Diss. ETH No. 18132

Management of Information System Risks

A dissertation submitted to
ETH ZURICH

for the degree of
Doctor of Sciences

presented by
DOMENICO SALVATI

lic. oec. publ., University of Zurich
born January 23rd, 1968
citizen of Cham, Switzerland and Italy

accepted on the recommendation of
Prof. Dr. Wolfgang Kröger, examiner
Prof. Dr. Adrian Gheorghe, co-examiner
Prof. Dr. Hans-Jakob Lüthi, co-examiner
Prof. Dr. Solange Ghernaouti-Hélie, co-examiner

Zurich 2008

Acknowledgements

Dear Reader,

I have developed this doctoral thesis in the stressful field of the banking world, the academic world and the world of fiery arts. The banking world perplexed me through four challenges which hinder a more rigorous management of risks in information systems. I named them Ambiguity, Likelihood, Influence and Decision Problems. To solve them I engineered an approach employing scenario techniques, the convolution of signals, Rough Sets Theory and Utility Theory. The world of fire artists helped me in reenergizing myself from the strenuous writing process I engaged in. To travel between the above worlds, I mainly used public trains. In fact, considerable parts of this thesis evolved during weekends on my journeys between Zurich, Innsbruck and Merano.

From the academic world I would like to thank Prof. Dr. Wolfgang Kröger (referee) and the co-referees Prof. Dr. Hans-Jakob Lüthi, Prof. Dr. Adrian Gheorghe and Prof. Dr. Solange Ghernaouti-Hélie who taught me academic rigor. I also thank my colleagues at the Laboratory for Safety Analysis of the Swiss Federal Institute of Technology in Zurich (ETHZ) namely, Monika Mortimer for her administrative support as well as Dr. Martin Diergardt and PhD candidate Manuel Kaegi for listening to and sound boarding my ideas. I would also like to mention Dr. Marco Laumanns of the Institute of Operations Research and Jürg Schelldorfer of the Seminar for Statistics of ETHZ for their invaluable support. Further, I would like to thank PhD candidates Markus Schlaepfer and Sigrid Wagner for two delightful evenings with an inspiring strategy game symptomatically called “Risk!”. I also thank Dr. Roman Bolinger, Dr. Nathalie Weiler and Dr. Gritta Wolf for reading the manuscript.

From the banking world I would like to thank Adolf Doerig of Pukall Doerig + Partner, Dr. Helmut Kaufmann and Dr. Stephan Murer of Credit Suisse as well as Dr. Ralph Holbein of Ernst & Young for overlooking the development of the case study in the context of a global bank. Further, I would like to thank Simon Zumstein of scip, Stephan Glaus of the Swiss Federal Police, Franco Cerminara of Infoguard as well as Claudia Gutermann, Ueli Baertschi, Martin Walder and Philippe Menotti of Credit Suisse for facilitating or

providing the raw data. I also thank Dr. Lukas Ruest, Rob Ife and Andrew Brice of Credit Suisse, Walter Widmer of UBS and Dr. Paul Schoebi of cnlab for the enthusiasm with which they followed up on my thesis. I thankfully mention that Credit Suisse sponsored my studies for three years.

Finally, I would like to thank my fiancée, Walburga Poehl, my parents Francesco and Leonilde, my sister Tilde and her family Lucas, Moreno and Michelle. They helped me in finding a balance between work and my precious spare time. My friends in Innsbruck and Merano, in particular, welcomed me to their world of wonders. Thanks goes to the Schwartz family, Rolf, Tamsn, Sonja, Clemens, Juls, Daewo, Patrizia, Verena, Cedar and Andrea, the fire artists of Tirasaru, Ninaruna and Spielvolk, the Elaminje drummers as well as Mucky, Maria, Sergio and Gerti, Günther and Greti, Leo, Elia and Moritz.

Zurich, October 2008 Domenico Salvati

Summary

Compliance Management aims at ensuring regulatory-compliant behaviour of a company's work force by requiring the presence of control processes which are intended to suppress illegal conduct such as money laundering, insider dealing, and corruption. The design of the individual controls is usually not prescribed by a legislator but is found in standards or best practices instead.

If executed diligently, compliance management is the pre-condition for a company to hope for milder sentences should culprit employees be convicted in a court of law. Since senior executives of a company can be held liable with their personal fortune, it is not surprising to see that compliance management has grown rapidly and control processes are spreading as industries become more regulated.

This spread is disconcerting as compliance management is, unfortunately, more and more applied to securing a company's information systems (IS). The approach misleads companies to emphasize costly control processes rather than rigorous IS risk management. This yields bureaucracy-like efficiency and effectiveness and its application is disputed, in particular, for the mitigation of Internet threats which are the focus of this work. Compliance management alone does not yield security and complacent control processes may even increase risk. Instead, it is desirable to rigorously apply IS risk management as it promises the efficient and effective selection of security mechanisms for counteracting, e.g., cyber attacks.

Regrettably, the application of IS risk management is impeded by four problems:

- The Ambiguity Problem: In IS, interpretations of risk differ from community to community
- The Likelihood Problem: In practice, likelihood ratings are poorly accepted and subject to controversial discussions
- The Influence Problem: The operational context of IS is not taken into account for considerations on likelihood
- The Decision Problem: Today's decision criteria for selecting security mechanisms do not fit a rigorous risk approach

To solve the aforementioned problems it was necessary to develop an understanding of the epistemological nature of IS risk. This revealed that today's IS risk concepts and terminology need revision. Consequently, more mature concepts from disciplines such as the management of risks in technical systems were investigated and adapted.

In particular, the Ambiguity Problem was mainly solved by introducing state-based event sequences for modeling scenarios. Furthermore, we learned that a threat and a security mechanism need to be treated as one indivisible entity. This concept yields an understanding where the threat is described by a distinct probability distribution and the security mechanism responding to the attack is represented by another. By convoluting the two curves – an approach which is commonly used in communications engineering for processing electric signals – the Likelihood Problem is solved. In essence, the Influence Problem is solved by pattern recognition. From a variety of tools and techniques available today, Rough Sets Theory was adopted mainly because of its ability to infer results based on little, incomplete and vague data. The Decision Problem was solved by adapting decision models from Utility Theory, which had not yet been applied for its use in IS risk management. In particular, the risk preferences of corporate decision makers were explored.

As a result, four modules are presented, each of which has been designed to solve one of the aforementioned problems:

1. the Process Module reduces ambiguity in describing risk
2. the Function Module introduces a general approach to estimate likelihood
3. the Influence Module evidences the influence of the context on likelihood
4. the Decision Module makes risk preferences the decisive criteria for developing security policies to counteract threats

The Four Modules form a new and comprehensive model for IS risk management. These are intended for adoption within large companies and supports senior executives in risk informed decision making. Finally, the practical applicability of the modules has been successfully verified by a case study at a global financial institution.

Zusammenfassung

Mit Compliance Management wird das gesetzeskonforme Verhalten der Belegschaft eines Unternehmens sichergestellt, indem Kontrollprozesse eingeführt werden. Diese Kontrollprozesse unterdrücken illegales Verhalten wie Geldwäscherei, Insidergeschäfte und Korruption. Der Entwurf der einzelnen Kontrollen ist üblicherweise nicht durch einen Gesetzgeber vorgeschrieben und wird typischerweise durch die Anwendung von Standards oder gängigen Praktiken im Unternehmen umgesetzt.

Sollten Mitarbeiter in einem Strafverfahren schuldig gesprochen werden, stellt das Compliance Management eine Vorbedingung für ein Unternehmen dar, um auf mildere Strafen zu hoffen. Weil zudem in gewissen Fällen das Management mit dem privaten Vermögen haften kann, erstaunt es nicht, dass Compliance Management schnell an Bedeutung gewonnen hat und die entsprechenden Kontrollprozesse ein Unternehmen durchwachsen, sobald eine Branche stärker reguliert wird.

Die Ausbreitung von Kontrollprozessen ist jedoch bedenklich, da diese zur Sicherung von immer mehr Informationssystemen (IS) zur Anwendung kommen. Dies kann bei grösseren Unternehmen in eine bürokratische Ineffizienz und Ineffektivität führen, insbesondere bei Gefahren aus dem Internet, welche den Schwerpunkt der vorliegenden Arbeit bilden. Das Betreiben von Compliance Management alleine stellt demnach keine akzeptable Grundlage für das Management von Sicherheit. Ausserdem erhöht der selbstgefällige Umgang mit Kontrollprozessen IS Risiken. Stattdessen ist es wünschenswert, ein rigoroses IS Risikomanagement zu betreiben, welches die effiziente und effektive Auswahl von Sicherheitsmechanismen ermöglicht.

Bedauerlicherweise wird die rigorose Anwendung von IS Risikomanagement in der Praxis durch vier Problemfelder erschwert:

- Das Mehrdeutigkeits-Problem: Es gibt viele verschiedene Interpretationen von Risiko. Diese ändern von Fachschaft zu Fachschaft
- Das Wahrscheinlichkeits-Problem: Wahrscheinlichkeitsschätzungen sind stark erwünscht, jedoch in der Praxis vielfach Gegenstand kontroverser Diskussionen
- Das Einfluss-Problem: Der betriebliche Kontext von IS wird in der Regel bei der Schätzung von Wahrscheinlichkeiten nicht berücksichtigt

- Das Entscheidungs-Problem: Entscheidungskriterien für die Auswahl von Sicherheitsmechanismen erfüllen nicht die Anforderungen eines rigorosen Risiko-Ansatzes

Um obige Probleme zu lösen, wurde der erkenntnistheoretische Gegenstand von Risiken in IS revidiert und es mussten gängige Konzepte sowie Terminologie überarbeitet werden. Es wurden auch explizit Konzepte angepasst und übernommen, die nicht im Kernbereich von IS liegen, wie z.B. das Risikomanagement von technischen Systemen.

Im Speziellen wird das Mehrdeutigkeits-Problem gelöst, indem statusbasierte Ereignisbäume zur Modellierung von Szenarien vorgeschlagen werden. Gefahren und Sicherheitsmechanismen werden als eine unteilbare Einheit betrachtet und werden in je eine unabhängige Wahrscheinlichkeitsverteilung dargestellt. Durch die Faltung beider Verteilungen — ein Vorgehen, welches in der Telekommunikation zur Bearbeitung elektrischer Signale angewandt wird — werden Wahrscheinlichkeiten berechnet. Im Grunde lässt sich das Einfluss-Problem durch Mustererkennung lösen. Aus einem grossen Fundus an Werkzeugen und Techniken wurden Rough Sets zu dessen Lösung herangezogen; hauptsächlich wegen ihrem Vermögen, Resultate aus wenigen, unvollständigen und vagen Daten abzuleiten, was in der Praxis einen Vorteil darstellt. Das Entscheidungs-Problem lässt sich durch Adaption entsprechender Modelle aus der Nutzentheorie lösen, einem Ansatz, welcher bisher noch nicht im IS Risikomanagement eingesetzt wurde. In diesem Rahmen wurden insbesondere die Risikopräferenzen von Entscheidungsträgern zur Entscheidungsfindung berücksichtigt.

Aus obigen Lösungsansätzen ergeben sich vier Module, welches jedes für sich die Lösung eines der oben genannten Probleme darstellt. Insbesondere sind dies:

1. das Prozess-Modul verringert Mehrdeutigkeiten beim Beschreiben von Risiken
2. das Funktions-Modul zeigt ein generelles Vorgehen zur Berechnung von Wahrscheinlichkeiten
3. das Einfluss-Modul zeigt Einflüsse der Umwelt auf Wahrscheinlichkeiten auf
4. das Entscheidungs-Modul erhebt Risikopräferenzen zum Kriterium für das Entwickeln von Sicherheitspolitiken zur Bekämpfung von Gefahren

Die Module formen einen umfassenden Rahmen für das Management von IS Risiken. Sie sind für den Einsatz in grosse Unternehmen ausgerichtet und unterstützen das höhere Management bei der „risiko-informierten“ Entscheidungsfindung. Die praktische Anwendbarkeit der Module wurde in einem global tätigem Finanzinstitut verifiziert.

Table of Contents

List of Figures

List of Tables

1. Problem Statement

To set a common ground, the Risk Terminology Primer is introduced in Chapter 1.1. Chapter 1.2 describes the general security context of global companies, which is determined by fierce international competition, mergers/acquisitions and the ubiquity of the insecure Internet. Such a context weakens IS security in companies and they face a variety of Internet threats, which are summarized in Chapter 1.3. An increasing number of threats are operated along very efficient, industry-like attack vectors while security mechanisms and processes in global companies are not. Consequently, the adoption of rigorous principles in IS risk management in order to respond more efficiently to Internet threats is advocated in Chapter 1.4. In Chapter 1.5, the Everglades of IS Risk are elaborated. They indicate four main areas where substantial progress is required in order for IS risk management to unfold its potential in:

- neutralizing ambiguity when describing IS risk to decision makers
- increasing confidence in likelihood estimates
- determining the degree of influence of the context on likelihood
- adopting risk preferences as the sole criterion for decision making

The objectives and benefits of such an approach to IS risk are reported in Chapter 1.6 and finally, in Chapter 1.7, an outline of the thesis is given.

1.1 Risk Terminology Primer

The following terms are introduced: information system, asset[1], threat[2] and security mechanism[2], vulnerability and control[2], threat event[1], impact and consequence[1], probability and frequency, and finally, information system risk and risk management.

---

[1] Asset, threat event, impact and consequence are refined in Chapter 2.

Description "Information System": An *INFORMATION SYSTEM (IS)* supports companies in co-ordinating an internal workforce, managing processes, offering products or customer services by pre-processing, transforming and allocating business intelligence. Given its outstanding importance, the term is not to be reduced to its digital counterpart, i.e. information technology (IT) as it may encompass other means for handling business intelligence such as paper and telephones.

Description "Asset": An *ASSET* is any component of an information system, which is of value to a company and worth protecting from threats. Given the convergence of many systems for handling business intelligence to a digital form, in this work, it is primarily digital assets as *DATA*, which are considered for protection. Data is ubiquitous and may be encountered in databases, files, or office applications.

Description "Threat", "Security Mechanism": A *THREAT* potentially causes harm to an asset. It is described by a *THREAT AGENT* performing *THREAT ACTIONS*. A threat is opposed by a *SECURITY MECHANISM*. A security mechanism displays *VULNERABILITY* towards a threat or is in *CONTROL* of it. Vulnerabilities may be closed out or reduced with additional security mechanisms.

Description "Threat Event": A *THREAT EVENT* forces an asset to change its state. It results from the clashing of a threat with its counteracting security mechanism. An event evokes an *IMPACT* on the asset in an information system context, i.e. in terms of confidentiality, integrity and availability. Impact must be distinguished from *CONSEQUENCE*, which describes a monetary, legal or reputation change of the asset in a business context, i.e. a shift of figures in a value system.

Description "Probability": *PROBABILITY* is referred to as numbered likelihood (between 0 and 1). It denotes either the subjective belief about the uncertainty of a statement or a statistical measure that a threat overcomes its opposing security mechanism. Probability is distinguished from *FREQUENCY*, which denotes how many times in a

[2] Threat, security mechanism, vulnerability and control are refined in Chapter 3.

specified period of time a threat is present and attempts to overcome the opposing security mechanism.

Description "IS Risk": In accordance with the vocabulary used by the International Standardization Organization (ISO) [1], *IS RISK* is described as the *combination of the probability of an event and its consequence*. However, this definition will be expanded (see Chapter 2).

Description "IS Risk Management": To describe the term *IS RISK MANAGEMENT*, emphasis is oftentimes put on its procedural aspects. For example, the ISO technical report 13335 [2]; the Communications Security Establishment [3] and Stoneburner, et al [4] recognize IS risk management as the process of identifying risks, evidencing a required security level and selecting appropriate security measures.

In contrast, it is necessary to emphasize the decisional aspects of IS risk management. In such an interpretation, IS risk management stands for a structured approach to risk informed decision making which aligns the operation of IS to the company's risk appetite.

1.2 General Security Context of Global Companies

Description "Global Companies": For the purpose of this thesis, *GLOBAL COMPANIES* are understood as consisting of one or more main offices or hubs with a variety of branches affiliated to them. Such companies often operate as entrepreneurs in a technologically advanced market in a globalized and highly competitive environment which relies heavily on complex IS.

In such an environment, when a company intends to enter a new market, to grow or to diversify risks, a business strategy based on mergers and acquisitions may be pursued. Such a strategy may yield heterogeneous IS landscapes with unevenly distributed levels of security along with discontinuous and inefficient security processes. These *DISCONTINUITIES AND INEFFICIENCIES IN GOVERNING THE COMPANY* generate vulnerability to Internet attacks. Such companies may turn out to represent a patchwork of smaller companies each with a distinct risk profile and risk appetite.

While the potential for vulnerability in IS has risen, the number of Internet attacks has also increased. For example, customers of companies in the finance industry are being victimized by an increasing number of so called phishing attacks, which erode trust in the Internet channel. Another example is the growing number of cases of industrial espionage by means of the Internet, which are fuelled by the aggressiveness of companies and governments while pursuing their goals in the globalized market.

According to Leiner [5], the actual emergence of modern IS occurred in 1969 when computers were first connected via the Internet by the *ARPANET*[3] project. During the late 1990s, the *WORLD WIDE WEB* (www) enabled simple and intuitive navigation, which expanded to become the most popular component of the Internet. Today, the Internet is a critical infrastructure which, in contrast to other critical infrastructures, grew extremely fast and serves global companies as major component for their *COMMUNICATION* (e.g., e-mail, voice over IP, etc.).

Dunn and Wigert describe the incapacitation or destruction of a critical infrastructure as having a debilitating impact on the national security and/or on the economic and social welfare of a nation [6]. According to Wejinen [7], a critical infrastructure shapes social changes at a much broader and complex level so that *the unintended consequences of infrastructure may be much larger than the outcomes for which it was designed*.

As the original idea was to make the Internet work for a small community of computer experts, engineers, scientists, and librarians it was designed with no or low security considerations in mind. Consequently, it offered insecure communication services such as the *SIMPLE MAIL TRANSFER PROTOCOL* (smtp), the unencrypted *FILE TRANSFER PROTOCOL* (ftp) and remote computer access (*TELNET*). As of today, despite the lack of security in its original design, the Internet accommodates a myriad of business applications. This has favoured organized Internet crime, which is highly underestimated and largely unknown in public. Freeh [8] and the U.S. Government Accountability Office [9] note in their reports to the U.S. Congress that the actors involved in organized Internet crime are company insiders, hackers, virus writers, criminal groups, terrorists or foreign intelligence services while the main victims are governments, businesses, and the private sector. In light of the above, concerns have increasingly been voiced, e.g., by Kröger [10]. Skoudis [11] identifies two general trends leading to the *Golden Age of Hacking*:

[3] Advanced Research Projects Agency Network (ARPANET) under the sponsorship of the United States Department of Defense

1. The first trend is concerned with the increased number of machines vulnerable to the same types of attacks. This is due to the homogenization of IS environments, which increases the interoperability of the Internet but leads to a decreased "biological" diversity
2. The second trend is concerned with the wide availability of automated hacker tools which have significantly enlarged the hacking community

1.3 Attackers and Attacks

1.3.1 Attackers

According to Schneier [12] and Schwarz [13], attackers are categorized along their *RESOURCES*, the *KNOWLEDGE* required for the attack, their *MOTIVATION* and *LOCATION*:

Description "Resources": According to research mandated by the European Union, attackers include *NATIONAL GOVERNMENTS, MILITARY AGENCIES* and *INDUSTRIAL SPIES* (Schmid [14] and Cachin et al [15]). According to the *Weltwoche* [16] and the *Frankfurter Allgemeine* [17], both renowned periodicals in German speaking countries, recent attacks aimed at spying on Swiss companies or to deny the service of government sites in the Georgian republic. Such powerful adversaries have plenty of resources.

Description "Knowledge": Attackers can be viewed as amateurs or professionals. According to Webroot Inc. [18], while amateur hackers are responsible for a substantial fraction of Internet threats, the increasing involvement of *PROFESSIONALS* and *ORGANIZED CRIME* is suggested by the rise both in the number and sophistication of attacks motivated by illegal profit. For example, Fyffe [19] reports that as of March 2008, an address costs $ 0.50 while bankruptcy details are worth up to $26.50 on the black market.

Description "Motivation": According to Shinder [20], the motivation for electronic attacks are abundant and ranges from fun and fame to extortion, profit, espionage, revenge, or a political agenda. Attacks for political purposes have become of growing concern since international attention has turned to terrorism while attacks oriented towards the *INVASION OF PRIVACY* or *IDENTITY THEFT* is a growing trend.

Description "Location": Attackers may be either internal or external to a company. Serdiouk [21] and Fyffe [19] report that 70% to 80% of successful attacks stem from *INSIDERS* rather than outsiders. Insiders may have the trust and knowledge of the target organization, which increases the chances of a successful attack. Moreover, in contrast to external attackers, insiders do not have to overcome perimeter defences.

1.3.2 Attacks

Attacks are composed of individual threat actions (also called attack vectors). **Table 1.1** gives an overview[4].

Threat Action	Description
SNIFFING	Sniffers eavesdrop on data transmitted over a local area network (LAN), e.g., user names or e-mails.
SESSION HIJACKING	Session hijacking compromises a user's remote login session, thus providing an attacker unauthorized access to a machine with the privileges of the legitimate user.
PASSWORD ATTACKS	*BRUTE FORCE ATTACKS* verify every possible password until the right one is found. *DICTIONARY ATTACKS* derive password candidates from a dictionary. More sophisticated attacks combine both techniques.
SOCIAL ENGINEERING	The attacker interacts directly with the victim to trick him into revealing personal information or user credentials.
EXPLOITS	An exploit is a piece of software that misuses vulnerabilities, e.g., in a program to cause unintended or unanticipated behaviour.
TROJAN HORSES	Trojan horses appear to be useful software. However, they contain malicious software which hides its existence from the user. It usually

[4] For a more detailed list see Chen and Davis [22].

Note: Wherever feasible, "references" as opposed to "literature" are reported in footnotes as they merely indicate secondary thoughts which are not used to evolve the contents of the text.

	allows an attacker to access remotely the victim's computer.
ADWARE and *SPYWARE*	Adware monitors and profiles online behaviour for the purpose of marketing and is often installed without the user's knowledge. In contrast, spyware records the user's activities. It records keystrokes, visited websites, screenshots, etc. and communicates this information back to the originator of the spyware.
WORMS and *VIRUSES*	Worms are stand-alone programs that replicate by spreading copies of themselves to other systems over a network. Worms may execute their payload, which opens backdoors for remote access, installs spyware, disables anti-virus software, etc. In contrast, viruses depend on the execution of a host program, which they modify (infect) with a copy of themselves. In this sense, viruses are not stand-alone programs.
SPAM EMAIL	Spam is the e-mail equivalent of unsolicited junk mail.
DENIAL OF SERVICE ATTACKS	A denial-of-service (DoS) attack aims at making a computer resource unavailable to its intended users. A commonly used technique involves flooding the victim's machine with a great amount of external communication requests, such that it cannot respond to legitimate traffic any more.
PHISHING (password fishing)	Phishing attacks aim at acquiring personal information from victims by adopting social engineering and/or technical subterfuge. The personal information is then misused to commit, e.g., an illegitimate money transfer from the victim's bank account for the benefit of the attacker.

Table 1.1: Internet Threats (Incomplete Overview).

1.4 Motivation

Managing IS risks promises an optimized allocation of resources to the companies exercising it. It balances the company's capability to realize profits with the expenditures required to pursue a "secure" course of action. This has been recognized by corporate decision makers. However, although much effort is put into an allegedly rigorous IS risk management, in light of closer examination, the efforts in many companies turn out to be

compliance management, commonly understood as the identification and closure of missing security mechanisms prescribed by some best practice or security standard.

For example, companies subject to the Sarbanes-Oxley Act [23], a US law prescribing the application of minimum security requirements for elaborating financial year end statements, display a best-practice-compliant rather than a risk-responsive business conduct when securing their IS. Other legislation such as the Gramm-Leach-Bliley Act [24], which was conceived to prevent money laundering, corruption, and insider dealing in global institutions as well as national and supra-national data protection laws reinforce this tendency. Examples include the Data Protection Directive of the European Union [25], the Bundesdatenschutzgesetz [26], the Tutela delle Persone e di altri Soggetti rispetto al Trattamento dei Dati Personali [27] and the Health Insurance Portability and Accountability Act [28]. In addition, important IS security standards also favour a compliance management approach to managing IS risks [29-31].

Ultimately, compliance management aims at ensuring a law-compliant, ethical behaviour of employees. It is the pre-condition for global companies whose employees are involved in fraudulent acts to hope for milder sentences in case of a court of law. Geissler [32] notes that compliance management has experienced a rapid growth in companies since the 1980s. This rise is disconcerting as it forces companies to emphasize costly control processes to reduce or discharge potential liability. Particularly in European companies, Dummer [33] has noted concerns being voiced as to the inefficiency of compliance management in reducing “criminal energy”; compliance management should not replace IS risk management.

