead, the box plot used gives the reader a better view of the entire distribution. Thus, the survey results, as presented in non-aggregate forms in Section 4 do not suffer from this particular threat to validity.

Still, some threats to validity remain: (i) Respondents may not have an accurate understanding of the enterprise-wide consequences of IT service outages. This is particularly true since respondents (as far as can be assessed from the channels used to approach them, and the instructions in the beginning of the questionnaire) typically come from the IT department. (ii) Respondents were explicitly encouraged to give approximate answers, and make a note of this in free text comments, rather than abstain, so answers should be assessed in light of this. (iii) The model used, epitomized in questions 5–12 in Section 3.1 above, may not be applicable or accurate for some (kinds of) enterprises. Following the insights from Section 5, enterprises with very non-linear variable cost or with very long outage durations may not get accurate results. As mentioned in Section 3.1, the survey instrument is agnostic with respect to whether incidents are non-malicious or attacks. This is both a strength and a weakness. The strength is that respondents are not forced to make a distinction that is simple in theory, but difficult in practice (as a skillful attacker can disguise an attack as a non-malicious failure). This clearly eliminates one source of errors. The weakness, however, is that the two kinds of incidents are different in nature. Specifically, as attacks depend on strategic interactions between attackers and defenders, they can be expected to exhibit more complicated (game-theoretical) features, compared to regular failures that can be expected to correspond to a more straightforward (decision-theoretical) problem structure. Thus, care has to be taken when analyzing them together. In particular, this is another reason for not presenting aggregate statistics (such as sums and averages) in Section 4, but rather qualitative descriptions and box plots.

It is important to note the case study nature of the results presented. While results are reasonably valid for each individual enterprise, statistically valid inferences *cannot* be made about the IT service outage costs of Swedish enterprises in general, or the particular sectors investigated, since the sampling is purposive, not random, as described in Section 3. Instead, as put by Yin, the aim of a case study such as the one reported here is analytical, not statistical, generalization (Yin, 2003).

## 6.2 Implications for cyber insurance

As a relatively novel and rapidly developing insurance line, cyber insurance faces a number of obstacles. For example, it is well known that a lack of incident statistics prohibits premiums from being actuarially set (ENISA, 2016; Franke, 2017), reflecting a wider problem of risk quantifiability (OECD, 2017, pp. 94–96). While the study presented in Sections 4 and 5 offers cost structure case studies rather than statistics, some implications and suggested courses of action for insurers can nevertheless be identified:

- It is well known that insurers use reasonable proxies such as revenue or industry to group similar insureds together so that a common pricing model can be applied (Woods et al., 2017; Romanosky et al., 2019). The results presented in Sections 4 and 5 show the limits of this approach: cost structures differ across organizations. On a very strict interpretation, the results thus suggest that cost structure information should be collected individually from each applicant. This is perfectly feasible on the market for large enterprises, with indemnity limits in the order of tens or hundreds of millions of dollars or euros, where each insured is thoroughly evaluated through a long and complicated underwriting process (such as most of the Swedish cyber insurance market from a few years ago (Franke, 2017)). But it is certainly not feasible on the mass market for SMEs, where insurers underwrite cyber risk based on just a few questions, or even without any cyber specific questions at all. Thus, efforts should be directed towards finding good ways to reliably infer cost structures *without* necessarily asking all the questions listed in Section 3.1.

- Using an instrument such as the questions presented in Section 3.1 (or a more efficient proxy), insurers can get an overview of how the potential losses look in a portfolio, thus addressing concerns over accumulation risk (Hofmann et al., 2018; OECD, 2017, pp. 96–98), in their portfolios. For example, a portfolio where some insureds have mostly fixed outage cost and some insured have mostly variable outage cost should, *ceteris paribus*, be considered less risky than a more homogeneous portfolio.
- Using an instrument as the questions presented in Section 3.1 in underwriting (on the large enterprise market) could have a pedagogical value. Lack of awareness of potential losses and misunderstandings about coverage are two well-known obstacles for cyber insurance becoming a mature tool for cyber risk management and transfer (OECD, 2017, pp. 101–104), and an instrument showing the different components of outage cost, some of which are covered, some of which are not, could help the demand-side better understand what it buys.
- Continuing this line of reasoning, there is also scope for product development. Some insurers might want to develop a policy that *does* cover lost productivity, as this is indeed a substantial cost driver for some potential customers, especially in the public sector, as seen in Fig. 3. However, such a product would also entail new challenges in terms of asymmetric information and accounting transparency.
- Insurers may also want to proactively nudge their customers towards more security, by influencing their availability investment strategies as outlined in Section 5. While this is certainly useful on the individual enterprise level, it might be even more useful if an insurer could get a bird’s-eye view of its portfolio, including interdependencies where IT service outages at one insured could spill over to others. In such situations, insurers might be able to resolve collective action problems by innovative new business models along the lines suggested by Wang (2019).