In contrast, given the market-orientation of organized Internet crime, a widely underestimated threat to companies and individuals, it becomes evident that Internet attacks have matured far quicker than approaches counteracting them. In short, while organized Internet crime implicitly adopts a risk-based approach for its own purposes (and consequently, has developed an industry-like efficiency), companies explicitly have not.

In the author’s opinion and personal experience, compliance management constitutes no durable basis for securing IS in global companies, because it focuses on simple questionnaires such as the ISO 17799 [34] and the Security Health Check by the Information Security Forum (ISF) [35]. The following problems are identified:

- the *exhaustive size* of the questionnaires requires a considerable effort to perform

- the *unknown relative importance* of the individual items of the questionnaire does not allow to prioritize which security gaps to address first and
- the *inconsistency of ratings* provided by a team of assessors and interviewees reduces the overall quality of the assessment results

A non-representative study[5] performed by Salvati [37] on the maturity level of IS risk management at one global bank showed the following: the compliance-based approach to IS security indeed ensures a minimum process discipline but a significant potential for projects to exceed cost and time estimates remains because the deliverables of risk analyses executed for the projects can be challenged. In addition, the results of such risk analyses are not necessarily reproducible and may yield resources being wrongly allocated or not allocated at all.

[5] The study was performed in terms of the Capability Maturity Model (CMM) a process model for large-scale software projects. CMM [36] was used to classify individual steps (processes) of IS risk analyses according to their maturity. The processes were rated with *level two* (repeatable) as opposed to *level three* (defined) because they are *not* qualitatively predictable.

1.5 Everglades of IS Risk

Burrell[6] and Morgan [38] assert that it is important to understand the theoretical concepts that form the basis of any methodological approach. Such an understanding indicates the underlying (philosophical) assumptions to a theory and smoothens the way for its systematic use. Applying theoretical rigour is important for risk because it is interpreted differently throughout various scientific disciplines. Althaus [39], for example, describes risk through the looking glass of engineering sciences and economics:

- engineering sciences understands risk as an objective reality, e.g., the potential catastrophe triggered by a severe accident in a nuclear power plant. This risk is measured, controlled and managed by the application of knowledge, remedial action or anticipatory measures
- in economics, risk is a decisional phenomenon, which is perceived as a mix of challenge and security: for entrepreneurs, to take the voluntary business risk for working a market means to be rewarded in case of success (reward-paradigm) while involuntary risks should be insured (insurance-paradigm). In the economic sense, diligently taking risks secures wealth and avoids loss.

Unfortunately, some approaches to IS risk have adopted the engineering interpretation while others have adopted the reward and insurance paradigms of economics without taking into consideration its precise epistemological foundation. For example, the approach *RETURN ON SECURITY INVESTMENT* (ROSI) obeys the reward-paradigm of economics and suggests *INVESTING* in security mechanisms if the Internal Rate of Return (IRR) or the Annualized Loss Expectancy (ALE) are favourable[7]. Although this approach is common in economic theory, it is only partly appropriate in IS as there is no market for IS risks. It is not possible to lower an organization's IS risk by trading it because IS risk is inherent to the peculiarities of the individual company. Consequently, the reward-para-

[6] Burrell and Morgan are the founders of a well known 2 by 2 matrix to classify sociological theories.

[7] For a proposed practical application see the Information Security Forum [40].

digm fails and is hardly used in practice as a non-representative study[8] by Altorfer [41] suggests: the study found that ROSI calculations for justifying security expenditures is applied only in 9% of the companies.

Moreover, Siponen and Willison [42] analyzed IS security literature for the period of 1990-2004 and found – among other results – that IS security lacks empirical research based on theories and that research methods are mostly subjective-argumentative. Consequently, practitioners *and* researchers lack a consistent overall IS risk model. As a result of the lack of conceptual rigor and based on the author's practical work experience at a global financial institution, four areas of discomfort in today's IS risk management have been identified. These areas are determined by the Ambiguity, Likelihood, Influence and Decision Problems.

1.5.1 Ambiguity Problem: Neutralize Ambiguity in Expressing IS Risk

According to the Oxford Online English Dictionary [43] ambiguity expresses the *capability of being understood in two or more ways*. Unfortunately, in IS risk practice, ambiguity prevails although the emergence of IS risk management and dates back to the mid-1970s [44]. According to Salvati and Diergardt [45], the reasons for this widespread ambiguity are manifold:

1. As hinted above, various interpretations of IS risk have been "imported" from other disciplines and industries of which two have been mentioned: engineering sciences and economics
2. Establishing a generally accepted risk terminology represents an incomplete endeavor as IS are continuously developing. For example, the young discipline of IS started off as comprising computer-based systems only. However, this limited interpretation evolved in time as IS are changing society at a cultural level. In light of this, it is legitimate to assume that the scope of IS risk has undergone similar fundamental changes
3. Finally, much of the ambiguity arises because today's standards and best practices do not offer methodological support for the inclusion of the *CONTEXT* where the IS operates in (e.g., the business context the IS supports)

[8] The study was conducted in 2006 and covered 43 companies of different sizes and industries.

1.5.2 Likelihood Problem: Increase the Confidence in Probability Estimates

The Likelihood Problem of IS risk management states that only few probability distributions have been analyzed and agreed upon. For example, the probability distributions of the following attacks have hardly been modeled:

- a brute force attack overcoming an encrypted passwords file
- a virus overcoming an anti-virus system
- a phishing attack overcoming the security awareness of users
- a spyware attack overcoming the detection mechanisms of computers.

Advances in estimating probabilities are hindered by inappropriate interpretations, ill-specified concepts and terminology. For example, in the field of IT security, threats and security mechanisms are treated as two distinct areas of research. This is incomprehensible when taking into account that a brute force attack triggers a different response from an encrypted password than a dictionary attack applied to the same encrypted password.

1.5.3 Probability Influence Problem: Determine the Influence of the Context

IS are not operated in an isolated environment. In fact, they are maintained by humans executing security processes. This context influences the security mechanism which, in turn, influences probabilities. For example, the encryption quality of a specific security mechanism may be influenced by periodic security awareness training or the diligent execution of access management processes. Consequently, determining the influence of the context helps to understand which security processes influence probability. This is called the Probability Influence Problem, or shorter, the Influence Problem.

1.5.4 Decision Problem: Select and Justify Security Mechanisms

In many large companies, the criteria for selecting and justifying security mechanisms has shifted from expert knowledge towards compliance. However, overemphasizing law-compliant behaviour is not desirable in an economy which ideally follows the entrepreneurship-paradigm. Therefore, risk-informed decision criteria are promoted as they better reflect a conduct in line with entrepreneurial business activities.

The four problems are depicted in **Figure 1.1**:

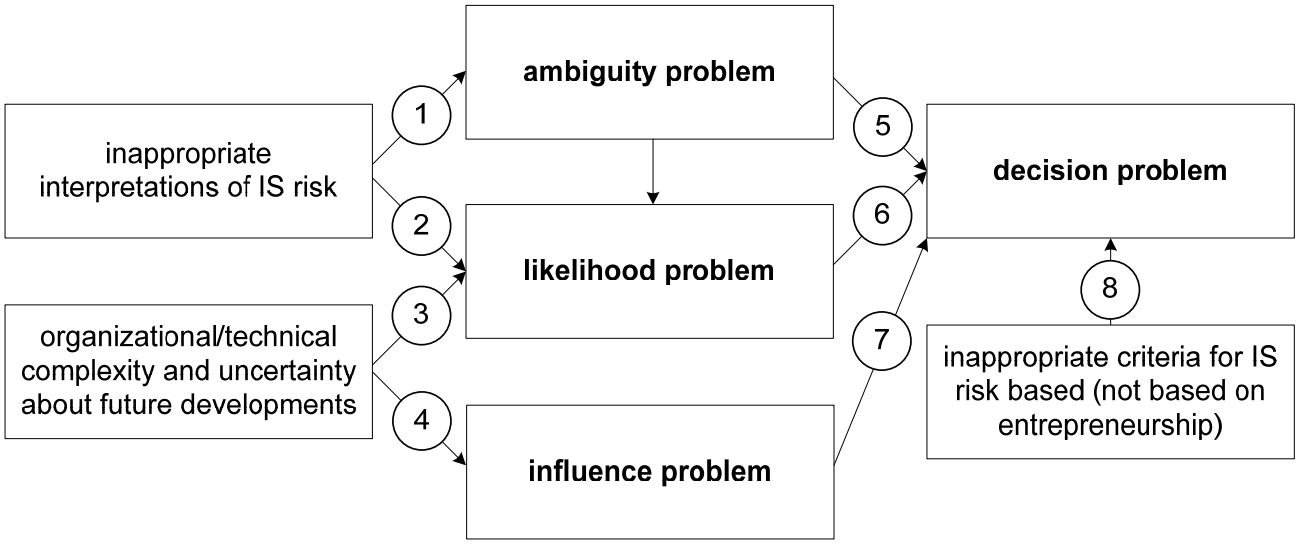

Figure 1.1: Ambiguity, Likelihood, Influence and Decision Problems.

Inappropriate interpretations of IS risk and their related concepts (1) lead to the Ambiguity Problem. Furthermore, they also cause the Likelihood Problem (2) because the terminology, which is grounded on the aforementioned inappropriate concepts, is not fit for a probabilistic approach. The Likelihood Problem is further enhanced by the complexity of IS and the related uncertainty about future developments on technical advances but also new threats (3). Complexity and uncertainty (4) evoke the Influence Problem.

Finally, the Ambiguity (5), Likelihood (6) and Influence (7) Problems contribute to an ill-specified decision situation, which evokes the Decision Problem. In addition, rigorous risk-based decision criteria are hardly applied in practice (8).

1.6 Objectives and Benefit

Given that a consistent and comprehensive modeling approach for risk is absent in today's IS management, the objective of this thesis is to solve the Four Problems by a coherent approach, i.e. by:

1. reducing ambiguity in describing IS risks
2. giving the possibility to measure success probabilities of threats
3. assessing the influence of security processes on probabilities
4. adopting decision criteria based on the entrepreneurship-paradigm.

In order to achieve this, four modules are introduced each addressing one of the above problems. They are the Process Module, the Likelihood Module, the Influence Module, and the Decision Module, which together are called the Four Modules.

The Four Modules benefit the management of IS risks as follows:

1. IS security professionals and corporate decision makers experience clearer communication
2. probability estimates become more credible
3. the influence of the context on probabilities is determined
4. decision makers obtain methodological support in setting up security policies.

1.7 Outline of Thesis

The remainder of this thesis is organized as follows:

Chapter 2 – Risk in Information Systems: The epistemological basis of IS risk is elaborated and the most relevant terms are refined. This foundation serves to infer modeling requirements. This chapter gives an overview on the state-of-the-art of IS risk and is organized along the Ambiguity, Likelihood, Influence and Decision Problems.

Chapter 3 – The Process and the Function Modules: This chapter introduces two modules addressing the Ambiguity and Likelihood Problems:

- the Process Module describes asset states and their changes into subsequent states which have been triggered by (threat) events
- the Function Module offers a generally valid approach to estimate success probabilities of threats

Chapter 4 – The Influence Module: This chapter addresses the Influence Problem by asserting the effect which the context exercises on the probabilities.

Chapter 5 – The Decision Module: This chapter addresses the Decision Problem by supporting decision makers in justifying security mechanisms.

Chapter 6 – Overall Model: In this chapter, the Four Modules are welded together to form one coherent model. It focuses on the interplay among the modules.

Chapter 7 – Case Study on Phishing: This chapter demonstrates the applicability of the integrated model. It is applied to phishing attacks, which aim at stealing personal and sensitive information from oblivious victims. This information is subsequently abused by an attacker for illegitimate personal gain.

Chapter 8 – Conclusions and Outlook: This chapter discusses the scalability of the model to global companies. It draws on the experience gained from executing the Case Study and addresses the challenges related to collecting data. Finally, it touches on the limitations of the approach and proposes further research topics.

2. Risks in Information Systems

In Chapter 2.1, the Nature of IS Risk is explored. The epistemological insights gained from this analysis lay the basis for modeling risk in the remainder of this thesis. These modeling requirements are also inferred through the consideration of the information needs of decision makers and the author's personal working experience in a global company.

In Chapter 2.2, the Risk Terminology Primer is revisited and the terms *THREAT*, *SECURITY MECHANISM*, and *RISK* are refined according to the insights gained from the nature of IS risk.

Finally, in Chapter 2.3, the state-of-the-art in addressing IS risks in selected information security communities is described by taking into consideration the Ambiguity, Likelihood, Influence, and Decision Problems.

2.1 Nature of IS Risks

2.1.1 Risk is a Game of the Mind

IS are interpreted as *SOCIALLY CONSTRUCTED ARTIFACTS* as their use invokes complex social processes. This interpretation contrasts with the notion of risk being a natural phenomenon. A natural phenomenon is expressed by a law of nature where all terms have a precise meaning. Conversely, choosing "man-made artifact" as the leading principle for interpreting IS risk rather than "natural phenomenon" has far-reaching implications: in an IS risk analysis, relevant threat events, their granularity and scope are selected rather than given and depend on the decision maker's perception of the environment. Moreover, as risk in IS is a mental construction, there is no need to enumerate all threats, which establish risk as an absolute figure. This means that the selection of threats and events to be investigated remains at the discretion of the decision maker. Accordingly, IS risk becomes a relative measure (relative to the decision maker). Consequently, adopting a scenario-based approach to modeling IS risk in order to corroborate the decision maker's

perception is advocated. Scenarios are introduced and described in the Risk Terminology Primer (Chapter 2.2).

Assumption 2.1 "Artefact": Risk is a socially constructed artefact, a game of the mind.

2.1.2 Bridging Engineering and Economics

In practice, IS bridges two worlds. One is the engineering world, which produces technical components of IS such as computers, firewalls, and routers. The other world is economics, where these technical devices support a company in doing its required business. Each world pursues its own risk paradigm which eventually accounts for specific aspects of IS risk. Each paradigm operates in its own context, namely, a systems context (engineering) and a business context (economics). To exemplify, imagine two attacks operated by different attackers. The attacks exploit the same vulnerabilities of an operating system, i.e. they cannot be distinguished in the system context; however, in the business context, the consequences of the attacks are entirely different. Different consequences mainly depend on the attacker, e.g., the motivation or the means at his/her fingertips.

Assumption 2.2 "Context": Each risk paradigm operates in an own context.

2.1.3 Interpretation of Probability in IS Risk

Turning to probability, two distinct interpretations have emerged: the frequentist and the subjectivist schools[9], see Bernstein [51] or Althaus [39]. The *FREQUENTIST*[10] school

[9] Other interpretations of probability exist such as classical probability, which is based on a finite number of equally likely outcomes (Marquis de Laplace [46]), logical probability (also called epistemic or inductive) where logical principles determine a unique probability for any body of evidence (Carnap [47]) and propensity probability where the probability of one single outcome is of interest (Popper [48]). See Hayek [50] in the Stanford Encyclopedia of Philosophy for an overview.

infers probability from observing relative frequencies in experiments or proportions in populations upon the repeatability and stability of experiments. The frequentist does not infer probability by observing one experiment only but by observing them all. Accordingly, probability does not relate to a single event but to all events. Having gathered the data pertaining to the experiments under investigation, a research hypothesis is formulated (for example, the failure rate of component *X* is positively influenced by component *Y*). The significance of the hypothesis is then tested by formulating a null-hypothesis (e.g., the failure rates do not depend on component *Y*). If the null-hypothesis is rejected, the research hypothesis is assumed to be true with the probability *P* (also called *P*-value).

In the *SUBJECTIVIST*[11] school (also called the Bayesian School), probability is a subjective degree of belief of an individual or group. Hence, expert judgment is added to the experimental data as an additional source of knowledge. Expert judgment reflects the ability to skillfully make assumptions about uncertain events. In contrast to the frequentist, the subjectivist does not necessarily require the experimental data to infer conclusions; in fact, it is necessary to rely on the user's current state of knowledge. Consequently, the subjectivist focuses on the representation of uncertainty related to an expert's opinion and on how to assess its quality and benefit of use.

The frequentist and subjectivist interpretations of probability are subject to contentious debates[12]. In the frequentist case, for example, the main criticism focuses on the general approach for testing statistical hypotheses where the probability with which a null-hypothesis is validated can arbitrarily be influenced by choosing an "adequately" large sample size. Other shortcomings are the large amount of raw data required to infer conclusions and that experiments must be performed under the exact same conditions. In the subjectivist case, the main criticism is centered around the fact that a single probability assignment cannot convey how well-grounded a personal belief is, i.e. how much evidence is available to support it, e.g., see Howson [59].

---

[10] Contributors to the frequentist approach are, e.g., Fisher [52], Pearson [53] with the "Neyman-Pearson lemma" and Neyman [54] with the "confidence interval".

[11] Representatives of the subjectivist school are Savage, Koopman or Abraham Wald; for reference material see Cooke [55] and Morgan, Henrion & Small [56].

[12] For an overview on the above criticism the reader is referred to the North Prairie Wildlife Research Center [57] or Kain [58].

For IS, both the frequentist and the subjectivist approaches are appropriate. On one hand, the frequentist interpretation is adopted to calculate the probability of threats overcoming security mechanisms where plenty of data is available (or can be generated). Examples include brute force attacks attempting to crack encrypted passwords files; viruses circumventing anti-virus software or stochastic failures of electro-mechanical components or software. On the other hand, the subjectivist interpretation of probability is used where personal conviction is expressed in degrees of belief or where personal preferences are expressed by preference probabilities. An example is the subjective probability that industrial espionage will take place against a company.

Assumption 2.3 "Frequentists and Subjectivists": Both the frequentist and subjectivist interpretation of probability have their legitimate use in IS risk. While the former is used where an **(a)** abundant amount of data is available, the latter expresses **(b)** personal conviction or preferences.

2.1.4 Interpretation of Frequency in IS Risk

Threats directed towards IS can repeatedly harm the same data asset. For example, out of a group of attackers threatening the same document file in an IS, an individual attacker can copy the document file independently from another attacker if no suitable security mechanism is activated (e.g., access control). This yields a situation where different attackers have the ability to repeatedly copy the same document file and then misuse it for their own purposes. This repeatability of attacks is a peculiarity of IS, which decision makers must take into account during the selection of security mechanisms.

This notion of frequency to a threat points the decision maker to the appropriate risk mitigation strategy. For example, threats like viruses and phishing attempts have a high frequency of occurring. As such they represent a known adversary which is an indication to the decision maker that an established process-based security mechanisms like patch management or the execution of periodic security awareness training will be required. Conversely, threats like earthquakes with a low frequency of occurring represent an unknown adversary. They can be treated by rolling-back IS into their original state or by adopting business continuity measures. Consequently, the frequency with which a threat

occurs must be distinguished from the (subjectivist or frequentist) probability that the threat will overcome a security mechanism.

Assumption 2.4 "Frequency and Probability": The notion of frequency is to be distinguished from probability. Frequency points to the appropriate mitigation strategy.

2.1.5 Indivisibility of Threat and Security Mechanism

Threats and security mechanisms mutually presuppose each other. For example, the outcome of brute force attacks attempting to crack encrypted files depends on the computational power of the attackers (the threat) but it also depends on the length of the password used for encryption (the security mechanism) among other things. Accordingly, in order to infer probabilities, a threat and its associated security mechanism need to be viewed as *ONE ENTITY*, one indivisible pair. Consequently, vulnerability does not reside in security mechanisms but rather results from their interplay with the associated threats.

Assumption 2.5 "Indivisible Entity": Threats and security mechanisms mutually legitimize their existence. To produce probability estimates, they must be regarded as one indivisible entity.

2.1.6 Decision Makers and Decision Making

A decision maker concerned with the selection of security mechanisms and the justification of their costs is not required to understand threat events in detail. As a result, omitting "superfluous" interactions among threats and security mechanisms greatly reduces complexity. Consequently, models capable of expressing intermediate system states resulting from dynamic, looping or cascading interactions, for example, are not of interest. For decision making in IS risk it is sufficient to display an asset's initial state and its end state, which are associated by a threat event.

In addition, since the decision makers act as entrepreneurs representing the interests of their companies, it is proposed to consider their personal risk preferences for making decisions. Consequently, the authoritative paradigm used for the selection and justification of security mechanisms should be one that requires acting according to the market's

assumed response to a proposal made by the entrepreneur. In such a business context, an entrepreneur may be right in his/her beliefs of what works in the market or not. Accordingly, arguments for the selection and justification of a security mechanism originating from within the system context are taken into consideration, but are not decisive enough to secure IS.

Assumption 2.6 "Decision Makers": **(a)** To reduce complexity in decision making only the initial state and end state of an asset in a business context are of interest. **(b)** A decision maker represents the interests of a company.

Finally, the above insights on the nature of IS risk are displayed in **Table 2.1** along with its resulting modeling requirements:

Nature of Risk	Modeling Requirements Resulting from the Nature of IS Risk
Assumption 2.1: IS are interpreted as social systems; IS risk is a mental construction.	A decision maker selects the scope and granularity of IS risk analyses according to scenarios which are of interest.
Assumption 2.2: IS bridges the worlds of engineering and economics.	Distinguish between different contexts - in particular, adopt a system and a business context.
Assumption 2.3a: The engineering view interprets IS as consisting of technical devices.	Interpret probabilities as defined by frequentists, i.e. consider all experiments within a specified period of time.
Assumption 2.3b: The social view interprets IS as consisting of individuals or groups.	Interpret probabilities as defined by subjectivists, i.e. consider the decision maker and expert opinion (degrees of beliefs and preference probabilities).
Assumption 2.4: Threats harm data repeatedly if no security mechanisms are installed.	Adopt frequency as a concept to guide the strategy of selection of security mechanism.
Assumption 2.5: Threats and a security mechanisms represent one indivisible entity.	Adopt a probabilistic concept for probability estimates which explicitly considers threats and security mechanisms as one indivisible entity.
Assumption 2.6a: Decision makers are not required to understand threat events in detail.	Only events that associate initial and end states of assets are displayed in the decision making process. Intermediate system states and loops are omitted.
Assumption 2.6b: Decision makers represent the interests of a company.	Risk preferences in the business context are proposed as the decisive criterion for making risk-based decisions in the system context.

Table 2.1: Summary of Modeling Requirements for IS Risk.

2.2 Risk Terminology

2.2.1 Risk and Scenario[13]

In the English language, the term risk has predominantly negative connotations. For example, the Oxford English Dictionary [43] refers to risk as *hazard, danger; exposure to mischance or peril* or *the chance or hazard of commercial loss*. Wharton [61] attributes the origin of the term "risk" either to the Arabic *RISK* or to the Latin *RISICUM*. The former denotes a fortuitous and favorable outcome while the latter denotes an equally fortuitous but unfavorable event: according to the Encyclopedia Britannica [62] it means *to dare something*. Risk is always oriented towards an uncertain future and Bernstein [51] notes that *the revolutionary idea that defines the boundary between modern times and the past is the mastery of risk: the notion that the future is more than a whim of the gods and that man and women are not passive before nature*.

The term "risk" has various semantic meanings. In scientific disciplines, risk is quint-essentially defined as a function of the consequence C of an undesired event E and its probability P, which is perceived as specified by Kolmogoroff [63] and refers to the occurrence of future events.

$$Risk = f(P(E), C(E)) \; . \qquad (\mathbf{2.1})$$

Kolmogoroff defined the probability P by three axioms[14] to lay today's generally accepted mathematical language of uncertainty. In principle, although the consequence C

[13] Except for equation **(2.3)**, the statements in this section are largely drawn from Salvati and Diergardt [45] and Diergardt [60].

[14] (1st axiom): the probability P of an event E of the sample set Ω is a non-negative real number,

(2nd axiom): the probability P that some elementary event in the entire sample set Ω will occur is 1, and

may be positive, it mostly refers to unwanted events and characterizes damage. A commonly applied form of this definition, e.g., in insurance practice, specifies risk as the product of the consequence *C* of an undesired event *E* and its probability *P*, i.e.

$$Risk = P(E) \cdot C(E) \ . \qquad (\mathbf{2.2})$$

For IS, the above actuarial definition used within the insurance-paradigm of the economic world (see Chapter 1.5) has three drawbacks. First, for risk-neutral measures, it puts a low-probability, high-consequence event on par with a high-probability, low-consequence event (e.g., *10*100 = 100*10*). Because these two types of events are different by their very nature valuable information is lost through the multiplication of probability and consequence. In IS practice, the first event is typically related to a business continuity incident (low probability, high consequence) while the second event usually relates to production problems (high probability, low consequence). Second, choosing probability as the sole measure to describe risk implicitly presupposes that the undesired event happens only once. Third, as risk is a mental construct, it cannot be sufficiently described by equation **(2.2)**.

The above shortcomings can be overcome by adapting a scenario based on the description of risk developed by Kaplan and Garrick [64] in 1981. In their definition, risk is a set of triplets consisting of *SCENARIO*, *PROBABILITY*[15] and *CONSEQUENCE*.

Definition "Risk": As distinguishing between frequency and probability is important for finding a general protection strategy (**Assumption 2.4**), *FREQUENCY* is explicitly included in Kaplan and Garrick's definition to describe risk as a set of quadruplets:

$$Risk = \left\{ \langle S_r, f_r, p_r, c_r \rangle \middle| r = 1,2,...N \right\} \qquad (\mathbf{2.3})$$

(3rd axiom): the probability *P* of any countable sequence of pair-wise disjoint events $E_1, E_2, ...$ is calculated by summing up the individual probabilities (also called σ-additivity).

[15] According to Kaplan and Garrick [64] probability is *a numerical measure of a state of knowledge, a degree of belief, a state of confidence* while frequency *refers to the outcome of an experiment of some kind involving repeated trials*.

where
- S_r is the r-th scenario
- f_r is the frequency in S_r of a threat occurring
- p_r is the probability in S_r of a threat overcoming its opposing security mechanism
- c_r is the consequence with respect to S_r

In an IS context, the goal of risk analysis is to determine whether the existing security mechanisms suffice with respect to scenarios specified by the decision maker. A scenario is an undesired course of action within a system caused by one or multiple triggers, which force its initial state into an end state. To describe a scenario, a discrete-event (event sequence) model is adopted as it is typically the most effective approach to describe the behaviour of complex IS, see, e.g., Diergardt [60]. **Figure 2.1** illustrates a generic scenario triggered by an initiating event, which may go through various intermediate states until it reaches an end state.

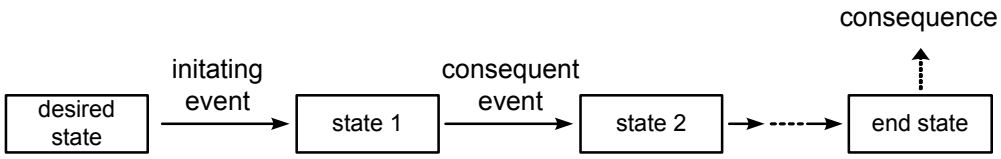

Figure 2.1: Composition of a Generic Linear Scenario.

The event-state pairs do not comprise all the possible events and states, but only the ones relevant for the specific scenario. Statemelatos [65] refers to those pairs as pivotal. They comprise occurrences between two system states which occur immediately or evolve in time.

So far linear scenarios have been elaborated, which consist of a single initiating event, a single end state and a sequence of event-state pairs in between. However, this linearity does not necessarily apply for all scenarios. Typically, scenarios have several end states (e.g., representing the possibility of different outcomes) and assume the possibility of concurrent events. This is depicted in the following graph:

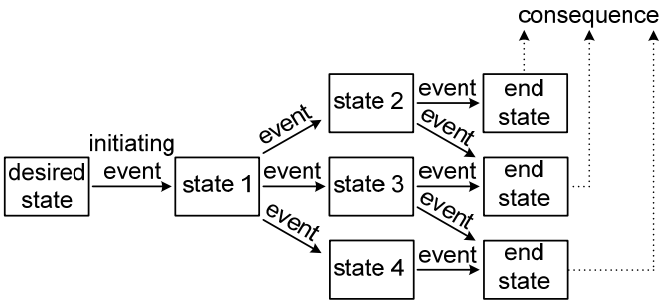

Figure 2.2: Scenarios with Multiple End States.

To model a scenario, two techniques are used: process modeling and function modeling (**Figure 2.3**). On the one hand, process modeling describes the course of events of an IS. Examples of process modeling techniques are event trees, cause consequence diagrams or event sequence diagrams. On the other hand, a function modeling technique describes the function oriented aspects of a system. In Chapter 3, both the process and the function models are adapted for use in IS.

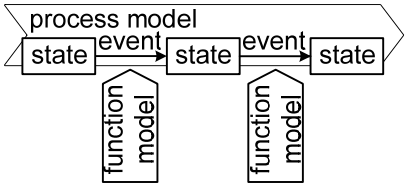

Figure 2.3: Scenario as a Combination of a Process Model and Function Models.

2.2.2 Threat

Definition "Threat": According to Kröger [66], a threat is a danger or hazard directed towards an asset. The Oxford English Dictionary [43] refers to danger as *power to inflict physical injury* while hazard is *chance, venture* and *risk of loss or harm; peril, jeopardy*. In this work, a threat is a set of tuples consisting of *THREAT AGENTS* and *THREAT ACTIONS*.

$$Threat = \left\langle t_i^a, T_i^c \middle| i = 0,1,...m \right\rangle, \qquad \textbf{(2.4)}$$

where t_i^a is the i-th threat agent t^a, $0 \le i \le m$

T_i^c is the set of threat actions available to the i-th threat agent

Definition "Threat Agent": A threat agent is any entity capable of triggering a threat action, e.g., an employee, a hacker, malicious code, a group of people, governmental agencies. Schwarz [13] describes a threat agent by its *MOTIVATION*, *KNOWLEDGE*, *RESOURCES* and *LOCATION*:

$$t_i^a = \left\langle m_{i.j}^a, k_{i.j}^a, r_{i.j}^a, l_i^a \right\rangle, \tag{2.5}$$

where t_i^a is the i-th threat agent, $0 \leq i \leq m$

m_i^a is the motivation of the i-th threat agent

k_i^a is the knowledge available to the i-th threat agent t_i^a

r_i^a are the resources available to the i-th threat agent

l_i^a is the location of the i-th threat agent relative to the company security perimeter

Schneier [12], Kyas [67] and Eckert [68] describe the extensions of motivation, knowledge, resources and location by:

- playing instinct and appreciation, discovering vulnerabilities, etc. (motivation)
- basic or expert knowledge in, for example, how to operate an operating system or how to execute a program
- low or significant budget, time or computational power (resources)
- internal or external to the security mechanism (location)

Definition "Threat Action": A threat action t^c exploits security mechanisms. Examples of threat actions are probing IP packets, scanning ports, creating password combinations. Threat agents *TRIGGER* a threat action, i.e. they *ACTIVATE* or *EXECUTE* a threat action. To trigger a threat action, a threat agent needs appropriate knowledge and resources.

Definition "Threat Strength": Colloquially, the knowledge and resources of the threat agent describe the *THREAT STRENGTH* of the threat action.

$$t_i^s = \left\langle k_i^a, r_i^a \right\rangle, \tag{2.6}$$

where t_i^s is the threat strength for the i-th agent, $0 \leq i \leq m$

k_i^a is the knowledge available to the *i*-th threat agent

r_i^a is the resources available to the *i*-th threat agent

2.2.3 Security Processes and Mechanisms

Definition "Security Process": *SECURITY PROCESSES* are recurrent activities to set up, maintain and decommission security mechanisms and mostly aim at influencing human behavior. Examples are the execution of security awareness programs, the implementation of change management procedures, defining housekeeping and maintenance of IS. For an overview refer to the International Standardization Organization (ISO) [29].

Definition "Security Mechanism": *SECURITY MECHANISMS* are designed to protect an IS by granting access, encrypting data, limiting or verifying memory boundaries of variables, validating input, etc. Examples are firewalls, encryption measures, authentication and authorization, security patches. For an overview refer to SANS [69].

A security mechanism has *PROPERTIES*, which are exploited by threat actions. An analogy for properties is the security mechanism "pipe[16]", which can be sabotaged in various ways. For example, an attacker (threat agent) applies acid (threat action) to cauterize a hole into the pipe, then applies a hammer to damage it or fire to bend it. The threat actions applied exploit different properties of the pipe namely, its resistance to acid, tension as well as pressure, and heat. A property is comparable to a key lock that determines whether a threat action qualifies (or matches) to exploit it. **Figure 2.4** visualizes the above.

[16] A pipe, in this context, is regarded as a security mechanism: its functionality is to prevent liquids from spilling.

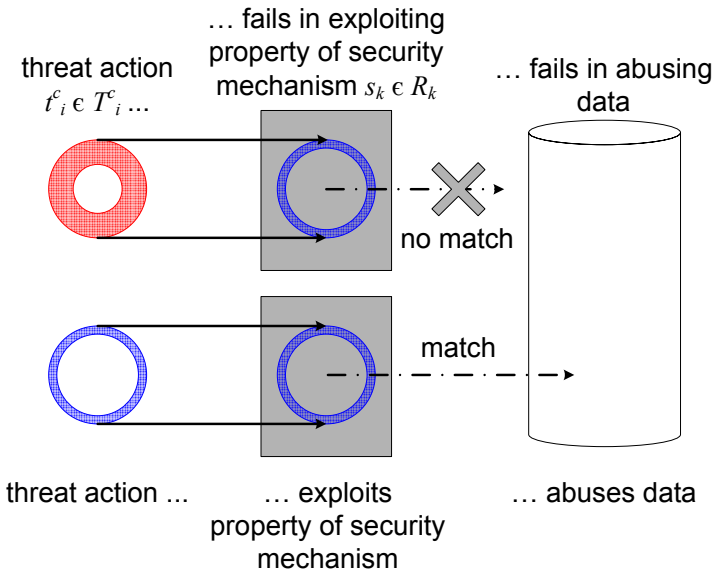

Figure 2.4: Matching Threat Actions with Security Mechanisms.

A security mechanism has two basic functions: First, it resists a threat action within some limits and second, it recovers from the threat action (imagine, e.g., the roll back functionality of a data base or a network recovering from a DOS attack). These two functions are attributed to the *RESISTANCE* and the *RESILIENCE* of the security mechanism respectively. In the pipe example, the properties "resistance to acid", "resistance to tension and pressure" and "resistance to heat" form the resistance tuple. In order to exploit a security mechanism, threat actions must overcome one or more values of its resistance tuples.

Definition "Resistance": The resistance of the *k*-th security mechanism is defined as a set of independent properties, i.e.

$$Resistance = \langle s_k, R_k \rangle | k = 0,1,...,q \,, \tag{2.7}$$

where s_k is the *k*-th security mechanism, $0 \le k \le q$

R_k is the set of properties related to the k-th security mechanism

Examples are the quality of a chosen password such as its length or the maximum number of tries a wrong password can be input. Resistance properties of a security mechanism are threat-dependent. For example, an encrypted password reacts differently to a brute

force attack than it would to a dictionary attack. The resilience properties are not followed up as focus is laid on initial and end states of assets rather than on intermediary states.

Definition "Susceptibility": The Oxford English Dictionary [43] defines susceptibility as the *capability of being, or disposition to be, affected by something*. The set of all threat actions, which potentially exploits a security mechanism s_k, is called the *SUSCEPTIBILITY set* S_k.

A threat agent triggers threat actions with the aim to exploit security mechanisms in order to enlarge the number of threat actions available to him. Threat actions can be used to exploit another security mechanism or to *ABUSE* an asset.

Susceptibility is distinct from vulnerability. Vulnerable stems from the Latin *vulnerare* which means *to wound*. According to Sans [70], vulnerable means *capable of being physically wounded*, *open to attack or damage*. Vulnerability originates from:

- an inappropriate specification of security requirements (e.g., when designing the IS, past, present or future threats were not considered),
- a faulty implementation of security requirements, or
- an unfavourable interaction among security mechanisms.

To counteract one or multiple vulnerabilities, one or more security mechanisms are applied, i.e. they enhance its resistance and/or reduce its susceptibility.

Definition "Perfect Vulnerability" and "Perfect Control": A security mechanism with an infinitely small resistance vector and a maximum susceptibility set displays a *PERFECT VULNERABILITY* towards a specific threat while a security mechanism with an infinitely large resistance vector and an empty susceptibility set is called a *PERFECT CONTROL*.

Vulnerabilities are specific to the interaction between threat and security mechanism. A threat designed to exploit a security mechanism succeeds if its resistance values are insufficient to withstand the threat strength, in which case the security mechanism has an Achilles' Heel, i.e. displays vulnerability. Otherwise, the security mechanism is in control.

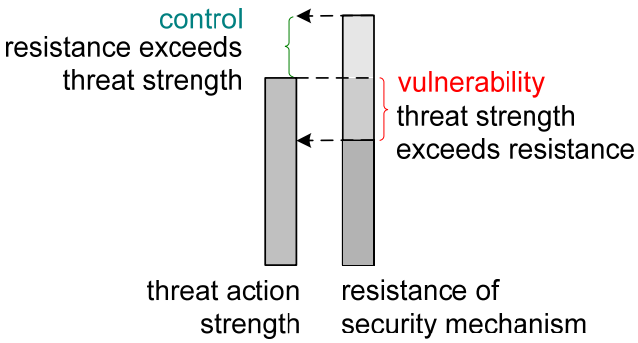

Figure 2.5: Threat Strength and Resistance of Security Mechanism.

To exemplify, **Figure 2.5** shows the threat strength in excess of the resistance of the security mechanism (vulnerability) and the resistance in excess of the threat strength (control). Hence, vulnerability and control are a measure to denote the surplus or insufficiency of resistance given specific threat strengths.

Definition "Vulnerability and Control": *VULNERABILITY* has two components: the first component describes its resistance which is lower than the threat strength, and the second component describes its susceptibility set. Analogously, *CONTROLS* display a resistance which exceeds the threat strength.

2.3 State-of-the-Art

2.3.1 Ambiguity Problem

The Ambiguity Problem encompasses:

Use of inconsistent risk terminology. Many standards or best practices name some threats after missing security mechanisms. For example, the German *IT Grundschutzhandbuch* [31] lists the "non-compliance with IT security measures" as a threat.

Use of inconsistent concepts. The *IT Grundschutzhandbuch* treats threats and security mechanisms as two different entities rather than as an indivisible pair. Other examples include NIST [4], ISF [35], ISO 17799:2005 [34].

Lack of coherent and comprehensive models. Jakobsson's model[17] [71] visualizes threats by using graphs and quantifies risks by economic analysis. However, economic analysis does neither deliver probabilities nor risk preferences for decision making. Another example of non comprehensive models is given by Schneier [72] and Moore, et al [73]. Their models visualize attacks in form of a tree structure[18].

Application of inappropriate modeling techniques. Examples include the occasional use of modeling techniques from the engineering sciences for use in IS risk management such as the Zurich Hazard Analysis [66], Fault Trees[19] [74, 75], Failure Mode and Effects Analysis [76], Markovian chains [77, 78] or Petri Nets [79]. Such modeling techniques do not address the different contexts as outlined by Cobit [30] or Garret and Apostolakis [80] in which threats and security mechanisms operate.

Remark: The insufficient consideration of the epistemological nature of IS is the root cause for the Ambiguity, Likelihood and Decision Problems. It yields inconsistent terminology and a patchwork of risk concepts as well as the adoption of inappropriate modeling techniques. It impedes coherent and comprehensive models.

2.3.2 Likelihood Problem

In contrast to probability, according to Wolfram's Mathworld [81], likelihood *differs from that of a probability in that a probability refers to the occurrence of future events, while likelihood refers to past events with known outcomes*. Accordingly, the Likelihood Problem looks at observations from the past which stand for probabilities denoting uncertainty about future events. While the use of probabilities as defined by Kolmogoroff (see Chapter 2.2.1) is widely accepted in the scientific community, in IS risk there is:

Disagreement on the use of probabilistic concepts/interpretation of probability. For example, in 1999, Garret and Apostolakis [80] question whether probabilistic techniques can be applied to assess risks related to software failing. They promote a context-based view to software risk assessment as failures result from both the input to the software as well as the error-forcing context it is operating in. Accordingly, failures do not occur randomly but as the result of encountering some environment for which the

[17] For more information on Jakobsson's model refer to **Appendix A**.

[18] For more information on Schneier's model refer to **Appendix A**.

[19] According to Lee, Grosh and Tillman [74], Fault Trees were discovered by H. Watson.

software was not properly designed. In contrast, the software reliability community[20] interprets IS as a black box, i.e. internal dynamics are not modeled and the probability of software failing is estimated via the failure rate of software. Pan [83], for example, assumes a distribution along the so-called bathtub curve (see **Appendix A**) which was originally used in engineering sciences; refer to Kröger & Mock [84], Pham [85].

Low confidence in probability estimates by practitioners. Based on the author's personal experience at a global bank, probability estimates are considered unreliable by decision makers and therefore are often disregarded in the decision making process.

Remark: Garret and Apostolakis postulate that if it was possible to analyze every single software failure to find a deterministic common cause there would be no need for a probabilistic interpretation while the approach used in the reliability community gives only limited insights into the dynamics leading to its failure. Both interpretations do not take into account the notion of threats exploiting security mechanisms and that probability is eventually determined by this interaction.

2.3.3 The Probability Influence Problem

Security processes set up, maintain and decommission security mechanisms, and influence their interaction with threats, i.e. the context in which an IS operates, influences the success probabilities of threats. Accordingly, security processes such as patch management and security awareness training are important for the selection and justification of security mechanisms. Unfortunately, the influence of security processes on security mechanisms is unknown in practice as the methodological basis to approach this problem is missing. Practitioners lack an important piece of information for decision making. Solving the *PROBABILITY INFLUENCE PROBLEM* is appealing as it promises an optimized allocation of resources.

Remark: Influence is measured by relating security processes with the probability of success of a threat. However, executing experiments in a rapidly changing IS context is challenging as the execution of the exact same experiment for two consecutive measurements is almost impossible. Approaches to evidencing relationships between security pro-

[20] For an overview on the field of software reliability refer to Goel [82].

cesses and security mechanisms such as Linear Regression, Principal Component Analysis and Clustering must overcome the aforementioned challenges.

2.3.4 Decision Problem

In practice, a wide variety of approaches have been brought forward for the selection of security mechanisms and justification of their costs. They are based on:

Checklists. Examples include the Code of Practice for Information Security Management [29], which identifies missing security processes, the Guidelines for the Management of IT Security [86], which proposes safeguards based on the type of IT system, security concerns and threats, the IT Grundschutzhandbuch [31], which prescribes security measures based on the type of system.

Economic principles. Examples include the Risk Management Guide for Information Technology Systems by Stoneburner [4], which is based on cost-benefit analyses, the approaches by Verhoef [87] and the Information Security Forum (ISF) [40], which focus around Return on Security Investments (ROSI) by evaluating indicators such as the discounted cash flow (DCF), the net present value (NPV), the pay back period (PBP), the internal rate of return (IRR) or the return on investment (ROI).

Maturity of IS. A prominent example is the Information Security Management Maturity Model by Canal [88] relating security controls to the maturity level of IS.

Compliance requirements. Examples of legal or regulatory requirements include the Sarbanes-Oxley Act [23] or the Basel II Accord [89].

Risk considerations. The Threat and Vulnerability Assessment proposed by the ISF [35] takes simple risk considerations into account.

The academic world has also developed a variety of approaches. For example, the Analytical Hierarchy Process (AHP) by Saaty [90] is used to rank alternatives based on the subjective assessment of objectives via a pair-wise comparison matrix; the Simple Multi-Attribute Rating Technique (SMART) by Edwards [91] views every outcome of an action having a value on a number of different objectives. These values are measured one objective at a time and then are aggregated across objectives. For an overview and an application example, the interested reader is also referred to Olson and Courtney [92]. For an application of Markov decision processes, refer to Goto, Lewis and Puterman [93].

Other approaches[21] focus on the uncertainty and ambiguity of information and use techniques based on fuzzy logic (Gheorghe and Mock [95]), Rough Sets (Pawlak and Slowinski [96]) or hybrids like the evidential reasoning approach (Yang [97], Huynh [98]).

Remark: Regardless of the variety of valid approaches both in practice and the academic world, the selection of security mechanisms and justification of their costs is not (well) integrated into the company decision process. In practice, determining the allocation of resources for IS risk management is generally not approached with sufficient rigour.

[21] For a comprehensive overview on the state-of-the-art in multi-criteria decision analysis see Figuera, Greco and Ehrgott [94].

3. Process and Function Modules

The Process Module provides a graphical notation to describe a chain of events. It visualizes states of assets by the means of scenario techniques which allow for a limited insight into a potential future of those assets. In Chapter 3.1, an overview of the Process Module is given while assets and events are described in Chapters 3.2 and 3.3 respectively. Scenarios are elaborated in Chapter 3.4 and their force of expression is demonstrated by an application example in Chapter 3.5. Chapter 3.6 devises on probabilities in scenarios.

The Function Module calculates event probabilities based on stochastic attack and defence behaviour of threats and security mechanisms. Chapter 3.7 gives an overview of the Function Module while its probabilistic concept is described in Chapter 3.8. In Chapters 3.9 through 3.11, the probability of a brute force attack overcoming an encrypted password file is calculated.

3.1 Overview of the Process Module

The value of an asset depends on the context that surrounds it; a system context and a business context are distinguished. The value of the asset is expressed by its state. The state of the asset is forced into a subsequent state by a threat event. A change of state in the IS context is referred to as *IMPACT* and as *CONSEQUENCE* in the business context. The Process Module provides a graphical notation to describe a chain of events with their related impacts and consequences. Events and assets are chosen by the decision maker according to information needs. To support a decision maker the risk analyst applies methods and techniques for the elicitation and representation of expert opinion[22].

[22] Methods and techniques for the elicitation of expert opinion are beyond the scope of this thesis. The interested reader is referred to Cooke [55], Zgurovskii [99] and Plous [100].

3.2 Assets and Events in a Business and Information System Context

According to the Merriam Webster Online English Dictionary [70], an asset is *an advantage, a resource* but also *the items on a balance sheet showing the book value of property owned*.

Definition "Asset": In a business context, asset is referred to as a producer's good, which creates direct and indirect monetary and non-monetary revenues.

On one hand, the value of an asset depends on the *CONTEXT* it is embedded in, for example, the industry, the company, etc. On the other hand, the value of an asset is also determined by the *PERCEPTION* of its intended use by the decision maker. Consequently, the asset is valued according to its acquisition price, implementation or maintenance costs, the value it has for a competitor, a hacker, etc.

Definition "Asset Value in the Business Context": The value of an asset is expressed in multiple dimensions and is commonly used in financial, reputation or legal settings. They are expressed by a value vector using monetary units[23]:

$$v_a^b = \left(f_a, r_a, l_a\right), \tag{3.1}$$

where v_a^b is the value vector in a business context for the a-th asset, $0 \le a \le b$

f_a is the financial value of the a-th asset (in monetary units)

r_a is the reputation value of the a-th asset (in monetary units)

l_a is the legal value of the a-th asset (in monetary units)

Definition "Consequence": A risk analyst is usually confronted with asset states reflecting a loss in value. The asset value is subject to a loss as *CONSEQUENCE* of an event defined as the difference in monetary units before and after an event:

[23] Monetary units are used for simplicity. It could also be conceived by choosing dimensions which are quantified or qualified in different ways.

$$c_a^b = v_{a.2}^b - v_{a.1}^b \ , \tag{3.2}$$

where c_a^b is the consequence vector in a business context for the a-th asset, $0 \le a \le b$

$v_{a.x}^b$ is the value vector in a business context for the a-th asset prior ($v_{a.1}^b$) and after the event ($v_{a.2}^b$)

Definition "Event in the Business Context": An *EVENT IN THE BUSINESS CONTEXT* is a probability function that describes the change of state of an asset, i.e:

$$t_a^e : v_{a.1}^b \mapsto v_{a.2}^b \ , \tag{3.3}$$

where t_a^e is the threat event forcing the a-th asset to change its state, $0 \le a \le b$

$v_{a.x}^b$ is the value vector in a business context of the a-th asset prior ($v_{a.1}^b$) and after ($v_{a.2}^b$) the event, i.e. $x \in \{1,2\}$

Remark: To avoid complexity, the number of events is limited to one per asset, i.e. to one initial state ($v_{a.1}^b$) and one end state ($v_{a.2}^b$). **Figure 3.1** visualizes the above function.

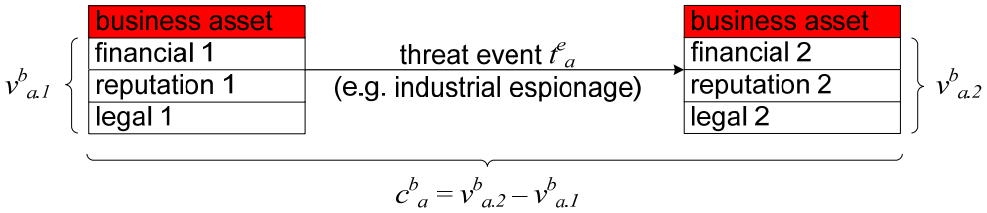

Figure 3.1: Event Forcing a Business Asset to Change its State.