## 7 Conclusions and future work

Relating back to the three research questions posed at the end of Section 1, some conclusions can be drawn:

1. *The instrument* developed in Section 3.1 works quite well, as elaborated in Section 6.1. For example, it can demonstrably be used in both the public and the private sector, and the respondents find the cost estimates reasonable. Still, the number of questions required may be too taxing when underwriting SMEs.
2. *The cost structures* of different enterprises differ, as illustrated in Fig. 2, and they differ substantially. For example, some enterprises incur only fixed outage cost, some incur (almost) only lost productivity, and some incur almost only lost revenue. This may seem trivial *a posteriori*, but it takes empirical investigation to know. The differing cost structures, furthermore, entail other consequences. For example, fixed cost and lost revenue are insurable, whereas lost productivity is not. *In the public sector, lost revenue is often negligible.* However, this does *not always* hold—some government agencies do experience substantial revenue losses. They may wish to pursue different risk management strategies than other agencies, e.g. using insurance to a greater extent. This is a good example of how a case study can make us better understand what is *possible*, even in the absence of statistics one exactly *how often* it is encountered in a population.
3. *Different cost structures matter* to availability investment strategies, as seen in Section 5. Though stylized, the availability investment allocation problem provides practical insight relating to the relative importance of investments in increasing MTTF and decreasing MTTR, respectively, the importance of fixed cost, and the non-linearity of outage cost.

In addition to these general conclusions, some more insurance-specific implications were discussed in the previous section.

The results also suggest some interesting future work. One such direction is to further investigate the notion of IT service outage cost consisting of different cost components, as illustrated in Fig. 2. In particular, it would be interesting to contrast the government agencies with, e.g., municipalities, to see whether lost revenue is typically negligible there as well, and to contrast the results from transportation and food with other private sector industries. Another direction is to conduct in-depth interviews to collect richer and more qualitative information on outage management and IT service availability decision-making. A third idea is to do a similar exercise on the insurance side—mapping rationales behind current cost calculations and assessing the opportunities for product development. A fourth worthwhile direction would be to work with government agencies responsible for introducing reporting regimes on service outages, e.g., those mandated by the NIS Directive in the EU (European Commission, 2017), thus obtaining more cost data and a correspondingly improved picture of outage cost. A fifth idea is to further investigate the differences between outages caused by (i) non-malicious failures or (ii) attacks, as these are typically treated differently in insurance policies. Finally, a sixth direction involves finding more efficient means to reliably infer cost structure, as discussed in Section 6.2.

## Appendix: Derivation of Eq. (7)

Denoting the net cost in Eq. (6) by  $f$  and the fixed budget constraint by  $g$ , we have:

$$\min_{K,L} f = K + L + \frac{t_{\text{op}}}{k_K K^{\beta_K}} (c_{\text{fix}} + c_{\text{var}} \cdot k_L L^{-\beta_L})$$

such that  $g = K + L - M = 0$

The Lagrangian is thus  $\mathcal{L} = f(K, L) - \lambda g(K, L)$ , and we solve the system of equations  $\nabla \mathcal{L} = \mathbf{0}$ :

$$\begin{cases} \frac{\partial \mathcal{L}}{\partial K} = 1 - \frac{t_{\text{op}}}{k_K} \beta_K K^{-\beta_K - 1} (c_{\text{fix}} + c_{\text{var}} k_L L^{-\beta_L}) - \lambda = 0 \\ \frac{\partial \mathcal{L}}{\partial L} = 1 - \frac{t_{\text{op}}}{k_K} K^{-\beta_K} c_{\text{var}} k_L \beta_L L^{-\beta_L - 1} - \lambda = 0 \\ \frac{\partial \mathcal{L}}{\partial \lambda} = K + L - M = 0 \end{cases}$$