Definition "Asset Value in the IS Context": In the IS context, examples of assets include data and technical IS services. The value of an IS asset is described by the well-known triplet: confidentiality, integrity, and availability:

$$v_a^I = \left\langle c_a^v, i_a^v, a_a^v \right\rangle \ , \tag{3.4}$$

where v_a^I is the value tuple of the a-th asset in an IS context, $0 \le a \le b$

c_a^v is the confidentiality value of the a-th IS asset expressed, e.g., in *public*, *internal*, *confidential*, *secret*

i_a^v is the integrity value of the *a*-th IS asset expressed, e.g., in *having integrity*, *not having integrity*

a_a^v is the availability value of the *a*-th IS asset expressed, e.g., in percentages

Definition "Impact": Let $v_{a.1}^I$ and $v_{a.2}^I$ be the initial and end state respectively of an asset *a* in the IS context. The IS asset is said to be *IMPACTED* by an event:

$$i_a^I = v_{a.2}^I \sim v_{a.1}^I , \tag{3.5}$$

where i_a^I is the impact tuple related to the *a*-th asset of an IS context, $0 \leq a \leq b$ and the operator "~" denotes the change in confidentiality, integrity and availability in words, e.g., "the confidentiality has changed from *secret* to *public*".

$v_{a.x}^I$ is the value tuple in an IS context prior ($v_{a.1}^I$) and after the event ($v_{a.2}^I$)

Definition "Event in the IS Context": Examples include network scanning or account cracking. In analogy to the business context, events in the IS context are represented by a probability function that describes the change of state of an asset, i.e.:

$$t_a^e : v_{a.1}^I \mapsto v_{a.2}^I , \tag{3.6}$$

where t_a^e is the threat event forcing the *a*-th asset to change its state, $0 \leq a \leq b$

$v_{a.x}^I$ is the value tuple in an Information System context of the *a*-th asset prior ($v_{a.1}^I$) and after ($v_{a.2}^I$) the event

Remark: To restrain complexity, the number of events is delimited to one per asset. **Figure 3.2** visualizes this function. Note that the impact is a tuple rather than a vector.

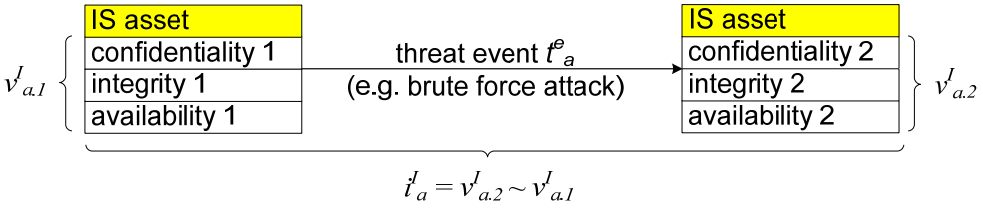

Figure 3.2: Event Forcing IS Asset to Change its State.

3.3 Scenarios in the Process Module

According to Merriam-Webster [70], *SCENARIO* in Italian means *outline or synopsis* and refers to the Commedia dell' Arte where it was used to indicate *a sequence of events, especially when imagined*. As a strategic planning tool, scenario techniques are firmly rooted in the military. In recent years, they have been popularized for long range business planning. In IS risk management, the use of scenario techniques supports decision makers in anticipating suitable security mechanisms and processes. The basic approach to scenario planning is to generate a decision situation in which known facts about threats are combined with trends. When disclosed in advance, future risks may be counteracted with timely actions rather than under the duress of an emergency.

There is wide-spread dissent in scenario literature about definitions and methodological approach[24]. As there is no obligation towards any particular approach, the following is proposed:

- the basic elements which constitute scenarios
- the chart in which scenarios related to the same type of attack or the same type of loss are placed.

3.3.1 Basic Elements

Definition "Scenario Elements in an Information System Context": In an IS context, the basic element of a scenario is defined as a triplet showing:

- the initial state (the initial value) of an asset, i.e. $v_{a.1}^{I}$
- the threat event forcing the change in state, i.e. t_a^e and
- the end state (the end value) of the asset, i.e. $v_{a.2}^{I}$

$$s_{s,z}^{I} =< v_{a.1}^{I}, t_a^e, v_{a.2}^{I} >, \tag{3.7}$$

where $s_{s,z}$ is the z-th element of scenario s, $0 \leq s, z \leq N$

[24] For a historical and foundational overview refer to Bradfield et al [101], for general reading on the subject refer to Fahey and Randall [102].

$v^I_{a.x}$ is the value vector in the IS context related to the *a*-th asset, $0 \le a \le b$, i.e. $x_{\cdot} \in \{1,2\}$

t^e_a is the probability function related to event t^e and asset *a*.

Definition "Scenario Elements in a Business Context": In a business context, the basic element of a scenario is defined as a triplet showing:

$$s^b_{s,z} = < v^b_{a.1}, t^e_a, v^b_{a.2} >, \quad \textbf{(3.8)}$$

where $s_{s,z}$ is the *z*-th element of scenario *s*, $0 \le s, z \le N$

$v^b_{a.x}$ is the value vector in the business context related to the *a*-th asset, $0 \le a \le b$, i.e. $x_{\cdot} \in \{1,2\}$

t^e_a is the probability function related to event t^e and asset *a*.

Definition "Scenario": A scenario is a set containing basic scenario elements, i.e.:

$$S_r = \{s^I_{s,z}, s^b_{s,z}\}, \quad \textbf{(3.9)}$$

where S_r is the *r*-th scenario, $0 \le r \le N$

$s^I_{s,z}, s^b_{s,z}$ is the *z*-th element of scenario element *s* in the IS or the business contexts respectively, $0 \le s, z \le N$.

To connect basic scenario elements XOR and OR gates are used:

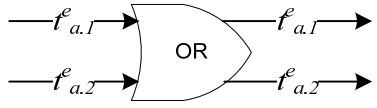	OR means that one or more events $t^e_{a.x}$ activate one of the following arrows. The OR gate shows at least two events connecting the initial state of an asset *a* with its end state(s).
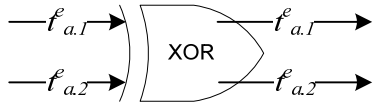	XOR means that one exclusive event $t^e_{a.x}$ occurs. The XOR gate depicts at least two events, which connect the initial state of asset *a* with its end state(s).

Table 3.1: XOR and OR Gates.

To control complexity, events may be branched or concurrent but not looped:

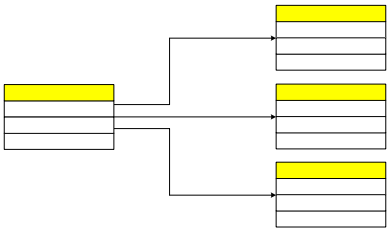	Branched events yield a different state of the same asset.
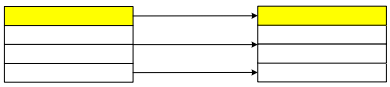	Concurrent events yield the same state of the same asset but via different threat events.
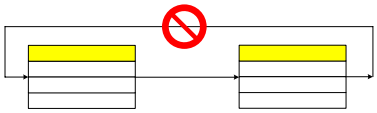	Loops are not allowed.

Table 3.2: Branched, Concurrent and Looped Events.

3.3.2 Scenario Chart

Assumption 3.1 "Closed World": Companies perceive attacks as elements of a *CLOSED WORLD* (closed world assumption).

In logic and knowledge management, the closed world assumption describes elements which are not currently known to be true as false. It has wide applicability in IS risk management as any attack, which is unknown in public, is generally assumed to be false by the vast majority of companies or users as they do not engage in threat research[25].

In a closed world, attacks are described by a coordinate system delimiting the degrees of freedom along which attack vectors evolve. In the IS context, the axes (usually) represent the attack deployment and execution respectively. Evolving attacks along the axes allows for limited foresight into one possible future. A particular scenario is chosen so it is both possible and uncomfortable. The simplest closed world situation consists of two axes showing the reciprocal relative position of five to nine pivotal scenarios. In a scenario chart, variants of the same attack are depicted. For example, phishing attacks[26] can be classified in classic phishing, SMiShing and vishing (refer to Chapter 7).

[25] In contrast, security companies or the military actively engage in threat research.

[26] Phishing attacks represent a form of Internet fraud.

3.4 Example of a Scenario – Jim Cracker

This fictitious example[27] depicts one attack with some alternatives for its deployment and execution. It shows the descriptive qualities of scenario techniques and their capability to reduce complexity. Consider Jim Cracker, a disgruntled employee of a company intending to disturb its business expansion by means of industrial espionage. Jim steals a secret document showing the company strategy for expansion in Europe, Middle East and Africa (the so-called EMEA region) to sell it to a competitor. The following scenario elements are identified:

Scenario Element 1.1: Jim Cracker performs a successful *Brute Force Attack* on the *Passwords File* and the confidentiality of the passwords is lost. The integrity and availability of the passwords file remain unchanged.

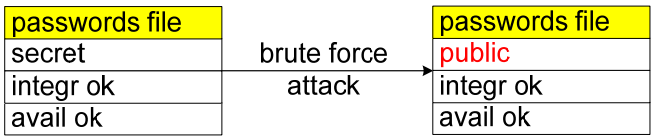

Figure 3.3: Scenario Element 1.1.

Scenario Element 1.2: As an alternative, Jim has the possibility to perform a *Social Engineering Attack* to obtain the passwords.

Scenario Element 2: Jim misuses the passwords to place a *Backdoor* on the firewall of the company network. Hence, the confidentiality and integrity of the *Firewall Configuration* are not preserved. The availability of the firewall remains within acceptable limits.

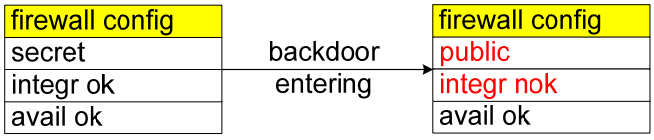

Figure 3.4: Scenario Element 2.

[27] An example of a scenario chart is not given and the reader is referred to Chapter 7.

Scenario Element 3.1: Upon bypassing the firewall, Jim uses the cracked passwords to gain access to the *Internal Network* and starts *scanning* for sensitive material. Jim eventually steals business plans contained in a document called *EMEA.doc*. Consequently, its confidentiality is lost but its integrity and availability are preserved.

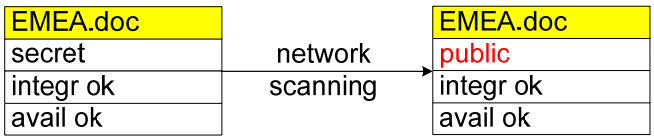

Figure 3.5: Scenario Element 3.1.

Scenario Element 3.2: Alternatively, Jim uses the cracked passwords to gain access to a portion of the network called *Security Zone* where he *scans* for *EMEA.doc*.

Scenario Element 3.3: Alternatively, Jim uses the cracked passwords to gain access to a portion of the network called *Back Up Zone* where he *scans* for *EMEA.doc*.

Scenario Element 4.1: Jim sells *EMEA.doc* to a competitor. As a consequence, the competitor *opens up* a business in the same area as Jim's employer. This affects profits and causes legal litigation among the competitors. The asset *EMEA branch* of Jim's employer shows financial and legal losses of 1,000,000 and 500,000 respectively.

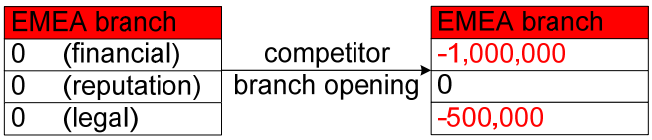

Figure 3.6: Scenario Element 4.1.

Scenario Element 4.2: Alternatively, Jim *unveils EMEA.doc* to a hacker colleague who has no further interest in the document. As a result, there are no consequences in the business context.

Figure 3.7 shows the concatenation of the aforementioned scenario elements. In the IS context, three pivotal assets are found: the *Passwords File*, the *Configuration File* of the Firewall, and the document *EMEA.doc*. In the upper left, the event *Brute Force Attack*

(1.1) is depicted. It forces the initial state of the *Passwords File* into its end state shown in the upper right. This event has a probability of success denoted by $p_{1.1}$.

The concurrent event *Social Engineering* (1.2) realizes with a probability of $p_{1.2}$. It forces the initial asset state of the *Password File* to change. Both the events *Brute Force Attack* and *Social Engineering* deliver the passwords Jim was looking for but only one attack must be performed. This is denoted by the XOR gate.

The event *Backdoor* (2), which is related to the asset *Firewall Configuration*, can only be realized if the *Passwords File* is known, i.e. the events (1.x) serve as a prerequisite. The event has a success probability of $p_{2.0}$.

Once the *Backdoor* is installed onto the firewall, Jim scans for the document *EMEA.doc*. To do so, he faces alternatives denoted by three concurrent events: *File Scan in the Internal Network* (3.1), *File Scan in the Security Zone* (3.2) and *File Scan in the Back Up Zone* (3.3). As only one of the three alternatives must be realized to find the sensitive document *EMEA.doc*, they are collected by an XOR gate. The concurrent events have a probability of occurring of $p_{3.1}$, $p_{3.2}$, and $p_{3.3}$ respectively.

The business context shows the alternative events *Competitor Branch Opening* (4.1) and *Unveiling* (4.2). The two events force the business asset *EMEA branch* into two different end states. The events have a probability of occurring of $p_{4.1}$ and $p_{4.2}$.

The thick arrows show the actual attack by Jim Cracker which has been devised into its deployment and execution. On the right hand side of **Figure 3.7**, the IS context and the business context are shown.

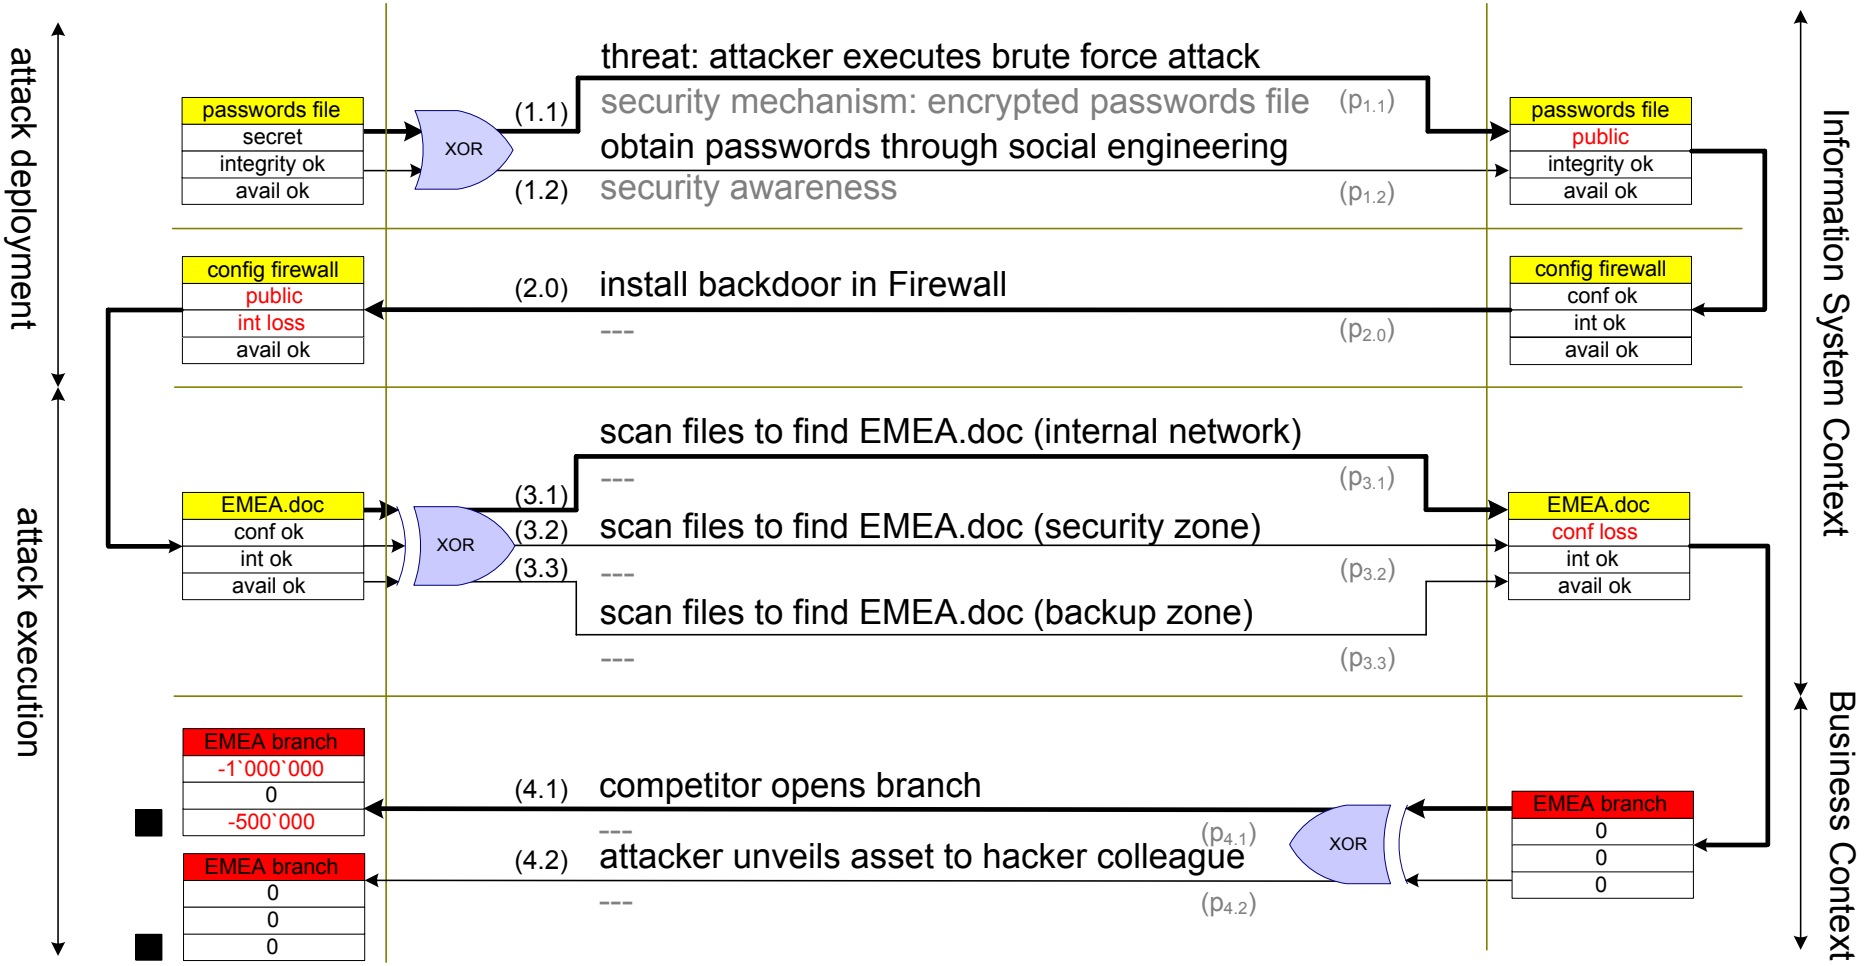

Figure 3.7: Scenarios for the Jim Cracker Example.

3.5 Probabilities in the Process Module

A scenario does not need to be complete in the sense that all or one of its elements must occur. For two *DEPENDENT* scenario elements, the conditional probability of an event *B* given that *A* has already occurred is commonly denoted as:

$$\Pr(A \text{ and } B) = \Pr(A) \cdot \Pr(B|A), \tag{3.10}$$

where *A, B* are two dependent events

Pr(.) is the probability.

In our setting, where scenarios are sequences of events, the fact that some scenario element *A* is a precondition for a successive element *B* implies that:

$$\Pr(B|\overline{A}) = 0, \tag{3.11}$$

where $\overline{A}$ denotes the event *not A*.

Moreover, for $\Pr(B|A)$ two special cases are distinguished. Firstly, if $\Pr(B|A) = 1$ then the events *A* and *B* can be considered *PERFECTLY CORRELATED*[28] since:

$$\Pr(A \text{ and } B) = \Pr(A). \tag{3.12}$$

Secondly, considering **(3.11)**, if $\Pr(B|A) = 0$ then:

$$\Pr(A \text{ and } B) = 0, \tag{3.13}$$

and the events *A* and *B* can be considered *EXCLUSIVE*[29].

[28] e.g., in case the security mechanism determining the event probability displays a perfect vulnerability, see **Assumption 3.3.**

[29] e.g., in case a security mechanism displays a perfect control, see **Assumption 3.3.**

3.6 Overview of the Function Module

In engineering sciences, Kröger and Mock [66] refer to a threat as *hazard directed towards an asset* while a security mechanism protects an asset. According to Salvati and Diergardt [45], a threat is refined into a threat agent and a threat action while a security mechanism is refined into vulnerability and control. A threat agent such as an employee, a hacker, malicious code, a group of people, a governmental agency triggers a threat action. A threat action exploits a security mechanism, for example, by probing IP packets, analyzing IP packets, creating password combinations. Eventually, the asset is abused.

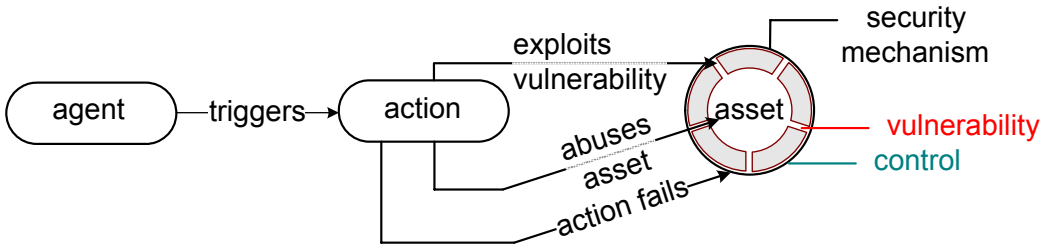

Figure 3.8: Threat Agent, Action, Security Mechanism and Asset.

As introduced in Chapter 2.2, in colloquial terms, a threat action possesses strength which a security mechanism opposes with resistance. Moreover, vulnerability and control are a measure to denote the surplus or insufficiency of a security mechanism in resisting a threat. In the case where the threat strength is greater than the resistance, the security mechanism displays vulnerability. The threat is then in a position to exploit or abuse it. If the threat strength is lower than the resistance, the security mechanism constitutes a control. The threat can neither exploit the security mechanism nor abuse the asset protected by it. In a colloquial interpretation, vulnerability (control) increases (decreases) the probability of a threat succeeding. **Figure 3.8** summarizes the above.

3.7 Probabilistic Concept of the Function Module

Assumption 3.2 "Stochastic Threats and Security Mechanisms": Unless one is a clairvoyant, they are unlikely to know when a threat attacks a security mechanism. Therefore, a *STOCHASTIC BEHAVIOUR OF THREATS AND SECURITY MECHANISMS* is assumed.

Whether or not a threat is successful depends, in many cases, on the timely update of the security mechanism counteracting that threat. Accordingly, threats can only exploit a

security mechanism within a specific time window. During this time window the security mechanism is vulnerable to the threat, otherwise the security mechanism is in control. Ideally, by displaying all timely occurrences of a threat and all timely occurrences of vulnerability in a specific security mechanism, the probability of the threat being successful is determined.

An example is the threat "new virus[30]". Any instance of this threat successfully exploits a security mechanism "anti-virus" only within a predetermined time window, namely when the security mechanism "anti-virus" has not yet been updated to recognize the "new virus". Likewise, hacking attempts, which are based on software vulnerabilities[31] are only successful while the software in question has not yet been updated with a security patch. Analogously, password cracking[32] attempts are only successful if they endure long enough to find the right password within an encrypted password file. Contrarily, the "better" the password quality the longer a password cracking attempt must last in order to crack it.

Assumption 3.3 "Perfect Vulnerability and Control": A security mechanism, which is never updated on time (i.e. a security mechanism with an infinitely small resistance) displays *PERFECT VULNERABILITY* towards the threat while a security mechanism, which is always updated on time (i.e. a security mechanism with an infinitely large resistance) constitutes *PERFECT CONTROL*.

Recall the statement in Chapter 1 requiring the distinction between the probability of a given threat exploiting a security mechanism and its (relative) frequency denoting how many times within an arbitrary time frame a threat is present and attempts to overcome the security mechanism.

In the case of perfect vulnerability the term "probability" looses its force of expression because any threat is successful when it attempts to overcome a security mechanism. Consequently, the only meaningful measure a decision maker can fall back on is the notion of frequency. Analogously, in case of a perfect control probability also looses its force of expression and the decision maker must fall back once more on frequency.

Assumption 3.4 "General Protection Strategy": To determine a general protection strategy a decision maker needs to know the frequency of occurrence of a threat:

---

[30] For further information on viruses refer to www.virus.org.

[31] For general information on attacks, which exploit software vulnerabilities, see www.exploitdatabase.com.

[32] For password cracking tools refer to http://sectools.org/crackers.html.

Frequency	General protection strategy	Examples
High / medium	Execution of periodic security processes	• Vulnerability patching • Software updates • Security awareness training
Low / very low	Business continuity measures	• Business continuity plans • Back up sites

Table 3.3: Frequency and General Protection Strategy.

3.8 Example: Brute Force Attacks on an Encrypted Password File

Reconsider the example of Jim Cracker in Chapter 3.4. Jim was fired due to irregularities in his conduct. Before leaving the company, he stole a password file on which he intends to perform a *BRUTE FORCE ATTACK* as previous *DICTIONARY ATTACKS* have failed. To estimate the probability of any attack succeeding all possible brute force attacks as well as all possibilities to encrypt passwords are taken into account.

3.8.1 Brute Force Attacks

Assumption 3.5 "Brute Force Attack": Let Ω_t represent the set of all brute force attacks and the random variable $T_t : \Omega_t \mapsto \Re$ its duration. The element $\omega_t \in \Omega_t$ is one specific brute force attack. The longer the duration of *independent*[33] brute force attacks on a time axis *t*, the fewer their number of occurrence[34]. This is modeled by the exponential probability density function *e(t)* **(3.14)**.

In a discrete interpretation, the values of *e(t)* represent the number of attacks (**Figure 3.9**). In a continuous interpretation (**Figure 3.10**), *e(t)* is normalized yielding a curve with area *1* **(3.15)** and an expected duration $E[T_t]$ **(3.16)**.

[33] In this context "independent" means that the individual brute force attacks do not exchange information among each other.

[34] This assumption is based on talks Salvati [103] has conducted with security professionals.

$$e(t) = \begin{cases} \lambda \cdot e^{-\lambda \cdot t} & , \quad t \geq 0 \\ 0 & , \quad else \end{cases}, \tag{3.14}$$

$$\int_0^{+\infty} e(t) \cdot dt = 1, \tag{3.15}$$

$$E[T_t] = \int_{-\infty}^{+\infty} t \cdot e(t) \cdot dt = \frac{1}{\lambda}. \tag{3.16}$$

where t describes the duration time of the attack, e.g., in days

$e(t)$ is the density function

$E[T_t]$ is the expected duration resulting from all brute force attacks, e.g., in days.

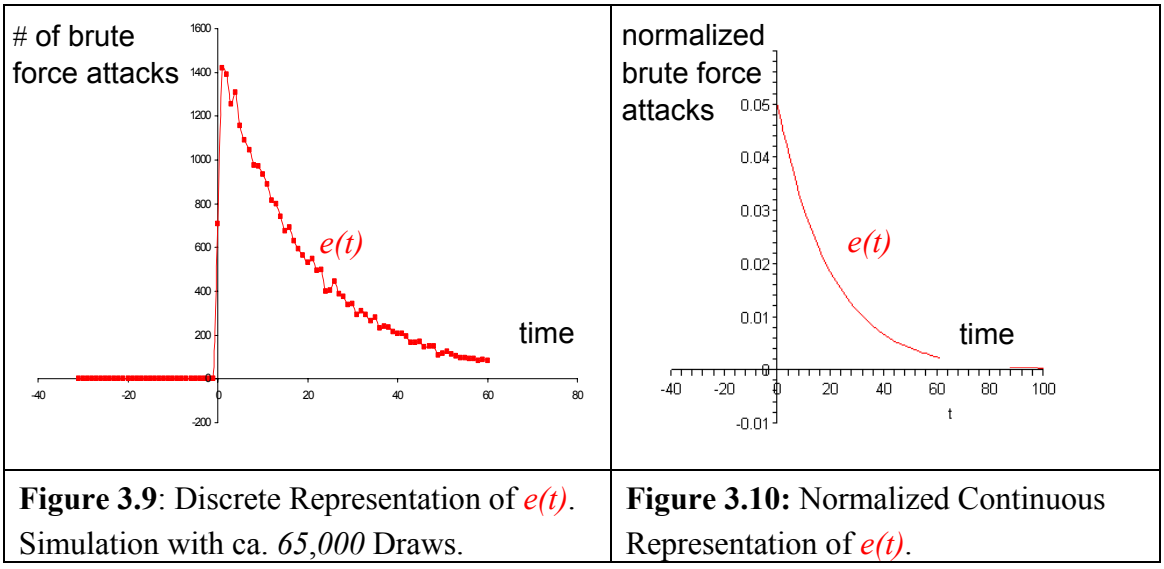

Figure 3.9: Discrete Representation of *e(t)*. Simulation with ca. *65,000* Draws.

Figure 3.10: Normalized Continuous Representation of *e(t)*.

Figure 3.9 shows a discrete representation of *e(t)* with $\lambda = 1/20$ and $E[T_t] = 20$. The graph shows a sample of *65,000* exponentially distributed and unconnected brute force attacks (simulation). The y-axis shows the number of occurrences which have a duration of *t*. **Figure 3.10** shows a normalized continuous representation of e(t) with $\lambda = 1/20$ and $E[T_t] = 20$.

Passwords are encrypted by using "this" password together with a salt[35] to generate a hash value. The hash value is then stored in the password file. According to Kaufman,

[35] A salt ideally comprises a sequence of random bits.

Perlman and Speciner [104] the salt and algorithm are known in practice (or easy to guess) and thus do not represent a challenge for the threat agent. A brute force attack attempts to recalculate and match the hash values of the passwords file. If it succeeds, the corresponding passwords in the file are cracked. One password cracking attempt comprises the calculation of the hash value, its comparison with all entries in the password file and lasts a specific amount of time τ. From τ, f is derived, which indicates, e.g., the processor frequency with which the password cracking attempts are performed. See equation **(3.17)**.

$$f = \frac{1}{\tau}, \tag{3.17}$$

where f describes the clock frequency of the processor, [MHz]

τ describes the time required for a complete cycle of a password cracking attempt

3.8.2 Encrypted Password File

The brute force attack computes a first hash value with the first *1*-character password of the password space and compares the result with the hash values stored in the password file. After exhausting all *1*-character passwords it tries all *2*-character passwords, then all *3*-character passwords, etc., until a matching hash value is found.

Assumption 3.6 "One Password": The security mechanism "encrypted passwords file" contains exactly one password of length n namely, the password the attacker is looking for.

The worst case for the security mechanism is when the valid password is at the lowest end of the n-character password space **(3.18)**, which causes the brute force attack to employ a minimum time t_{min} **(3.19)** to crack the password. The best case for the security mechanism is when the brute force attack needs to check the entire *1*-character, *2*-character, …, and n-character password space before finding a valid password **(3.20)**, which causes it to employ the time t_{max} **(3.21)**. t_{min} and t_{max} are determined by the number of attempts the brute force attack is capable to perform in a specific time frame, i.e. by the clock frequency of the processor.

$$n_{\min} = \sum_{i=1}^{n-1} z^i, \tag{3.18}$$

$$t_{\min} = \tau \cdot n_{\min}, \tag{3.19}$$

$$n_{\max} = \sum_{i=1}^{n} z^i, \tag{3.20}$$

$$t_{\max} = \tau \cdot n_{\max}, \tag{3.21}$$

where n is the length of the password, $0 \le i \le n$

z is the number of characters available for each position within the password (e.g., $z = 95$ for the printable ASCII set)

n_{min} is the minimum number of password guessing attempts before the password is found (worst case for the security mechanism)

n_{max} is the maximum number of password guessing attempts (best case for the security mechanism)

t_{min} is the minimum time the security mechanism "encrypted password file" resists the brute force attack, e.g., in days

t_{max} is the maximum time the security mechanism "encrypted password file" resists the brute force attack, e.g., in days.

Assumption 3.7 "Encrypted Passwords File": Let Ω_r represent the set of possible passwords and the random variable $T_r : \Omega_r \mapsto \Re$ their resistance to brute force attacks. The element $\omega_r \in \Omega_r$ represents one specific password and it is unknown to the attacker where the password is hidden in the password space. This is modeled by a uniformly distributed density function *u(t)* **(3.22)** whose area under the curve is *1* **(3.23)**.

In a discrete interpretation, *u(t)* shows the number of possibilities the encrypted passwords file accommodates to hide a single password (see **Figure 3.11**). The number of possibilities is equivalent to the number of attempts resisted over the time *t*. Accordingly, the time the passwords file resists starts at t_{min} and ends at t_{max}. The expected time the security mechanism resists is reported in **(3.24)**.

$$u(t) = \begin{cases} \dfrac{1}{t_{\max} - t_{\min}} & , \quad t_{\min} \le t \le t_{\max} \\ 0 & , \quad else \end{cases}, \tag{3.22}$$

$$\int_{-\infty}^{+\infty} u(t) \cdot dt = 1, \tag{3.23}$$

$$E[T_r] = \int_{-\infty}^{+\infty} t \cdot u(t) \cdot dt = \frac{1}{2} \cdot (T_{\max} - T_{\min}), \tag{3.24}$$

where t is the time the security mechanism "encrypted passwords file" is resisting the brute force attack, [s]

$u(t)$ is the uniformly distributed density function, [probability / time]

$E[T_r]$ is the expected time the security mechanism resists, [s].

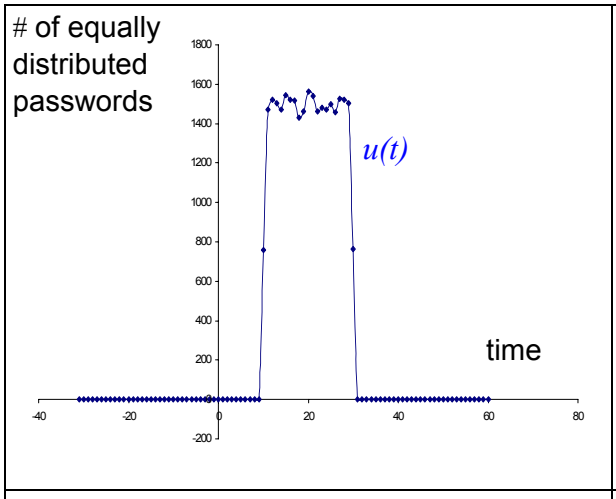

Figure 3.11: Discrete Representation of *u(t)*. Simulation with ca. *32,000* Draws.

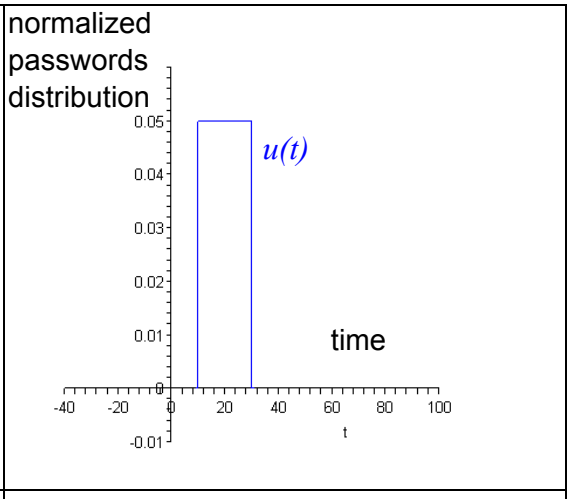

Figure 3.12: Continuous Representation of *u(t)*.

Figure 3.11 shows the discrete representation of *u(t)* with $t_{min} = 10$, $t_{max} = 30$, and $E[T_r] = 20$. The graph shows a sample of *32,000* uniformly distributed possibilities to hide passwords. The y-axis shows the number of occurrences of attacks which have a duration *t*. The uniformly distributed density function *u(t)* in **Figure 3.12** shows the time *t* (on the x-axis) the security mechanism "encrypted passwords file" resists to the brute force attack with $t_{min} = 10$, $t_{max} = 30$ and $E[T_r] = 20$.

3.9 Success Probability of Threats Overcoming Security Mechanisms

To calculate the cumulative probability of brute force attacks succeeding in overcoming an encrypted passwords file, all possible occurrences of brute force attacks are combined with all possibilities to hide a password in the file. A random brute force attack

is successful in cracking a password when it duration lasts longer than the resistance time of the password. Then, the cumulative probability is composed of:

$$\Pr(Z \le x) = \begin{cases} 0 & t_{\max} + x < 0 \\ \dfrac{\frac{1}{\lambda} \cdot \left(1 - e^{-\lambda(t_{\max}+x)}\right) - x - t_{\max}}{t_{\min} - t_{\max}} & -t_{\max} \le x \le -t_{\min} \\ \dfrac{t_{\min} - t_{\max} + \frac{1}{\lambda} \cdot \left(-e^{-\lambda(t_{\max}+x)} + e^{-\lambda(t_{\min}+x)}\right)}{t_{\min} - t_{\max}} & t_{\min} + x > 0 \end{cases} \quad \textbf{(3.58)}$$

where Z is the probability of the event "brute force attacks on encrypted passwords file" being successful

x is the time, e.g., in days

t_{min} is the minimum time the security mechanism "encrypted password file" resists the brute force attack, e.g., in days

t_{max} is the maximum time the security mechanism "encrypted password file" resists the brute force attack, e.g., in days

λ is the number of characters available for each position within the password.

In round terms, **(3.59)** is obtained by conversely convoluting (**3.14)** with **(3.23)**. The following briefly shows how this was done. (Comprehensive calculations are in **Appendix B**).

Suppose that T_t and T_r are two discrete random variables, which represent the stochastic occurrence of threats and security mechanisms respectively. The duration $T_t(\omega_t)$ of a random brute force attack ω_t out of the considered universe Ω_t is compared with the resistance $T_r(\omega_r)$ of a random password ω_r of the considered universe Ω_r. An arbitrary brute force attack decrypts the password if and only if:

$$T_t(\omega_t) \ge T_r(\omega_r), \quad \textbf{(3.25)}$$

i.e. its duration is greater or equal to the time the password resists. Conversely, if:

$$T_t(\omega_t) < T_r(\omega_r), \tag{3.26}$$

then the password prevails over the brute force attack. This is illustrated for *e(t)* and *u(t)* in **Figure 3.13**.

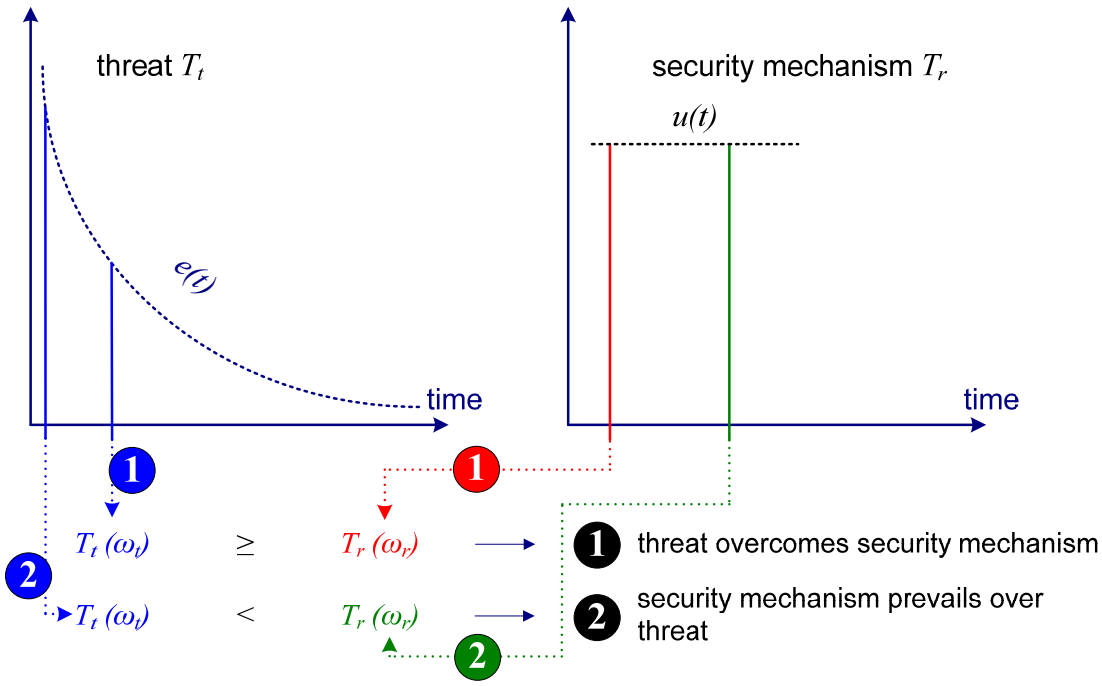

Figure 3.13: Stochastic Attack and Defense Behavior.

The combination of the random variables T_t and T_r represents the event *Z brute force attack on encrypted passwords file* with $Z:\Omega_t \times \Omega_r \mapsto \Re$ and a successful brute force attack is obtained by postulating $Z(\omega_t,\omega_r) = T_t(\omega_t) - T_r(\omega_r) \geq 0$, $\omega = (\omega_t,\omega_r)$.

The distribution function *e(t)* pertains to the threat with random Variable T_t and *u(t)* pertains to the security mechanism with random variable T_r. It is of interest to determine the distribution function *z(t)* which represents the random variable *Z*. Since T_t and T_r are independent it suffices to determine $\Pr(Z(\omega_t,\omega_r))$ which is equal to $\Pr(T_t(\omega_t))$ multiplied by $\Pr(T_r(\omega_r))$. All individual probabilities for T_t and T_r are then summed to obtain **(3.27)** which yields **(3.28)** as $Z = T_t - T_r$.

$$\Pr(Z(\omega)=t) = \sum_{(\omega_t,\omega_r)\in\Omega_t\times\Omega_r \mid t=T_t(\omega_t)-T_r(\omega_r)} \Pr(T_t(\omega_t)=t_t)\cdot\Pr(T_r(\omega_r)=t_r), \tag{3.27}$$

$$\Pr(Z(\omega)=t) = \sum_{t_t=-\infty}^{\infty} \Pr(T_t(\omega_t)=t_t)\cdot\Pr(T_r(\omega_r)=t_t-t), \tag{3.28}$$

$$z(t) = \sum_{t_t=-\infty}^{\infty} e(t_t) \cdot u(t_t - t)\,, \tag{3.29}$$

where t_t signifies the duration of the brute force attack (control variable with …, -2, -1, 0, 1, 2, …)

t_r signifies the resistance time of "this" password

t is a fixed point in time.

Equation (**3.29**) resembles the well-known convolution[36] of discrete random variables. The distinction lies in the argument $(t_t - t)$ of **(3.30)** which would show as $(t - t_t)$ in the convolution. Accordingly, for the continuous case, the probability density *z(t)* is calculated as:

$$z(t) = \int_{-\infty}^{+\infty} e(t_t) \cdot u(t_t - t)dt_t\,, \tag{3.30}$$

where t_t signifies the duration time of "this" brute force attack

t_r signifies the resistance time of "this" password

z(t) is the density function obtained by conversely convoluting *e(t)* and *u(t)*.

Integrating *z(t)* over the range of interest where $T_t \geq T_r$, i.e. $0 \ldots \infty$, (**3.31**) is obtained:

$$\Pr(T_t \geq T_r) = \Pr(Z \geq 0) = \int_0^{\infty} z(t)dt\,, \tag{3.31}$$

where t signifies one specific occurrence of the event $(Z = t)$, i.e. the pair T_t and T_r

T_t is the random variable representing the duration of brute force attacks

[36] The interested reader is referred to Crutchfield [105] for an animated demonstration of the (commutative and associative) convolution.

T_r is the random variable representing the encrypted passwords file

Z is the random variable of the event "brute for attacks on encrypted passwords file".

Conversely, in the case where the probability that the security mechanism prevails is of interest, then *z(t)* for $T_t < T_r$ is integrated over the range $-\infty \ldots 0$ and **(3.32)** is obtained:

$$\Pr(T_t < T_r) = \Pr(Z < 0) = \int_{-\infty}^{0} z(t)dt \,, \tag{3.32}$$

where t, T_t, T_r, same as above
Z.

Consequently, the probability density function *z(t)* is calculated by conversely convoluting **(3.14)** with **(3.22)**, which yields **(3.33)** and the cumulative distributed function *Z(t)* is derived by integrating *z(t)* over the area of interest.

$$z(t) = \int_{-\infty}^{+\infty} \left[\lambda e^{-\lambda t_t} \cdot \mathrm{I}_{[t_t \geq 0]}(t_t) \right] \cdot \left[\frac{1}{t_{\max} - t_{\min}} \cdot \mathrm{I}_{[-T_{\max} \leq t - t_t \leq -T_{\min}]}(t_t - t\ ) \right] \cdot dt_t \,, \tag{3.33}$$

where *z(t)* is the density function of the random variable T_t - T_r.

$\mathrm{I}_{[a \leq x \leq b]}(x)$ is the indicator function with the control variable x and value *1* in an interval $a \leq x \leq b$ and value *0* else, i.e.

$$\mathrm{I}_{[a \leq x \leq b]}(x) = \begin{cases} 1 & a \leq x \leq b \\ 0 & else \end{cases} \tag{3.34}$$

Applying case differentiation[37] for $t_{\max} + t < 0$, $-t_{\max} \leq t \leq -t_{\min}$, $t_{\min} + t > 0$ yields

[37] The reader is referred to **Appendix B** for all intermediate steps of the calculation.

$$z(t) = \begin{cases} 0 & t_{max} + t < 0 \\ \frac{\lambda}{t_{max} - t_{min}} \cdot \int_{-\infty}^{+\infty} e^{-\lambda t_t} \cdot \mathrm{I}_{[0 \le t_t \le t_{max}+t]}(t_t) \cdot dt_t & -t_{max} \le t \le -t_{min} \cdot \\ \frac{\lambda}{t_{max} - t_{min}} \cdot \int_{-\infty}^{+\infty} e^{-\lambda t_t} \cdot \mathrm{I}_{[t_{min}+t \le t_t \le t_{max}+t]}(t_t) \cdot dt_t & t_{min} + t > 0 \end{cases} \quad \textbf{(3.37)}$$

Integration and evaluation of the boundaries yields:

$$z(t) = \begin{cases} 0 & t_{max} + t < 0 \\ \frac{1}{t_{min} - t_{max}} \cdot \left[e^{-\lambda(t_{max}+t)} - 1 \right] & -t_{max} \le t \le -t_{min} \cdot \\ \frac{1}{t_{min} - t_{max}} \cdot \left[e^{-\lambda(t_{max}+t)} - e^{-\lambda(t_{min}+t)} \right] & t_{min} + t > 0 \end{cases} \quad \textbf{(3.41)}$$

The function *z(t)* is obtained by combining *e(t)* of **Figure 3.9** with *u(t)* of **Figure 3.11** (for the discrete case, simulation is applied) and **Figure 3.10** with **Figure 3.12** (for the continuous case) as indicated in equation **(3.30)** resulting in **Figure 3.14** and **Figure 3.15** respectively. Alternatively, *z(t)* is also found by numerical or by graphical methods (not followed up).

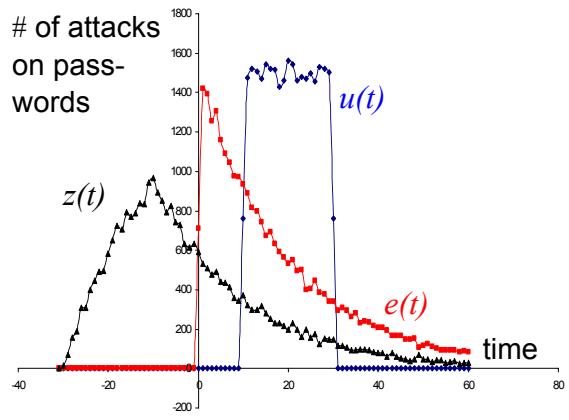

Figure 3.14: Discrete Representation of *z(t)*, *e(t)* and *u(t)*.

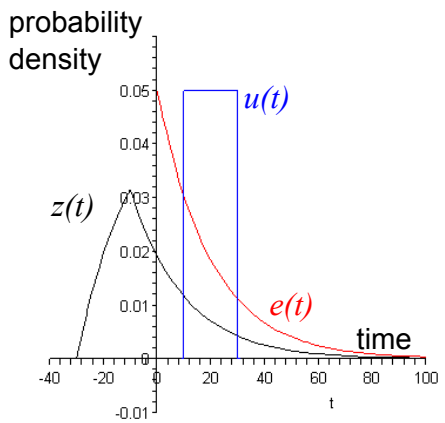

Figure 3.15: Continuous Representation of *z(t)*, *e(t)* and *u(t)*.