Using the first two equations, we eliminate  $\lambda$  and obtain:

$$\frac{t_{\text{op}}}{k_K} \beta_K K^{-\beta_K - 1} (c_{\text{fix}} + c_{\text{var}} k_L L^{-\beta_L}) = \frac{t_{\text{op}}}{k_K} K^{-\beta_K} c_{\text{var}} k_L \beta_L L^{-\beta_L - 1}$$

Dividing by  $\frac{t_{\text{op}}}{k_K} K^{-\beta_K - 1} \beta_L$  we have

$$\frac{\beta_K}{\beta_L} (c_{\text{fix}} + c_{\text{var}} k_L L^{-\beta_L}) = c_{\text{var}} k_L K L^{-\beta_L - 1}$$

Dividing by  $c_{\text{var}} k_L$  and multiplying by  $L^{\beta_L}$  we have

$$\frac{\beta_K}{\beta_L} \left( \frac{c_{\text{fix}}}{c_{\text{var}}} \frac{L^{\beta_L}}{k_L} + 1 \right) = \frac{K}{L}$$

which is the first-order optimality condition for  $K^*$  and  $L^*$  in Eq. (7).

## Acknowledgments

This work was supported by the Swedish Civil Contingencies Agency, MSB, agreement no. 2015-6986. Not only did MSB function as the funding agency, but was also instrumental in securing access to the public-private fora where many respondents were recruited. The author is grateful for this, in particular to Johan Turell, who facilitated these contacts. Thanks are also due to Professor Shaun S. Wang of the Nanyang Technological University in Singapore, for discussions about availability investment allocation problems.

## References

- Andrade E, Nogueira B, Matos R, Callou G, Maciel P (2017) Availability modeling and analysis of a disaster-recovery-as-a-service solution. *Computing* pp 1–26, DOI [10.1007/s00607-017-0539-8](https://doi.org/10.1007/s00607-017-0539-8), URL <http://dx.doi.org/10.1007/s00607-017-0539-8>
- Bharadwaj A, Keil M, Mähring M (2009) Effects of information technology failures on the market value of firms. *The Journal of Strategic Information Systems* 18(2):66 – 79, DOI <http://dx.doi.org/10.1016/j.jsis.2009.04.001>
- Biener C, Eling M, Wirfs JH (2015) Insurability of cyber risk: An empirical analysis. *Geneva Papers on Risk and Insurance: Issues and Practice* 40(1):131–158, URL <http://dx.doi.org/10.1057/gpp.2014.19>
- Bosse S, Splieth M, Turowski K (2016) Multi-objective optimization of IT service availability and costs. *Reliability Engineering & System Safety* 147:142 – 155, DOI <http://dx.doi.org/10.1016/j.res.2015.11.004>
- Carfora M, Martinelli F, Mercaldo F, Orlando A (2019) Cyber risk management: an actuarial point of view. *Journal of Operational Risk* (4):77–103, DOI [10.21314/JOP.2019.231](https://doi.org/10.21314/JOP.2019.231)
- Cobb CW, Douglas PH (1928) A theory of production. *The American Economic Review* 18(1):139–165, URL <http://www.jstor.org/stable/1811556>
- Daffron J, Ruffle S, Andrew C, Copic J, Quantrill K, A S, Leverett E (2019) Bashe attack: Global infection by contagious malware. Tech. rep., Cambridge Centre for Risk Studies, Lloyd’s of London and Nanyang Technological University, URL <https://www.lloyds.com/news-and-risk-insight/risk-reports/library/technology/bashe-attack>, accessed February 4, 2019.
- Durkee D (2010) Why cloud computing will never be free. *Queue* 8(4):20, URL <http://dx.doi.org/10.1145/1755884.1772130>
- Edwards B, Hofmeyr S, Forrest S (2015) Hype and heavy tails: A closer look at data breaches. In: *The Workshop on the Economics of Information Security (WEIS)*
- Eling M, Loperfido N (2017) Data breaches: Goodness of fit, pricing, and risk measurement. *Insurance: mathematics and economics* 75:126–136
- Eling M, Schnell W (2016) What do we know about cyber risk and cyber risk insurance? *The Journal of Risk Finance* 17(5):474–491, URL <http://dx.doi.org/10.1108/JRF-09-2016-0122>
- ENISA (2016) Cyber insurance: Recent advances, good practices and challenges. Tech. rep., European Union Agency for Network and Information Security, URL <http://dx.doi.org/10.2824/065381>
- European Commission (2017) The directive on security of network and information systems (nis directive). URL <https://ec.europa.eu/digital-single-market/en/network-and-information-security-nis-directive>
- European Commission (2019) The Digital Economy & Society Index (DESI). URL <https://ec.europa.eu/digital-single-market/en/desi>, accessed on September 27, 2019.
- Falco G, Eling M, Jablanski D, Gordon LA, Wang SS, Schmit J, Thomas R, Elvedi M, Maillart T, Donovan E, Dejung S, Weber M, Durand E, Nutter F, Scheffer U, Arazi G, Ohana G, Lin H (2019) A research agenda for cyber risk and cyber insurance. *Workshop on the Economics of Information Security (WEIS)*

- Florêncio D, Herley C (2013) Sex, lies and cyber-crime surveys. In: Schneier B (ed) *Economics of Information Security and Privacy III*, Springer New York, New York, NY, pp 35–53, DOI 10.1007/978-1-4614-1981-5\_3
- Franke U (2012) Optimal IT Service Availability: Shorter Outages, or Fewer? *Network and Service Management*, IEEE Transactions on 9(1):22–33
- Franke U (2014) Enterprise Architecture Analysis with Production Functions. In: *IEEE 18th International Enterprise Distributed Object Computing Conference (EDOC 2014)*, IEEE, pp 52–60, DOI 10.1109/EDOC.2014.17
- Franke U (2017) The cyber insurance market in Sweden. *Computers & Security* 68:130–144, DOI 10.1016/j.cose.2017.04.010
- Franke U, Holm H, König J (2014) The distribution of time to recovery of enterprise IT services. *IEEE Transactions on Reliability* 63(4):858–867, DOI 10.1109/TR.2014.2336051
- Goldstein J, Chernobai A, Benaroch M (2011) An event study analysis of the economic impact of IT operational risk and its subcategories. *Journal of the Association for Information Systems* 12(9):1
- Hitt LM, Brynjolfsson E (1996) Productivity, business profitability, and consumer surplus: three different measures of information technology value. *MIS Quarterly* 20(2):121–142
- Hofmann DM, Wilson S, Carter RA (2018) Advancing accumulation risk management in cyber insurance. Tech. rep., The Geneva Association, URL <https://www.genevaassociation.org/research-topics/cyber-and-innovation/advancing-accumulation-risk-management-cyber-insurance>, accessed February 25, 2019.
- IBM Global Services (1998) Improving systems availability. Tech. rep., IBM Global Services
- Ibrahimovic S, Franke U (2016) A probabilistic approach to IT risk management in the Basel regulatory framework: A case study. *Journal of Financial Regulation and Compliance* 25:176–195, DOI 10.1108/JFRC-06-2016-0050
- Insurance Sweden (2019) Hur försäkrar vi oss mot cyberrisker och databrott? [How can we insure ourselves against cyber risks and cyber crimes?]. URL <https://www.svenskforsakring.se/aktuellt/nyheter/2019/hur-forsakrar-vi-oss-mot-cyberrisker-och-databrott/>, accessed on September 27, 2019.
- Jammal M, Hawilo H, Kanso A, Shami A (2017) Mitigating the risk of cloud services downtime using live migration and high availability-aware placement. pp 578–583, DOI 10.1109/CloudCom.2016.0100
- Kapur P, Pham H, Aggarwal AG, Kaur G (2012) Two dimensional multi-release software reliability modeling and optimal release planning. *IEEE Transactions on Reliability* 61(3):758–768
- Lerner A, Ganguli S, Bhalla V (2016) How to Reduce Network Downtime in the Era of Digital Business. Tech. rep., Gartner, Inc., g00317252
- Li X, Qi Y, Chen P, Zhang X (2017) Optimizing backup resources in the cloud. pp 790–797, DOI 10.1109/CLOUD.2016.107
- Logeswaran L, Bandara H, Bhatthiya H (2017) Performance, resource, and cost aware resource provisioning in the cloud. pp 913–916, DOI 10.1109/CLOUD.2016.133
- Meland PH, Tøndel IA, Moe M, Seehusen F (2017) Facing uncertainty in cyber insurance policies. In: *International Workshop on Security and Trust Management*, Springer, pp 89–100