The probability of successful brute force attacks is the area described below *z(t)* in the range where $T_t \ge T_r$ ($t_{diff} \ge 0$), i.e. $0 \ldots \infty$. This value is given by $1 - F_Z(0)$ where F_Z is the

cumulative distributed function. F_Z is obtained by integration from $-\infty$ to the point x of interest, i.e:

$$F_Z(x) = \Pr(Z \le x) = \int_{-\infty}^{x} z(t) \cdot dt \,. \tag{3.42}$$

For $t_{\max} + x \le 0$ of the above case differentiation **(3.42)** results in:

$$\Pr(Z \le x) = 0 \,. \tag{3.43}$$

Taking the case of differentiation into account for $-t_{\max} \le x \le -t_{\min}$ yields:

$$\Pr(Z \le x) = \frac{1}{t_{\min} - t_{\max}} \cdot \int_{-t_{\max}}^{x} \left[e^{-\lambda(t_{\max}+t)} - 1 \right] \cdot dt \,, \tag{3.44}$$

$$= \frac{\frac{1}{\lambda} \cdot \left(1 - e^{-\lambda(t_{\max}+x)}\right) - x - t_{\max}}{t_{\min} - t_{\max}} \,. \tag{3.50}$$

Taking the case of differentiation into account for $t_{\min} + x > 0$:

$$\Pr(Z \le x) = \frac{1}{t_{\min} - t_{\max}} \cdot \left[\int_{-t_{\max}}^{-t_{\min}} \left[e^{-\lambda(t_{\max}+t)} - 1 \right] \cdot dt + \int_{-t_{\min}}^{x} \left[e^{-\lambda(t_{\max}+t)} - e^{-\lambda(t_{\min}+t)} \right] \cdot dt \right]. \tag{3.51}$$

$$\frac{t_{\min} - t_{\max} + \frac{1}{\lambda} \cdot \left(- e^{-\lambda(t_{\max}+x)} + e^{-\lambda(t_{\min}+x)} \right)}{t_{\min} - t_{\max}} \,. \tag{3.57}$$

In summary,

$$\Pr(Z \le x) = \begin{cases} 0 & t_{\max} + x < 0 \\ \dfrac{\frac{1}{\lambda} \cdot \left(1 - e^{-\lambda(t_{\max}+x)}\right) - x - t_{\max}}{t_{\min} - t_{\max}} & -t_{\max} \le x \le -t_{\min} \\ \dfrac{t_{\min} - t_{\max} + \frac{1}{\lambda} \cdot \left(- e^{-\lambda(t_{\max}+x)} + e^{-\lambda(t_{\min}+x)} \right)}{t_{\min} - t_{\max}} & t_{\min} + x < 0 \end{cases} . \tag{3.58}$$

3.10 Results

For the brute force attack, the following data is used: $E[T_t] = 20$, $\lambda = 1/20$. For the encrypted passwords file $E[T_r] = 20$, $f = 333\ MHz$, z = 95 characters (the printable ASCII set), $n = 7$, i.e. $t_{min} = 10$, $t_{max} = 30$ is calculated. With the above data, **(3.59)** yields:

for $-t_{max} \leq x \leq -t_{min}$: $Pr(Z \leq x) = 0.5 + exp(-1.5 - 0.05x) + 0.05x$

for $t_{min} + x > 0$: $Pr(Z \leq x) = 1 + exp(-1.5 - 0.05x) - exp(-0.5 - 0.05x)$

The green portion of **Figure 3.16** shows the curve for $-t_{max} \leq x \leq -t_{min}$ and the red portion of the curve indicates $t_{min} + x > 0$. In more detail, the following curve is drawn:

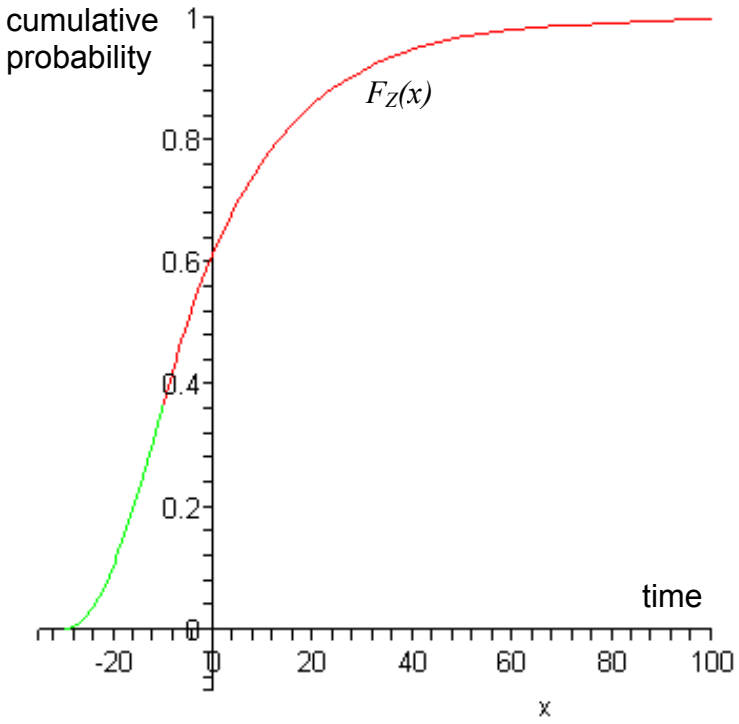

Figure 3.16: Probability Plot (Brute Force Attacks vs. Encrypted Passwords File).

As it is not possible to know which brute force attack will hit the passwords file and how a password in the file has been encrypted, all possible occurrences of brute force attacks are compared to every potential password. Consequently, a random brute force attack will crack a password if its strength is greater than the resistance of the password.

On the *x-axis*, **Figure 3.16** shows the time difference between the duration of a brute force attack and the resistance of the password where:

- $x = 0$ means that a brute force attack lasts as long as a password resists
- $x < 0$ means that a brute force attack lasts for less time than a password resists
- $x > 0$ means that a brute force attack lasts longer than a password resists.

Interpreting the above curve for application in practice yields a way of reasoning where decision makers either:

- accept the fact that around 40% of all brute force attacks present in the Internet may crack the passwords of the company should they be directed against it, or
- reject the fact that the passwords in the company are only successful in around 60% in resisting all brute force attacks.

4. Influence Module

This module solves the general Influence Problem by answering the question to what extent security processes influence success probabilities of attacks. It is applied to solve the more specific Governance Problem of global companies, which seeks to strike a balance between a centralized and a *laissez-faire* approach to governing their branches. The Influence and Governance Problems are introduced in Chapter 4.1. In Chapter 4.2 the strengths and weaknesses of various approaches for addressing the Influence and, in particular, the Governance Problem are analyzed (State-of-the-Art). This analysis yields requirements in order to solve both problems which are addressed by means of Rough Sets Theory (RST). In particular, set approximation of RST is used to display security information (Chapter 4.3), to the classification of branches (4.4), to evidence dependencies among security processes (4.5), to determine their dispensability (4.5), and to evaluate their significance (4.6). The interested reader is referred to **Appendices C** and **D** for an introduction to Rough Sets Theory.

4.1 Influence and Governance Problems

The *CONTEXT* entails security processes, which set up, maintain, or decommission security mechanisms.

Assumption 4.1 "Influence of the IS Context": Omitting the IS context leads to an incomplete decision set as vulnerabilities displayed by security mechanisms largely depend on security processes. This entails that security processes influence success probabilities, i.e. the interaction between threat and security mechanism.

Assumption 4.2 "Influence of the Business Context": Analogously, security processes in the business context influence probabilities.

INFLUENCE includes the:

- dependency of security mechanisms on security processes
- dispensability of security processes
- significance of individual security processes.

A prominent example of an Influence Problem is Public Key Infrastructure (PKI), which is used for encryption and authentication over the Internet. PKI indeed rests upon (technical) mechanisms such as encryption algorithms, key lengths, etc. However, failure to, e.g., diligently execute the correct identification of entities applying for PKI certificates (security process) will inevitably jeopardize the entire PKI (regardless of the "strength" of the security mechanisms).

The *GOVERNANCE PROBLEM* represents one specific instance of the Influence Problem. Suppose that a particular head office of a global company aims to influence the success probability of specific Internet attacks, which are directed against its branches. In terms of the Governance Problem, the head office is confronted with *how to design security processes for execution at its branches such that they favourably interact with local security mechanisms and their corresponding Internet attacks*.

For head offices, the Influence Module puts the probability density functions (*pdf*) of threats and security mechanisms in the context of the above security processes which allows balancing between rigorously exercised central governance and a *laissez-faire* approach. Evidencing influence boils down to identifying which processes interact with which security mechanisms and/or threats[38]. Practical experience, however, shows that the Influence and Governance Problems are challenged by the quality and quantity of the data describing security processes. This data is:

- vague
- incomplete and inconsistent
- unavailable (in the sense of unavailability of *a priori* data).

To perform a data analysis in terms of the Governance and Influence Problems, there is a variety of potentially applicable methods to choose from. Next, in "Standard Methods for Correlation Analysis", the Design of Experiments, Principal Component Analysis, Clustering, Bayesian Belief Networks and Variance Analysis are analyzed. Their elicitation has contributed in gaining a better understanding of the problems at hand.

[38] In some cases this relationship may also exist between security processes and threats. For an example, in Chapter 7 the process "security awareness training" is employed to influence the threat "employees responding to phishing emails".

4.2 Standards Methods for Correlation Analysis

Design of Experiments (DOE): DOE is an approach to plan and evaluate experiments where input variables (factors) are modified to identify the most significant ones and to relate them to the output variables. DOE approaches are superior[39] to One-Factor-At-a-Time (OFAT) experiments as they help to identify the main factors affecting the output variable(s) in a structured way. Each factor assumes a predefined number of values and is then combined with all other factors. Finally, the output variable is measured.

Let the *3* factors of an experiment each assume *2* values. Consequently, the maximum number of experiments is *2^3 = 8*. There are *3* main factors, *3* two-factor and *1* three-factor potentially affecting the output variable. By executing the maximum number of experiments, all the effects are visible. By reducing the number of experiments, the individual effects of the main, two and one-factor are superposed and interactions cannot be analyzed thoroughly any more. However, practice shows that (1) the main factors and their interactions are dominant, (2) as a rule of thumb, interactions of higher order (i.e. > *2*) can be neglected, (3) reduced tables still yield good results. Basically, DOE provides the means to meaningfully reduce the number of experiments. Then, by applying linear regression, statistical statements on the reliability of the results are made. Finally, linear equations describe the results of the experiment.

In principle, DOE prescribes a set of measurements for evidencing interdependencies between security mechanisms and processes. However, prescribing a structured approach to measuring security processes requires repeatedly executing experiments in branches of a global company, which is not feasible in practice. For example, pairing the fictitious security process *1* with process *2* and pairing process *2* with process *3* to measure a potential change on a security mechanism is time consuming and not acceptable. Consequently, we are looking for an approach that allows for measurements to be executed in one go.

Principal Component Analysis (PCA): PCA is used to classify objects described by a large set of multidimensional data. One area of application of PCA is automated face

[39] According to Czitrom [106] DOE shows clear advantages over OFAT as it requires fewer resources, estimates the effects of each factor more precisely, gives the opportunity to estimate the effects among the factors systematically and improves the prediction in the response because the entire factor space is taken into consideration.

recognition. A face is described by many dimensions; however, not all dimensions are needed to classify them because some are relevant while others are not. For example, two faces may considerably differ in the sections of the eyes, nose and mouth while they are not discernible by the forehead or cheeks. Consequently, from a complete set of measurements, the dimensions eyes, nose and mouth suffice. The goal of the PCA is to reduce the number of dimensions in the data.

Imagine the available data as a scatter plot, which has *n*-principal components and a center of gravity. These *n*-principal components describe the axes of an *n*-dimensional coordinate system. The center of the coordinate system is first moved to the center of gravity of the scatter plot representing the data. Then, it is placed into the scatter plot such that the first axis shows the direction of the largest variance of the scatter plot. Next, the coordinate system is turned such that the next axis (uncorrelated with the first axis) shows the direction of the largest variance remaining. This step is repeated until all *n*-dimensions of the scatter plot fit the *n*-dimensions of the coordinate system. Each principal component makes a contribution to the overall variance. This contribution is an indicator for the importance of each individual component to classify the data. By applying the PCA approach, the loss in information when reducing the dimensions of the scatter is minimal in terms of the least square fit.

If it is assumed that faces correspond to security mechanisms and the eyes, nose, etc. correspond to security processes, then this approach is basically applicable to solve the Governance Problem. However, in practice, an exhaustive set of multidimensional measurements on security mechanisms and processes is hardly the problem and inference is usually based on incomplete data.

Moreover, it is not of interest to create new security processes (new dimensions) because this potentially entails loosing information on an already low amount of data. In practice, the number of security processes is relatively stable as they usually evolved over time or are required by regulations such as the Sarbanes/Oxley Act [23].

Clustering: Clustering is a core task in structuring knowledge, which partitions data into subsets (clusters), so that the elements in each subset share some common clustering criteria. As indicated by Estivill-Castro [107], a clustering criterion is the mathematical representation of an inductive principle, which is in part in the eye of the beholder. Choosing one inductive principle over the other manifests fundamental differences in structuring knowledge. There is a magnitude of induction principles which can be used for selecting clustering criteria. Two examples are presented: (1) the Maximum Likelihood principle, which indicates the probability that an element belongs to a specified class and (2) the Homogeneity of a Class in terms of variance of multivariate analysis. Both principles are realized by algorithms. While the first one indicates

choosing a clustering algorithm that maximizes the probability of fit of the data being generated, the second one picks (out of a set of *k* representatives) one that minimizes the total squared error. As an example for the second inductive principle, the *PAIRWISE AGGLOMERATIVE CLUSTERING* as described by Faber [108] can be used. Imagine a number of values in ascending order. The interest is in clustering the data such that the values are successively merged starting from the lowest and closest two points. The two points are replaced by one single point (which represents the non-weighted average of the two points). The next step repeats this algorithm (find the two closest points, calculate their average, merge the two points) until a stop criterion is reached (which may be defined arbitrarily by the beholder).

The nature of clustering is exploratory. Accordingly, if it is intended to structure data about security processes or mechanisms, then clustering is appropriate. For the Governance Problem, specifying criteria of similarity for clustering data corresponds to a preprocessing step to explore potential discretization of security processes or mechanisms. Dealing with the Governance Problem is a non-trivial process, which postulates to identify valid, novel, potentially useful and ultimately understandable patterns (structure) in data. However, this is but one step towards indicating how to account for a relationship between security processes and security mechanisms.

Bayesian Belief Networks[40] (BBN): BBNs are used to represent dependencies between information variables and hypothesis variables, which are related by conditional probabilities or conditional independence statements. BBNs model a system in terms of a directed acyclic graph where the nodes represent states of random variables and the edges represent a probabilistic dependency relationship. BBNs are mainly used for statistical inference where they merge subjective belief (e.g., an expert opinion) and quantitative probability statements. Upon "feeding" a BBN with all the available information, if there is evidence that an event specified in the BBN has occurred, then the probability of other events happening or having happened can be inferred by using probability calculus and Bayes Theorem.

BBNs model expert knowledge in a descriptive rather than an exploratory manner. They tend to overemphasize expert knowledge when precise quantitative statements on probability distributions of security mechanisms or processes states are not available. As the Influence Module supports the discovery of potential relationships between security processes and security mechanisms, the descriptive character is hindering for the Influen-

[40] For an introduction to Bayesian Belief Networks, refer to Wooldridge [109].

ce Problem. Moreover, BBNs tend to become unmanageable for a larger number of security processes (e.g., for 32 security processes each having multiple possible states). Finally, the need to draw a separate BBN for each security mechanism (to evidence the relationship between security processes and "that" security mechanism) indicates limitations of the approach for exploring the Governance Problem.

Variance analysis[41] (VA): VA analyses the effects of one (or multiple) independent variable(s) onto one (or multiple) dependent variable(s). While regression analysis searches for a relationship between metric[42] independent variables and metric dependent variables, in variance analysis, the independent variables may occur in nominal[43] scale and dependent variables assume measurements either at the interval or ratio[44] levels. Precondition for applying VA is an assumption on the relationships among the variables (a model).

Imagine experimental data, which has been sampled in comparable branches to explain the effects of some independent variables upon one dependent variable. To explain these effects a model is set up which considers the mean of the dependent variables for each location and the overall mean of the (same) dependent variables for all locations. The *overall mean* is assumed to reflect the "true" value where all internal and external influences compensate for each other. VA presupposes that accidental influences external to a model occur for all branches with the same value and form a difference to the mean, which is indeed measured but cannot be explained. The *mean of each individual branch* expresses the peculiarities of its independent variables. This difference may contain influences of variables external to the model, which show in the measurements but, as before, cannot be explained by it. Conversely, if events considered by the model influence the branches in a systematic way then it is possible to assess this influence by an analysis of the mean sum of the squares. By comparing the mean sum of the squares among the effects, which take place between and within the branches, the statistical hypothesis of an occurrence (H_0) can be confirmed or rejected with a probability.

---

[41] For an introduction to variance analysis refer to Backhaus in [110].

[42] Metric measurements occur with interval or ratio scales. Such measures are distinguishable and can be brought in order.

[43] Measurements in the nominal scale are distinguishable among each other; however, they cannot be brought in any order. Examples are measurements such as "male" – "female", "Yellow" – "Green" – "Red", etc. Measurements in nominal scale may be interpreted as describing states, e.g., the implementation quality of a security process.

[44] The interval scale is partitioned in intervals of equivalent size and lacks a natural zero point. In addition to the interval scale, the ratio scale has a natural zero point.

In principle, VA is suited for solving the Governance Problem if a conversion is found transforming the measurements related to the security mechanisms and processes into interval/ratio levels and nominal levels respectively. However, operating with nominal scales for both security mechanisms and processes would accommodate better the usually vague measurements.

In light of this brief analysis, the author turns to Rough Sets Theory (RST) by Pawlak [111]. RST groups branches that display congeneric security processes into equivalence classes. Likewise, branches that display similar security mechanisms[45] are also assembled into equivalence classes. To build equivalence classes the use of ratio or interval scales is not necessary as the branches may occur with nominal values or values that lie within a specific range (roughness). In essence, RST determines how well an equivalence class (formed by security processes) approximates a second equivalence class (formed by a security mechanism). If the branches of the first equivalence class are also predominantly present in the second equivalence class then a relationship is assumed between security processes and security mechanisms.

On one hand, RST offers advantages over the aforementioned methods because:

- it fits well the initial requirements on vagueness of data, the low amount of data, and the unavailability of *a priori* data
- in contrast to DOE, it requires a singular execution of the measurements rather than multiple measurements of the same security mechanism or threat
- in contrast to PCA, it selects existing security processes, which "best" describe the relationship to security mechanisms (rather than creating a lower number of new security processes)
- in contrast to Clustering, it identifies relationships between processes and mechanisms / threats rather than exploring a potential representation of the data
- in contrast to VA, it handles nominal values for security processes and mechanisms.

On the other hand, the shortcomings[46] of RST lie in:

- its exploratory nature where the choice of the nominal values and the range of values (roughness) for building the equivalence classes is arbitrary

[45] and/or threats

[46] Limitations to RST from a practical point of view are reflected in Chapter 8.1.3.

- the computational effort required to build subsets of *n* security processes for subsequent verification of their relationship to a security mechanism. The effort rises exponentially given that 2^n subsets of security processes exist (the problem is NP-hard according to Rauszer [112]).

Despite the above challenges, the aim is to show the applicability of RST for use in IS risk management to determine the influence of security processes on security mechanisms/threats. To this end, a fictitious example is presupposed where data is available on:

- the quality of five security processes, as well as,
- the probability density functions of a brute force attack and a an encryption mechanism (password quality).

4.3 Displaying Security Information in Data Tables

Suppose that the following security processes are required by a fictitious head office to be implemented at the branches of a global company:

1. adoption of security policies
2. creation of local security procedures
3. execution of technical security assessments
4. allocation of resources to security functions
5. regular follow ups on security in local projects.

Further assume that the aforementioned five security processes have been assessed for ten branches located in *Frankfurt, Madrid, Paris, Milan, Guernsey, Luxembourg, New York, Nassau, Singapore* and *Sydney*. In practice, the implementation quality of security processes is usually assessed by assigning discrete values (e.g., *D+, C, D-, N, N/A*)[47].

The head office is interested in understanding how the above security processes influence the probability of successful brute force attacks. Consequently, the *pdf*s of brute force attacks and encryption mechanisms are measured for the above branches. Further

[47] *D+* signifies that the security process exceeds the implementation quality required by the head office, *C* means that the security process is compliant with expectations, *D-* signifies below expectation and *N* signifies the absence of a security process. *N/A* indicates that the security process is not eligible for implementation at the branch.

suppose that those *pdf*s show the duration of brute force attacks to be exponentially falling and the resistance of passwords is uniformly distributed (as in Chapter 3).

The following problems need to be addressed before processing the raw data further. Input data must be:

- corrected as it may be affected by errors in measurement or display
- completed as it may be lacking attribute values
- transformed to become discrete. In particular, the probability density of brute force attacks and the resistance of passwords are descretized along *D+, C, D-, N* and *N/A*
- acknowledged with respect to its inconsistency, i.e. some branches have the same values for brute force attacks yet different values for passwords.

In the ongoing example, the above problems have been addressed by the specialist functions concerned with the assessment. As a result, **Figure 4.1** displays the values of the security processes (yellow/set A), the security mechanisms (red/set Dec) and the branches (grey/Set U):

set U	set A					set Dec	
	Adoption of Policies	Creation of Procedures	Technical Assessments	Resource Allocation	Security in Projects	Brute Force Attack	Password Quality
Frankfurt	D-	C	C	D-	D-	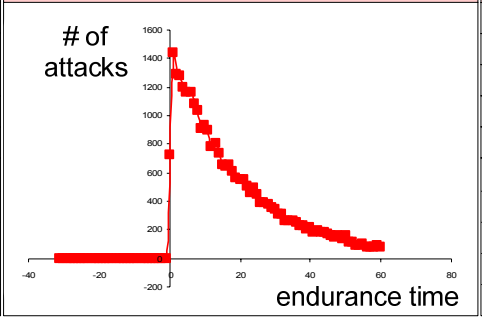	D-
Madrid	C	D-	N	N	N		N
Paris	C	D-	N	N	N		D-
Milan	D-	C	D-	D-	D-		D-
Guernsey	D-	D-	N	N	N		D-
Luxembourg	C	C	C	D-	D-		C
New York	C	C	N	D-	D-		C
Nassau	D-	D-	D-	N	N		D-
Singapore	C	D-	C	N	N		C
Sydney	D-	C	D-	D-	N		D-

Figure 4.1: Fictitious Security Information in a Data Table.

More generally, the security information contained in a data table *I* forms a 7-tuple:

$$I = \langle U, A, Dec, V_a, V_d, \rho_a, \rho_d \rangle, \quad \textbf{(4.1)}$$

where	U	is a set of branches (universe U of objects[48])
	A	is a set of security processes (conditional attributes, see footnote)
	Dec	is a set of threats and security mechanisms (decision attributes, see footnote)
	V_a	is a set of values C that security processes are described with; e.g., *D+, C, D-, N* and *N/A*
	V_d	is a set of values $\{v_{d.1}, v_{d.2}, \ldots, v_{d.q}\}$ that threats and security mechanisms are described with; e.g., *D+, C, D-, N* and *N/A*
	ρ_a	represents a function that assigns values to security processes; i.e. the assessment by the risk analyst
	ρ_d	represents a function that assigns values to threats and security mechanisms; i.e. the assessment by the risk analyst.

In **Figure 4.1**, the brute force attacks are shown symbolically with the same frequency curve for all branches while the frequency curve for the *password quality* of each branch has been measured and discretized in terms of *D+*, *C*, *D-*, *N* and *N/A*.

The security information in **Figure 4.1** may be unnecessarily large as indiscernible branches can be represented by equivalence classes. Furthermore, some security processes may not be needed to build equivalence classes as they depend on other security processes. They are rejected and called (totally or partially) redundant. These simplifications to the data table are discussed next.

[48] In terms of RST, branches are *OBJECTS*, security processes are *CONDITIONAL ATTRIBUTES* and security mechanisms are *DECISION ATTRIBUTES*. These terms are not followed up.

4.4 Classifying Branches by Set Approximation

Consider the data table in **Figure 4.1** where *B = {policies, projects}*[49] is a subset of *A*. Discerning branches of *U* by the security processes in *B* yields the equivalence classes:

B$_1$: *{Fr, Mi}*[50], where *policies* = *D-* and *projects* = *D-*

B$_2$: *{Ma, Pa, Si},* where *policies* = *C* and *projects* = *N*

B$_3$: *{Gu, Na, Sy},* where *policies* = *D-* and *projects* = *N*

B$_4$: *{Lu, Ne},* where *policies* = *C* and *projects* = *D-*.

Generalizing, this is equivalent to applying an indiscernability relation *IND(.)* to a subset *B* of *A* which yields a structure of equivalence classes *[x]*$_B$ where:

$$IND(B) = \{(o_i, o_j) \in U^2 \,|\, \rho_b(o_i) = \rho_b(o_j); \forall b \in B\}, \tag{4.2}$$

$$[o]_B = \{B_1, B_2, ..., B_f\}, \tag{4.3}$$

and *o* is a branch

B is a subset of security processes of *A*

B_i is the *i*-th equivalence class induced by *B* with $0 \leq i \leq f$

[o]$_B$ is the equivalence class structure induced by *B*.

The indiscernability relation is the basis of RST. It is applied to the Governance Problem to build sets $X \subseteq U$ of branches where security mechanisms or threats show common characteristics. To illustrate this, let the arbitrary target set *X* contain all branches with an insufficient password quality (*password quality* = *D-*), i.e.

X = {Fr, Pa, Mi, Gu, Na, Sy}.

[49] *Policies* means *adoption of policies*; *projects* means *security in projects*.

[50] *Fr* means *Frankfurt*, i.e. each location is abbreviated by its first two letters.

In the ongoing example, it is of interest to know how well X is approximated by the equivalence classes of $[o]_B$. If this succeeds, a relationship between the context (security processes of B) and the security mechanism *password quality* is inferred. Describing X by the B-indiscernible equivalence class structure yields an approximate set:

$$\{Fr, Mi\} \cup \{Ma, Pa, Si\} \cup \{Gu, Na, Sy\}.$$

This approximate set also contains the branches *Ma* and *Si*, which are not present in X. In fact, X cannot be expressed exactly in terms of $[o]_B$ because the security processes of B also discern other branches not contained in X. This is called the *UPPER APPROXIMATION* of X induced by B and is constructed by letting:

$$\overline{B}X = \left\{o \in U \middle| [o]_B \cap X \neq \{\ \}\right\}. \tag{4.4}$$

Conversely, the *LOWER APPROXIMATION* is denoted as:

$$\underline{B}X = \left\{o \in U \middle| [o]_B \subseteq X\right\}, \tag{4.5}$$

which in the ongoing example yields:

$$\{Fr, Mi\} \cup \{Gu, Na, Sy\}.$$

The lower approximation correctly classifies all branches of X at the expense of omitting those equivalence classes, which add branches not contained in the target set X (as is the case for B_2).

Next, the *ACCURACY OF APPROXIMATION* of X by B is defined as the ratio:

$$\mu_B(X) = \frac{\left|\underline{B}X\right|}{\left|\overline{B}X\right|} \leq 1. \tag{4.6}$$

In the ongoing example, the accuracy of approximation is:

$$\mu_B(X) = \frac{\left|\underline{B}X\right|}{\left|\overline{B}X\right|} = \frac{\left|Fr, Mi, Gu, Na, Sy\right|}{\left|Fr, Mi, Ma, Pa, Si, Gu, Na, Sy\right|} = \frac{5}{8} = 0.625.$$

From the upper and lower approximation the *BOUNDARY REGION* $BN_B(.)$ is determined. The boundary region contains all branches that neither can be ruled in nor ruled out as members of *X*. It is denoted by:

$$BN_B(X) = \overline{B}X - \underline{B}X \ . \tag{4.7}$$

In the ongoing example, the boundary region is:

$$BN_B(X) = \{Fr, Mi, Ma, Pa, Si, Gu, Na, Sy\} - \{Fr, Mi, Gu, Na, Sy\} = \{Ma, Pa, Si\}.$$

The negative region is represented by U - $\overline{B}X$, i.e.:

$$\text{negative region} = U - \overline{B}X \ . \tag{4.8}$$

In our example, the negative region is:

$$\{Lu, Ne\}.$$

In the ongoing example, in 5 out of 8 branches, the security processes *policies* and *projects* correctly describe (classify) the target set *X*, i.e. their influence on *password quality* of is moderate. In particular, for the branches located in *Madrid*, *Paris* and *Singapore* (the boundary region) it is unclear whether a relationship exists between the security processes and the password quality. This result indicates that other security processes may be better suited to influence the password quality at branches.

Consequently, the search for more appropriate descriptions of *X* is continued (where *X* consists of branches with *password quality = D*-). By calculating the lower and upper approximations of *X* induced by all possible combinations of the security processes an accuracy of approximation of *5 / 7 = 0.714* is found for:

{policies, technical, resources} or

{policies, procedures, technical} or

{policies, technical, resources, projects} or

{policies, procedures, technical, projects} or

{policies, procedures, technical, resources, projects}

(out of all possible combinations of security processes).

In terms of RST, a significant relationship has been identified between the above security processes and branches and insufficient password quality. Next, the methods for feature selection and extraction are refined by considering dependencies among security processes, their dispensability as well as their significance.

Remark: An accuracy of approximation of *1* would indicate that the security processes *policies* and *projects* correctly describe all branches where the password quality is insufficient. An accuracy of approximation of *0.5* would show a random relationship and an accuracy of approximation of *0* would show no relationship.

4.5 Dependency among Security Processes

In terms of RST, the degree k of dependency $(0 \leq k \leq 1)$ between two disjoint subsets of security processes, e.g., B and C is computed by:

$$k = \gamma(B,C) = \frac{\sum_{i=1}^{f} |\underline{C}B_i|}{|U|} \leq 1, \tag{4.9}$$

where $\gamma(B,C)$ denotes the degree k of dependency of B on C

B_i is the i-th equivalence class of $[x]_B = \{B_1, B_2, \ldots B_f\}$ induced by the B-Indiscernability relation with $0 \leq i \leq f$

$\underline{C}B_i$ is the lower approximation of B_i in terms of C.

The ratio in **(4.9)** is called *QUALITY OF APPROXIMATION*. The numerator shows all equivalence classes B_i with $0 \leq i \leq f$, which are described in terms of the lower approximation induced by C. Accordingly, if every equivalence class B_i is matched perfectly by its lower approximation $\underline{C}B_i$ then the security processes in B are totally dependent on the security processes in C.

4.5.1 Interpretation of Dependency

First, let B = *{procedures}* and C = *{resources}*. Accordingly, referring to **Figure 4.1**,

$[x]_B = \{\{Fr, Mi, Lu, Ne, Sy\}, \{Ma, Pa, Gu, Na, Si\}\}$,

$[x]_C = \{\{Fr, Mi, Lu, Ne, Sy\}, \{Ma, Pa, Gu, Na, Si\}\}$.

Applying **(4.9)** to $[x]_B$ and $[x]_C$ results in:

$$\gamma(B,C) = \frac{5+5}{10} = 1.$$

Both security processes create identical equivalence classes and B = *{procedures}* is totally dependent on C = *{resources}*. This relationship indicates that assessing the values of the security process in C potentially generates the values for the process in B over all branches. As this dependency is symmetric, it does not matter which of the two security processes is assessed. However, dependency between two disjoint sets of security processes is not induced by the values of the security processes but rather by their equivalence class structures. To exemplify, the values for the security processes *procedures* and *resources* are:

	branches	Fr	Ma	Pa	Mi	Gu	Lu	Ne	Na	Si	Sy
B:	*procedures*	*C*	*D-*	*D-*	C	*D-*	*C*	*C*	*D-*	*D-*	*C*
C:	*resources*	*D-*	*N*	*N*	D-	*N*	*D-*	*D-*	*N*	*N*	*N*

Table 4.1: Values for *Procedures* and *Resources* (all Branches).

The colour code in **Table 4.1** denotes branches of the same equivalence class. The *light orange* and *light green* show the equivalence class structure for *resources* while the colours *dark orange* and *dark green* denote the equivalence class for *procedures*. Accordingly, *procedures* depend on *resources* (or vice versa) although their values differ for each branch. This total dependency is evidenced by a square box around "branches" (grey). In practice, this implies that an initial assessment is needed to obtain starting values for the two processes. In a next step, a relation mapping one value set into the other must be constructed.

Second, let B = *{procedures}* and D = *{policies, resources}*. Accordingly,

$$[x]_B = \{\{Fr, Mi, Lu, Ne, Sy\}, \{Ma, Pa, Gu, Na, Si\}\},$$
$$[x]_D = \{\{Fr, Mi\}, \{Ma, Pa, Si\}, \{Gu, Na, Sy\}, \{Lu, Ne\}\}.$$

Applying **(4.9)** results in:

$$\gamma(B,D)=\frac{4+3}{10}=0.7\,.$$

B = *{procedures}* is described to 70% by the lower approximation induced by *D* = *{adoption of policies, resource allocation}*. In total, 7 branches of $[x]_B$ share a relationship with branches of $[x]_D$ namely, *Frankfurt, Milan, Luxembourg, New York* and *Madrid, Paris, Singapore*. Accordingly, for these branches, the values of *procedures* follow from *policies* and *resources*, see T**able 4.2**.

	branches	Fr	Ma	Pa	Mi	Gu	Lu	Ne	Na	Si	Sy
B:	*procedures*	*C*	*D-*	*D-*	*C*	*D-*	*C*	*C*	*D-*	*D-*	*C*
D:	*policies*	*D-*	*C*	C	*D-*	*D-*	*C*	*C*	*D-*	C	*D-*
	resources	*D-*	*N*	*N*	*D-*	*N*	*D-*	*D-*	*N*	*N*	*N*

Table 4.2: Values for *Procedures*, *Policies* and *Resources* (all Branches).

However, as the equivalence class structure built by *D* is more granular than the structure built by *B*, the above dependency is not symmetric. Third, switching the succession order of $\gamma(B,D)$ into $\gamma(D,B)$ results in:

$$\gamma(D,B)=\frac{0+0+0+0}{10}=0\,.$$

Accordingly, no dependency can be shown between *D* and *B* as the equivalence class structure built by *D* cannot be described in terms of the lower approximation induced by *B* and thus they share no relationship.

	branches	Fr	Ma	Pa	Mi	Gu	Lu	Ne	Na	Si	Sy
B:	*policies*	*D-*	*C*	C	*D-*	*D-*	*C*	*C*	*D-*	C	*D-*
	resources	*D-*	*N*	*N*	*D-*	*N*	*D-*	*D-*	*N*	*N*	*N*
D:	*procedures*	*C*	*D-*	*D-*	*C*	*D-*	*C*	*C*	*D-*	*D-*	*C*

Table 4.3: Values for *Policies*, *Resources* and *Procedures* (all Branches).

4.5.2 Applying the Quality of Approximation to the Governance Problem

The concept of quality of approximation is applied to formulate a relationship between the five security processes in **Figure 4.1** and the security mechanism *password quality* (i.e. to the Governance Problem). The equivalence class structure of the password quality is (see **Figure 4.1**):

$$[x]_X = \{\{Fr, Pa, Mi, Gu, Na, Sy\}, \{Ma\}, \{Ne, Lu, Si\}\}.$$

There are various combinations of security processes which induce a good approximation of the password quality. However, C_1 and C_2 shown below perform a qualitative approximation with a minimal number of security processes:

$C_1 = \{policies, technical, resources\}$ yields

$[x]_{C1} = \{\{Fr\}, \{Ma, Pa\}, \{Mi, Sy\}, \{Gu\}, \{Lu\}, \{Ne\}, \{Na\}, \{Si\}\}$ and

$C_2 = \{policies, procedures, technical\}$ yields

$[x]_{C2} = \{\{Fr\}, \{Ma, Pa\}, \{Mi, Sy\}, \{Gu\}, \{Lu\}, \{Ne\}, \{Na\}, \{Si\}\}$.

Accordingly, applying **(4.9)** to the above yields:

$$\gamma(X, C_1) = \frac{5+0+3}{10} = 0.8,$$

and

$$\gamma(X, C_2) = \frac{5+0+3}{10} = 0.8.$$

4.6 Dispensability of Security Processes

As previously seen, the security processes *procedures* and *resources* describe the same equivalence class structure. Consequently, instead of using a set $E = \{procedures, resources\}$ to describe this equivalence class structure, a reduced set R consisting of *{pro-*

cedure} or *{resources}* is sufficient. This set is called a *REDUCT*. Generalizing, a reduct *R* is a subset of *E* such that:

$$[o]_R = [o]_E, \tag{4.10}$$

where *RED* is a minimal set.

By subtracting any further security process from *R*, the equivalence class structure induced by *E* is not preserved. In a data table of security processes, there may be more than one reduct for the same equivalence class structure as $[o]_E$. Moreover, security processes may exist, which are common to every reduct and are said to form the *CORE*, i.e.:

$$CORE(E) = \bigcap_{R \subseteq E, [o]_R = [o]_E} R \quad . \tag{4.11}$$

Such security processes are more important than others because removing them causes an irrevocable loss of information. They are called indispensable for representing the structure of the data table while other security processes are called dispensable. Searching for indispensable processes in *E* means to search for *CORE(E)* (which is equivalent to finding a minimal reduct with the maximum number of reducts being $2^{|E|}$). Searching for the *CORE* is *NP*-hard. Rauszer [112] and Skowron and Rauszer [113] showed that the computational power required for finding all reducts rises exponentially[51]. Therefore, many algorithms[52] have been proposed to find approximate results. In IS Risk, however, $|E|$ is usually low. For example, by using the *ROSE 2* Software [117] with *B = {policies, procedures, technical, resource, projects}* yields *K = {policies*, *technical}* as the core. It is possible for the core to be empty.

Referring to the Governance Problem, the core is interpreted as a minimal set of security processes, which inevitably needs attention by head offices. Consequently, for balancing security processes between rigorously exercised central governance and a *laissez-faire* approach, head offices may decide to keep the indispensable security processes while the dispensable ones can be taken off its radar. The quality of approximation for the core *K* is calculated by letting:

$[o]_X$ = *{{Fr, Pa, Mi, Gu, Na, Sy}, {Ma}, {Ne, Lu, Si}}* and

[51] As cited in Delic, Lenz and Neiling [114].

[52] An introduction on heuristics for calculating the *CORE* is given by Zhong, Dong and Oshuga [115]. See also Nguyen [116].

$[o]_K$ = *{{Fr}, {Ma, Pa, Ne}, {Mi, Na, Sy}, {Gu}, {Lu, Si}}.*

Applying **(4.9)** results in:

$$\gamma(X,K)=\frac{5+0+2}{10}=0.7\,.$$

A quality of approximation of $\gamma(X,C)=0.7$ is not significantly worse than $\gamma(B,C)=0.8$. Therefore, head office may as well choose to concentrate on the core *K* (two security processes) instead of assessing three security processes (as in C_1 or C_2) for each branch.

4.7 Significance of Security Processes

An alternative to the aforementioned feature selection based on dispensability considers significant security processes $a \in B$ which are evaluated by removing *a* from *B* in $\gamma(B, C)$. The difference shows the change in dependency when removing the security process *a*. Normalization by $\gamma(B, C)$ yields the *ERROR OF CLASSIFICATION* $\sigma_{(B, C)}(a)$, i.e.:

$$\sigma_{(B,C)}(a)=\frac{\gamma(B,C)-\gamma(B-\{a\},C)}{\gamma(B,C)}\le 1, \qquad \textbf{(4.12)}$$

where σ is the error of classification which occurs when the security process *a* is dropped.

For example, by letting *B = {policies, technical}, C = {policies, technical, resources}* and by applying **(4.12)**, the error of classification is obtained:

$$\sigma_{(B,C)}(resources)=\frac{0.8-0.7}{0.8}=0.125\,.$$

Analogously, a set *F* instead of *{a}* can be used where *F* is a subset of *B* and any subset *B* of *C* can be treated as an approximate reduct of *C* yielding the *ERROR OF REDUCT APPROXIMATION*, i.e.:

$$\varepsilon_{(B,C)}(F)=\frac{\gamma(B,C)-\gamma(F,C)}{\gamma(B,C)}\le 1, \qquad \textbf{(4.13)}$$

where ε is the error of reduct approximation.

The idea of an approximate reduct is useful in the cases when a smaller number of security processes is preferred over the quality of classification. This type of feature selection maps an approximate reduct to an arbitrarily defined threshold figure (the error of reduct approximation).

4.8 Results

The Governance Problem of a fictitious head office interested in knowing which of five security processes "best" describe the password quality at its ten branches has been devised. To achieve this, the security processes best describing a password quality of *D*- over all branches were assessed. As a result, the security procedures *policies, technical, resources* or *policies, procedures, technical* were found to describe this relationship with an *ACCURACY OF APPROXIMATION* of *0.714*.

Five security processes were assessed as to whether total dependencies can be spotted among them. A total dependency between security processes may encourage the head office to assess only part of the processes present at the branches because the values of one security process determines the values of the other. Such an approach saves time and money. Occasionally, to "recalibrate" the Influence Module, a check of all security processes could be envisaged to confirm the dependencies. For the ongoing example, the *procedures* were found to be totally dependent on *resources*. Also, *procedures* is described by the security processes *policies* and *resources* with a *QUALITY OF APPROXIMATION* of *0.7*. Conversely, no dependency between *policies, resources* and *procedures* has been found. Next, the quality of approximation for *policies*, *technical*, *resources* in describing the security mechanism *password quality* has been found to be *0.8*. The same figure also applies for the relationship between *policies*, *procedures*, *technical* and the security mechanism *password quality*.

Then, the dispensability of security processes was looked at which leads to the concepts of *REDUCT* and *CORE*. The core has been proposed as another criterion to strike a balance between centrally exercised security governance and a *laissez-faire* approach. In the ongoing example, this approach significantly reduced the number of security processes relieving the head office into assessing only two as opposed to five security processes (*policies* and *technical*).

Imagine adding an arbitrary number of security processes to the core to form an expanded set of security processes. Further assume that this set describes a relationship to a specific security mechanism. This expanded set of security processes may yield a better quality of approximation than the core alone. Adding security processes to the core can be repeated either until an arbitrary threshold of the quality of approximation is reached or until adding security processes to it does not increase the approximation any more. However, there may be multiple possibilities for adding security processes to the core, so that each of the expanded sets delivers the same approximation quality. Consequently, the question is which set to choose in order to describe the relationship between security processes and security mechanisms. If two expanded sets of security processes are faced, the notion of significance of a security process may yield a useful concept for choosing among the sets. This criterion for feature selection is now used to calculate the *ERROR OF APPROXIMATION* when omitting one or more security processes from the enhanced set. In fact, it could be decided to choose the enhanced set where omitting one security process from the core causes the other core security processes to preserve as much as possible of their approximation quality.

5. Decision Module

Chapter 5 introduces the Decision Module in which decision makers accept or reject risks based on personal preferences. Chapter 5.1 outlines the Decision Problem and in Chapter 5.2 the Value at Risk and the Analytical Hierarchy Process are presented as well as standard methods for approaching it; in Chapter 5.3 the decision situation of IS risk management is elaborated while in Chapter 5.4, the graphical notation of the Process Module is adapted for use in the Decision Module. In Chapter 5.5, the Five Axioms[53] of Utility Theory are transposed for use in IS risk management while in Chapter 5.6 the maximum price for a risk analysis is discussed. Chapter 5.7 shows an application example and those results are discussed in Chapter 5.8.

5.1 Decision Problem

Description "Decision Problem": A corporate decision maker intending to equip "the" company with the "appropriate" level of security faces the following two challenges:

- First, what is the "appropriate" level of security?
- Second, what are the "right" security mechanisms to obtain the desired level?

The aforementioned questions are answered by observing the entrepreneurship paradigm. Accordingly, decision criteria for the selection of security mechanisms are anchored into the company's business context rather than in its IS context. Such a paradigm offers advantages over compliance management as it yields a less checklist oriented and a less bureaucratic approach.

A variety of methods are available[54] to approach the above questions. This thesis highlights two for closer investigation: the *VALUE AT RISK* and the *ANALYTICAL HIERARCHY PROCESS*. Their elicitation has contributed in gaining a better understanding of the Decision Problem.

[53] For an introduction the reader is referred to **Appendix E.2**.

[54] For an extensive list see Chapter 2.3.4

5.2 Value at Risk and Analytical Hierarchy Process

Value at Risk[55] (VaR): In the financial industry, VaR expresses the market risk of portfolios of financial instruments in terms of a single measure showing the worst loss due to "usual" market movements. More precisely, VaR is a quantile measurement indicating the maximum loss of a portfolio after a specific holding period at a predefined confidence level. Precondition for the calculation of the VaR is knowledge about the basic market factors of the instrument under investigation. These market factors are used to feed pricing formulas, which determine the value of the instrument.

Assume the VaR of one FX forward contract[56] in $t_0 + t$ needs to be calculated where t_0 denotes the current day. The VaR of such a portfolio typically expresses its potential maximum loss due to changes in foreign exchange rates for one to ten days and the confidence level usually ranges from 95% to 99%. To calculate the VaR, three methods are available: historical simulation, variance-covariance, and Monte Carlo statistical techniques.

The *HISTORICAL SIMULATION* method is based on data of gains and losses that the existing contract has experienced each day in a past period, say for the past *100* days. These changes in asset value are ordered from the lowest to the highest gain and loss, grouped and displayed in a frequency curve. If multiple instruments are assessed then this step is repeated for all other instruments and the changes in value are all summed up. Then, the confidence interval is determined in the frequency curve and the VaR denotes the maximum loss value left of that interval (quantile). This method is simpler compared with the variance-covariance method (coming up next), since it does not require breaking down the probability of, and determining correlations between, market factors. The disadvantage is the need for a sufficient amount of historical data.

The *VARIANCE-COVARIANCE* method presumes normally distributed market factors to estimate the distribution of profit and losses of the FX contract with a mean $\mu = 0$ because the change in portfolio value for a short holding period is reasonably low. The standard deviation is obtained, e.g., through curve fitting of statistical data. The advantage of this

[55] For an introduction to the financial aspects of VaR the author refers to Linsmeier and Pearson [118] while an overview of its mathematical armamentarium is given by Holton [119].

[56] In an FX forward contract, counterparties agree to exchange a specified amount of different currencies at a fixed exchange rate at some future date. FX forwards are used to mitigate volatilities in foreign exchange currency for future transactions.

method is that, once the normally distributed curve has been obtained, then it is straightforward to calculate the VaR by multiplying the standard deviation with a constant reflecting the quantiles of interest. This approach captures the notions of variability and comovement of price changes in a variance-covariance matrix by the statistical concepts of standard deviation (variance) and correlation (covariance). A disadvantage of this method is that market prices do not necessarily have a normal distribution and may exhibit heavy tails with the occurrence of more frequent extreme values.

The *MONTE CARLO METHOD* is founded on the generation of a large number of simulations (single events). Each simulation is created through a combination of randomly generated values of market factors from their assumed probability distribution which eventually yield the estimates for the VaR. The advantage of the Monte Carlo approach lies in the freedom to choose an assumed probability distribution of the market factors while its drawback can be the computation time for the simulations.

Simons [120] states that VaR holds an important promise to construct one firm-wide measure of risk. For IS risk such a measure summarizes the maximum loss due to security breaches over a target horizon at a given level of confidence. This thesis does not follow up on VaR as estimating monetary losses is not in focus but future work faces the following challenges. First, the above methods yield varying results. Consequently, there is no "correct" VaR. Second, potential losses due to IS risks can neither be estimated nor mitigated by trading them in a market; they usually signify more than just a financial consequence and are specific to the company. Third, the original VaR was designed for measuring short term risks due to "usual" markets movements rather than estimating long-term risks rooted in a quickly changing IS environment. Fourth, long-term forecasts are naturally subject to model and parameter uncertainties (Dowd, Blake and Cairns [121]). Fifth, portfolios of instruments with probability distributions other than "normal" lack sub-additivity[57].

Analytic Hierarchy Process (AHP): The Analytic Hierarchy Process which goes back to Saaty [90] facilitates decisions given a set of alternatives and multiple objectives. In a first step, AHP generates a weight w_i indicating the relative importance of the i-th objective compared to others. All weights are displayed in a pair-wise comparison matrix adopting a scale from, e.g., 1 to 10. Upon estimation the matrix is normalized such that the sum of all weights is *1*. Then, a consistency index is calculated to verify whether the estimates are transitive. In a second step, a pair-wise comparison matrix is built reflecting alternatives which can be acted upon. Each alternative is assessed as to how well it "satisfies" or "scores" on the objectives. In step three, the alternative with the highest overall score is chosen.

[57] Refer to Artzner et al [122] for characteristics required to obtain Coherent Measures of Risk addressing sub-additivity.

AHP provides decision support where securing a product or process is one objective among others. As such it determines the relative importance of each objective to estimate how well it is fulfilled by a given set of alternatives. This selects an optimal solution in terms of relative importance; however, it does not provide mechanisms for asserting it and the assertion process is left entirely to the decision maker. Consequently, the question related to the "appropriate" level of security is unanswered.