- Morency JP (2014) Managing IT Resilience Is Much More Than Simply Failing Over Applications. Tech. rep., Gartner, Inc., updated September 2016. G00233822
- Nguyen TA, Kim DS, Park JS (2016) Availability modeling and analysis of a data center for disaster tolerance. *Future Generation Computer Systems* 56:27 – 50, DOI <http://dx.doi.org/10.1016/j.future.2015.08.017>
- OECD (2017) Enhancing the Role of Insurance in Cyber Risk Management. DOI 10.1787/9789264282148-en
- Olsson T, Franke U (2019) Risks and Assets: A Qualitative Study of a Software Ecosystem in the Mining Industry. In: Proceedings of the 2019 27th ACM Joint Meeting on European Software Engineering Conference and Symposium on the Foundations of Software Engineering, ACM, ESEC/FSE 2019, pp 895–904, DOI 10.1145/3338906.3340443
- Patterson D (2002) A simple way to estimate the cost of downtime. In: Proc. 16th Systems Administration Conf.—LISA, pp 185–8
- Ponemon (2016) 2016 cost of data center outages. Tech. rep., Ponemon Institute and Emerson Network Power
- Rachev ST, Chernobai A, Menn C (2006) Empirical examination of operational loss distributions. In: Perspectives on Operations Research, Springer, pp 379–401
- Rapoza J (2014) Preventing virtual application downtime. Tech. rep., Aberdeen Group
- Romanosky S, Ablon L, Kuehn A, Jones T (2019) Content analysis of cyber insurance policies: how do carriers price cyber risk? *Journal of Cybersecurity* 5(1):1–19, DOI 10.1093/cybsec/tyz002
- Sohlberg I, Jansson S (2012) Dolda it-kostnader i verksamheten. Försäkringskassan och Pensionsmyndigheten. [Hidden enterprise IT costs. The Swedish Social Insurance Agency and the Swedish Pensions Agency]. Swedish Social Insurance Inspectorate, report 2012:5
- Sousa E, Lins F, Tavares E, Maciel P (2017) Cloud infrastructure planning considering different redundancy mechanisms. *Computing* pp 1–24, DOI 10.1007/s00607-016-0533-6, URL <http://dx.doi.org/10.1007/s00607-016-0533-6>
- Statskontoret [The Swedish Agency for Public Management] (2017) Den offentliga sektorn i korthet 2017]. URL <http://www.statskontoret.se/globalassets/publikationer/2017/offentliga-sektorn-korthet-2017.pdf>, No. 2017/20-5.
- Tonn G, Kesan JP, Zhang L, Czajkowski J (2019) Cyber risk and insurance for transportation infrastructure. *Transport policy* 79:103–114, DOI 10.1016/j.tranpol.2019.04.019
- Vecchio D (2016) How to Derive Business Value From DevOps. Tech. rep., Gartner, Inc., G00317166
- Wang SS (2019) Integrated framework for information security investment and cyber insurance. *Pacific-Basin Finance Journal* 57:101173, DOI 10.1016/j.pacfin.2019.101173
- Woods D, Agrafiotis I, Nurse JR, Creese S (2017) Mapping the coverage of security controls in cyber insurance proposal forms. *Journal of Internet Services and Applications* 8(1):8
- Woods D, Moore T, Simpson A (2019) The county fair cyber loss distribution: Drawing inferences from insurance prices. Workshop on the Economics of Information Security (WEIS)
- World Economic Forum (2016) The 10 countries best prepared for the new digital economy. URL <https://www.weforum.org/agenda/2016/07/countries-best-prepared-for-the-new-digital-economy/>, accessed on January 9, 2017.

World Economic Forum (2018) Cyber resilience playbook for public-private collaboration. Tech. rep., World Economic Forum, URL [http://www3.weforum.org/docs/WEF\\_Cyber\\_Resilience\\_Playbook.pdf](http://www3.weforum.org/docs/WEF_Cyber_Resilience_Playbook.pdf), accessed March 9, 2018. REF 110117.

Yin RK (2003) Case Study Research: Design and Methods. Applied Social Research Methods, Vol. 5, SAGE Publications, Inc