In light of this brief analysis the author turns to Utility Theory, in particular, to Howard's Theory[58] of *RISK PREFERENCES* [123-125] to treat the outcome of uncertain events. First, many companies act as entrepreneurs in their market. They are subject to economic risk, which encompasses the risk-reward and insurance paradigms discussed in Chapter 1. As both paradigms base on the risk "appetite" of decision makers it is plausible to adopt such risk preferences as the main criterion for decision making in IS risk. Second, risk preferences implicitly reflect the appropriate level of security (rather than today's compliance management).

[58] The reader is referred to **Appendix E** for an introduction on Howard's Risk Preferences.

5.3 Decision Situation in IS Risk Management

The Process and Function Modules describe a scenario, i.e. an adverse course of events including its associated probabilities and consequences. In the Decision Module, a decision maker accepts or rejects a scenario where:

- accepting its associated risk leaves the scenario as it is, and
- rejecting its associated risk means to adopt security mechanisms[59] to alter a scenario.

Assumption 5.1 "Prospects": The adoption of a security mechanism offers two distinct prospects to the decision maker: it either succeeds or fails in protecting an information system.

Assumption 5.2 "Money Allocation": Allocating money for a security mechanism is irreversible and — should the security mechanism fail — the allocated money is lost in addition to the adverse consequence of the scenario.

Decision makers oftentimes face an ill-specified Decision Problem because the:

- scenarios are not fully specified or developed,
- consequences relative to a specific scenario have not been explored,
- success probabilities of threats are based on hear-say,
- relationship between security mechanisms and security processes is unclear
- risk preferences of the decision makers are unknown.

Assumption 5.3 "Commissioning Risk Analyses": A decision maker mandates a risk analysis to enhance the specification of the Decision Problem. However, risk analyses are costly and decision makers only buy the additional information if it is worth the extra expense.

In turn, the risk analyses will:

[59] The security processes that set up the security mechanisms maintain and decommission them are also implied.

- evolve scenarios both in the information system and business contexts,
- assert the consequences and probabilities of threats succeeding, i.e. of security mechanisms failing,
- assert the influence of security processes onto the probabilities,
- assert the risk preferences of the decision maker,
- create a list of alternative security mechanisms, which could be used for mitigation.

Assumption 5.4 "Single Decision Maker": Finally, in the realm of his or her functional area, a single decision maker is assumed. Groups of decision makers (bodies) are explicitly excluded from consideration as entrepreneurial practice shows that such governance bodies and steering committees blur the individual accountabilities of its members.

5.4 Using the Graphical Notation of the Process Module

In terms of the graphical notation of the Process Module, a standard decision situation for IS risk management is depicted as follows: In **Figure 5.1**, the XOR gates represent decisions on whether or not to adopt a specified security mechanism, e.g., security mechanism A. Accordingly, the event *1* signifies that the security mechanism was not adopted while the event *2* signifies that security mechanism A was adopted.

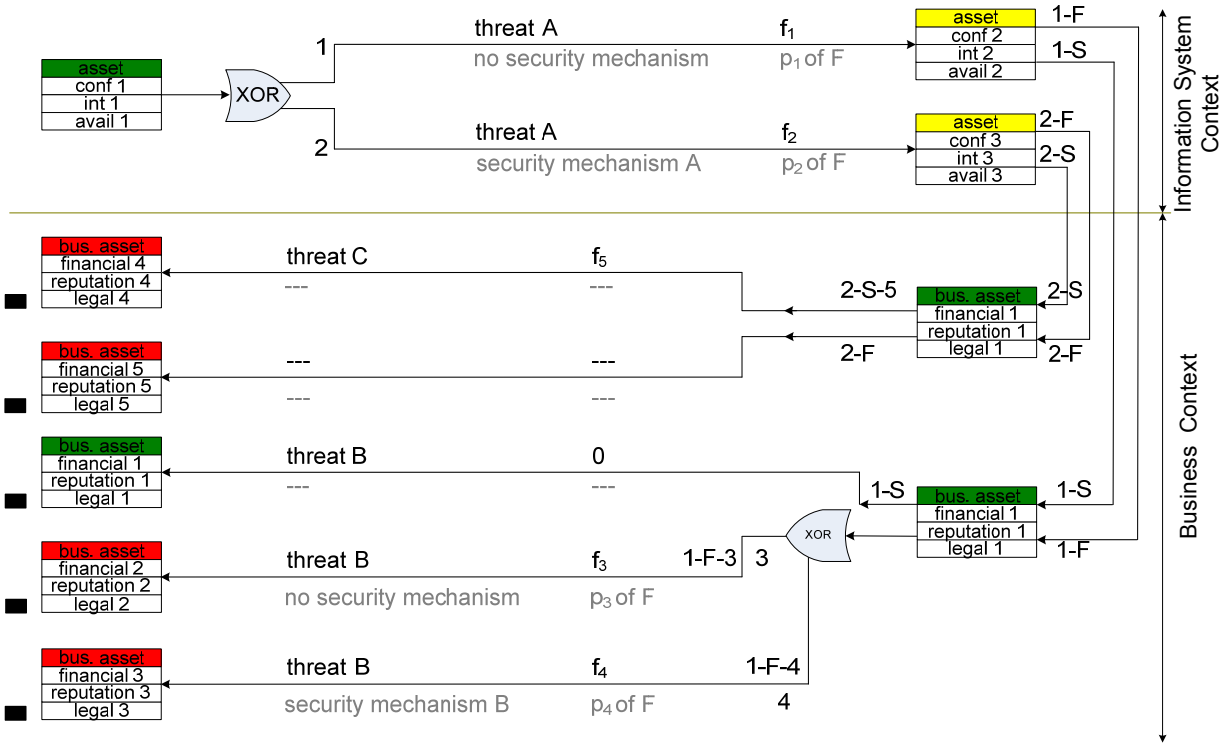

Figure 5.1: Decision Tree Expressed by Means of the Process Module.

In the upper part of **Figure 5.1** the green colour indicates an acceptable initial asset state while the yellow colour indicates that the asset has been compromised in the IS context. If event *1* is followed up, then not adopting security mechanism A may yield a failure with probability p_1 (the failure is denoted by event *1-F*, i.e. **threat A** with a frequency of occurrence of f_1 does prevail) or a success with probability $1 - p_1$ (the success is denoted by event *1-S*, i.e. **threat A** does not prevail).

Analogously, if event *1-S* is followed up into the lower part of **Figure 5.1**, then the business context is entered where **threat B** on a business asset does not occur (the frequency is *0*), no security mechanism is present and no further decision is required (no XOR gate). The other events can be interpreted likewise.

With some loss of information the above situation is depicted by decision trees, e.g., as used by Howard [126] (see **Figure 5.2**). They neither display the initial and end states of the assets under investigation nor the threats nor the security mechanisms counteracting them.

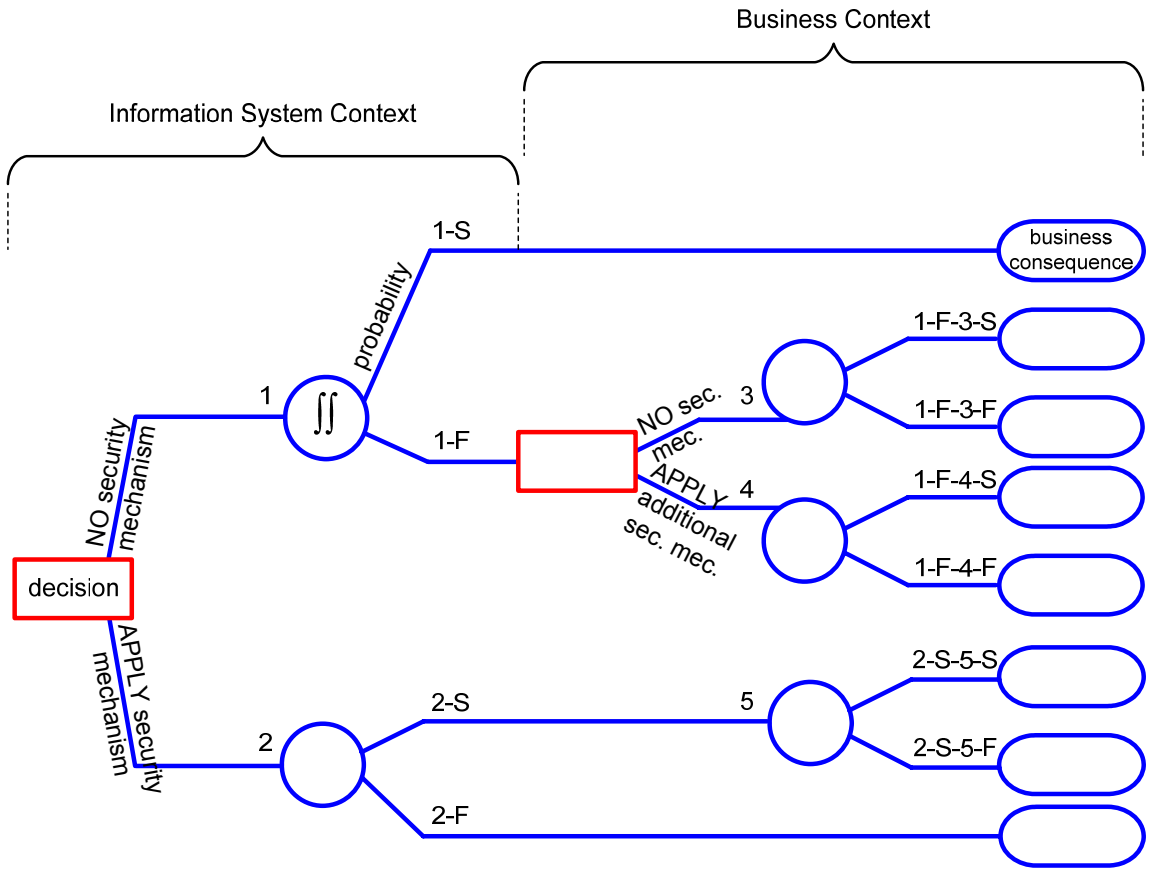

Figure 5.2: Decision Tree for IS Risk Management.

The above decision tree displays both the IS and the business contexts, proposes the selection of a security mechanism instead of alternatives, considers the probabilities of events[60] and denotes whether a security mechanism was successful in controlling its opposing threat where "S" means "successful" and "F" means "failure". The business consequence is reported on the right hand side of **Figure 5.2**.

Both notations can be used in the Decision Module depending on the information needs of its intended audience. On one hand, Howard's Notation may be better suited for communication of decision alternatives to senior executives as it only unveils the business consequence. On the other hand, the extended notation of **Figure 5.1** may be more apt for use with IS personnel and risk analysts.

[60] The symbol " ∬ " denotes the convolution with subsequent integration of probability densities for threat and security mechanism.

5.5 Using the Five Axioms of Utility Theory in IS Risk Management

The Five Axioms[61] by Howard [124] lay the basis for Utility Theory to be applied to a decision situation in IS risk management, i.e. to elicit the risk preferences of decision makers and the scenario information gathered by the risk analyst. In the following the applicability of the Five Axioms to IS risk management is discussed.

Orderability Axiom: The *ORDERABILITY AXIOM* requires a decision maker to arrange consequences *c* (prospects) of a scenario in descending order and that transitivity holds (if $c_A > c_B$, $c_B > c_C$, then $c_A > c_C$).

Intuitively, this axiom is easily fulfilled and thus applicable as consequences in IS risk are oftentimes expressed in financial terms.

Continuity Axiom: By accepting the *CONTINUITY AXIOM*, the decision maker must be able to compare two consequences in terms of a theoretical construct called certain equivalent. Let a security mechanism succeed (or fail) with a probability p (or $1 - p$) when counteracting a threat. The consequence of the security mechanism succeeding (or failing) is also given, c_S (or c_F). The decision maker must now be in a position to select a consequence c_E given it realizes with probability 1 such that the decision maker is indifferent in choosing between the certain equivalent c_E and c_S with probability p (c_F with a probability of $1 - p$ respectively). Alternatively, the decision maker may assign preference probabilities p_p ($1 - p_p$ respectively) with which s/he would like the c_S (c_F respectively) to occur such that s/he would be indifferent in choosing between c_S with probability p_p (c_F with a probability of $1 - p_p$ respectively) and a certain equivalent c_E.

Being a theoretical construct this axiom is easily fulfilled because c_E (or c_F) or p_p (or $1 - p_p$) can easily be obtained from the decision maker.

Substitutability Axiom: The *SUBSTITUTABILITY AXIOM* specifies that c_S (c_F respectively) be substitutable by c_E.

This axiom is fulfilled.

Decomposability Axiom: The *DECOMPOSABILITY AXIOM* states that only the final consequence pertaining to choosing a specific security mechanism be taken into consideration for decision making.

This axiom is easily fulfilled as most consequences are expressible in financial terms.

[61] The original Five Axioms are reported in detail in **Appendix E**.

Monotonicity Axiom: The *MONOTONICITY AXIOM* requires the decision maker to choose the security mechanism displaying the highest probability of achieving the best prospect. S/He is considered a *homo oeconomicus* pursuing economic goals and maximizing utility.

5.6 Determining the Maximum Price for a Risks Analysis

Prior to making a decision, the decision maker ascertains whether the decision situation is well or ill specified. If the decision situation is ill specified then additional information is required and the decision maker needs to assess how much s/he is willing to pay for it. Such an assessment is performed by using the concept of *CLAIRVOYANCE*; an analytical construct introduced by Howard [124]. Clairvoyance symbolizes an ideal, perfectly specified decision situation where all information is complete and reliable. The value of clairvoyance is that it indicates the maximum price a decision maker should pay to achieve this perfect state. Next, the concept of clairvoyance is adapted to suit decision situations of IS risk management.

Description "Clairvoyant": A clairvoyant identifies all relevant scenarios (or decision trees) and tells with certainty whether the events therein occur or not, i.e. the frequency of occurrence of events, their probability and their consequence. A clairvoyant does not work for free; s/he wants to be paid a price p_c for his services.

In IS risk the services of a clairvoyant correspond to risk analyses of analysts. Consequently, to find the maximum price of a risk analysis its decision tree is changed to reflect the possibility of buying clairvoyance:

- If the decision maker decides not to buy clairvoyance then the decision situation remains unchanged. This is depicted in the upper part of the decision tree in **Figure 5.3**.
- If the decision maker decides to buy clairvoyance, then the individual prospect values diminish by the amount charged by the clairvoyant (e.g., $v_{SA} - p_C$, see lower part of **Figure 5.3**).
- Let "A-S_A", "A-F_A", "B-S_B" and "B-F_B" be events investigated by the clairvoyant. From the point of view of the decision maker, these events are included in the same decision situation prior to consulting the clairvoyant (upper part of decision tree) and are thus marked with the probabilities p_A, p_B and $1 - p_A$, $1 - p_B$ which were previously known. This is denoted by the bold green text in **Figure 5.3**.
- Once the clairvoyant is going to report, e.g., "A-S_A" then A-SA is going to happen, i.e. its probability in the decision tree is *1*. This is denoted in bold blue text. An analogous procedure applies for the other events in the lower part of the decision tree.

The decision situation now reflects the possibility to buy clairvoyance. Next, the following steps are applied:

1. The utilities for the two main alternatives (no clairvoyance vs. buy clairvoyance) are calculated by determining the utility of the individual prospects (e.g., u_{SA}, u_{FA}, etc.).
2. They are then multiplied with their probabilities and the values are summed up for each lottery (e.g., the bold red "*u(A)*" in **Figure 5.3**). This is repeated for all alternatives yielding *u(no clairvoyance)* and *u(clairvoyance)*.
3. By comparing *u(no clairvoyance)* of the upper tree with *u(clairvoyance)* in the lower tree an indication is obtained on whether or not to buy clairvoyance at the price p_C .
4. Increasing/decreasing the price p_C until the decision maker is indifferent in choosing between *u(A, B)* and *u(clairvoyance)* determines the value of clairvoyance.

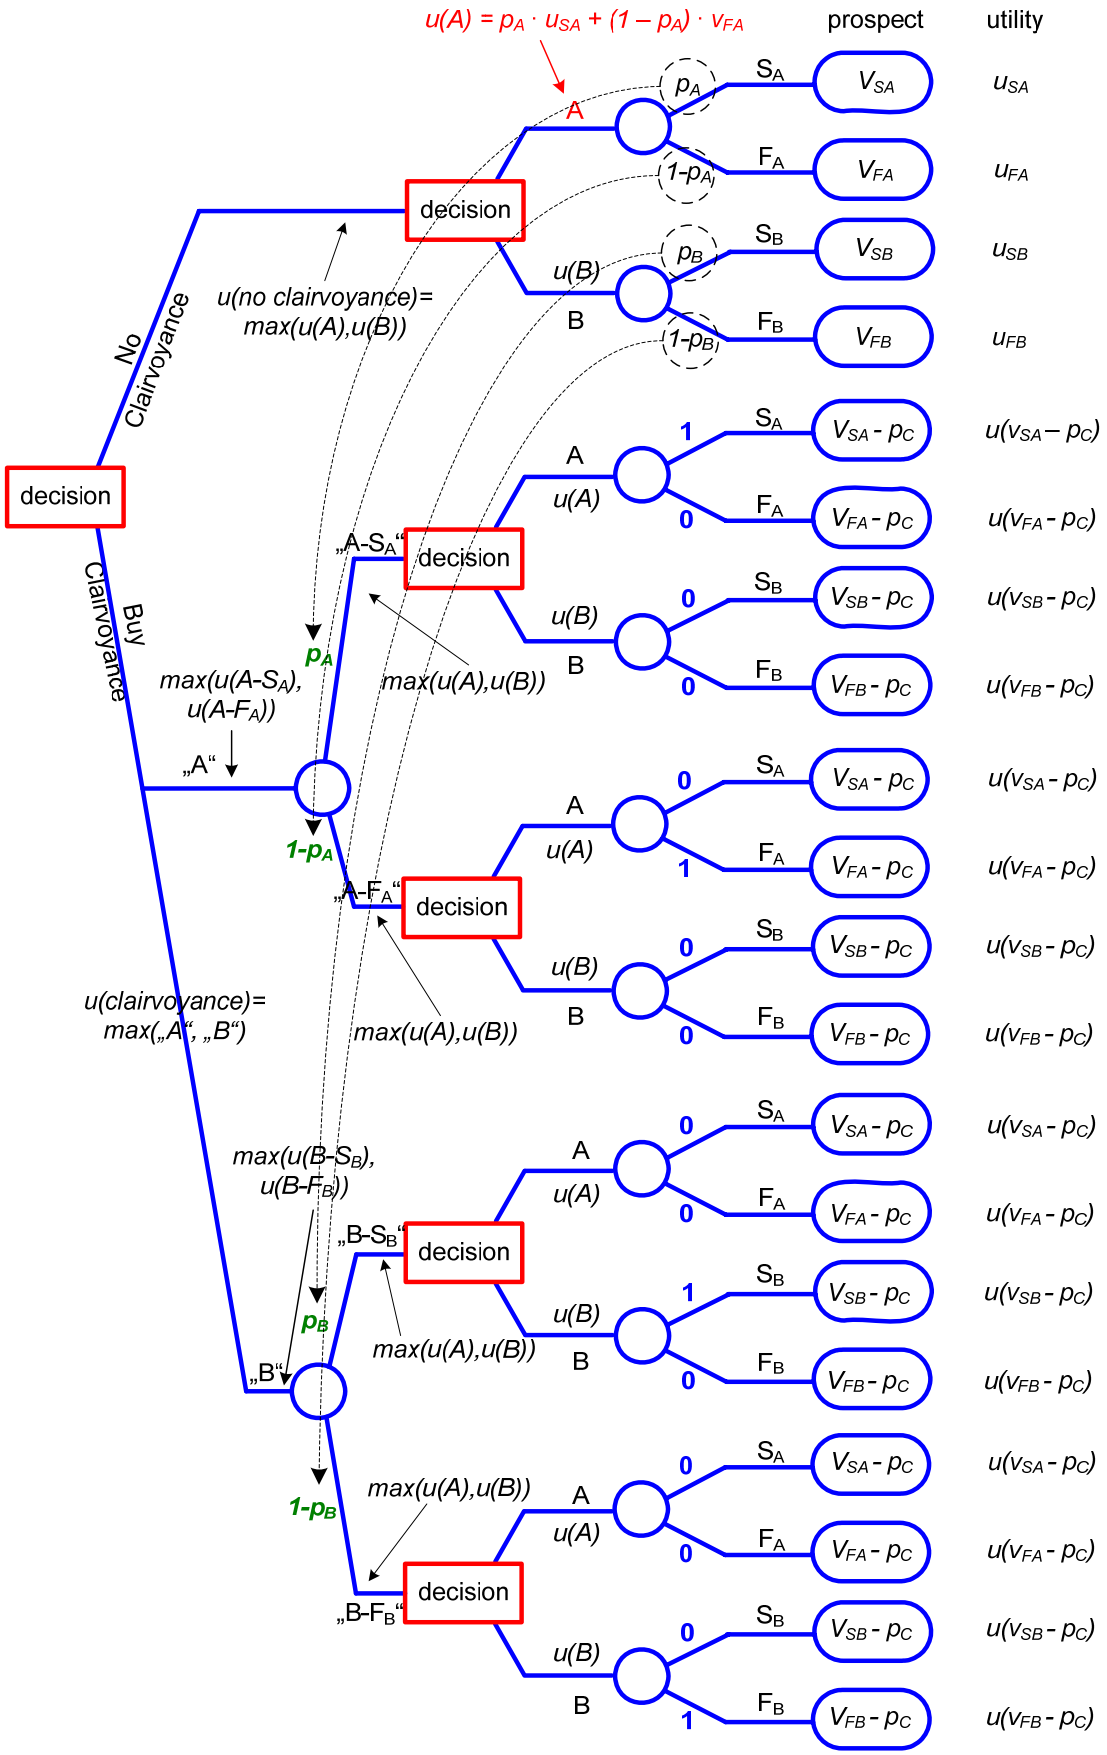

Figure 5.3: Valuing Risk Analyses by Means of Clairvoyance.

5.7 Application Example

5.7.1 Biometric, Weak or Strong Authentication

To secure a company from the threat of inwardly disclosing business information, a corporate decision maker asks security experts to elaborate on authenticating employees by means of:

1. biometric authentication,
2. user ID and password (weak authentication),
3. user ID, password and a physical token (strong authentication).

After fruitful consideration, the security experts present the decision maker with the following information:

- biometric authentication (*B*) costs *1.50 monetary units*; it fails with probability *0.05*
- weak authentication (*W*) costs *0.30* and saves *1.20* compared with (*B*); it fails with probability *0.32*
- strong authentication (*S*) costs *1.20* and saves *0.30* compared with (*B*); it fails with probability *0.01*.

In case an authentication mechanism fails the amount of the loss is assumed at *-1000* for all scenarios. In addition, the price for the respective authentication mechanism in place at the time of the failure must be added to the loss figure. In the example, losses in reputation and legal are deemed not relevant. Overall, the following loss figures are obtained should the respective authentication mechanism fail:

- *-11.50* for biometric authentication
- *-10.30* for weak authentication
- *-11.20* for strong authentication

The above figures are transposed into the notation depicted by **Figure 5.4**:

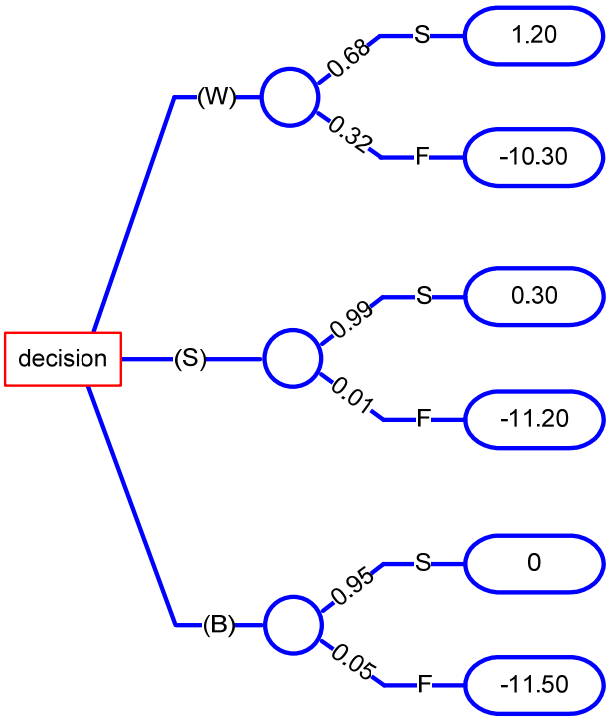

Figure 5.4: Selection of Authentication Mechanisms (Howard's Notation).

5.7.2 Utility Curve of the Decision Maker

The utility curve of a fictitious decision maker "Bossert" has been assessed. Her utility curve in **Figure 5.5** shows a slight risk preferring attitude for positive monetary values ≤ *6* while her attitude is risk averse for negative monetary values between *-6* and *-12*.

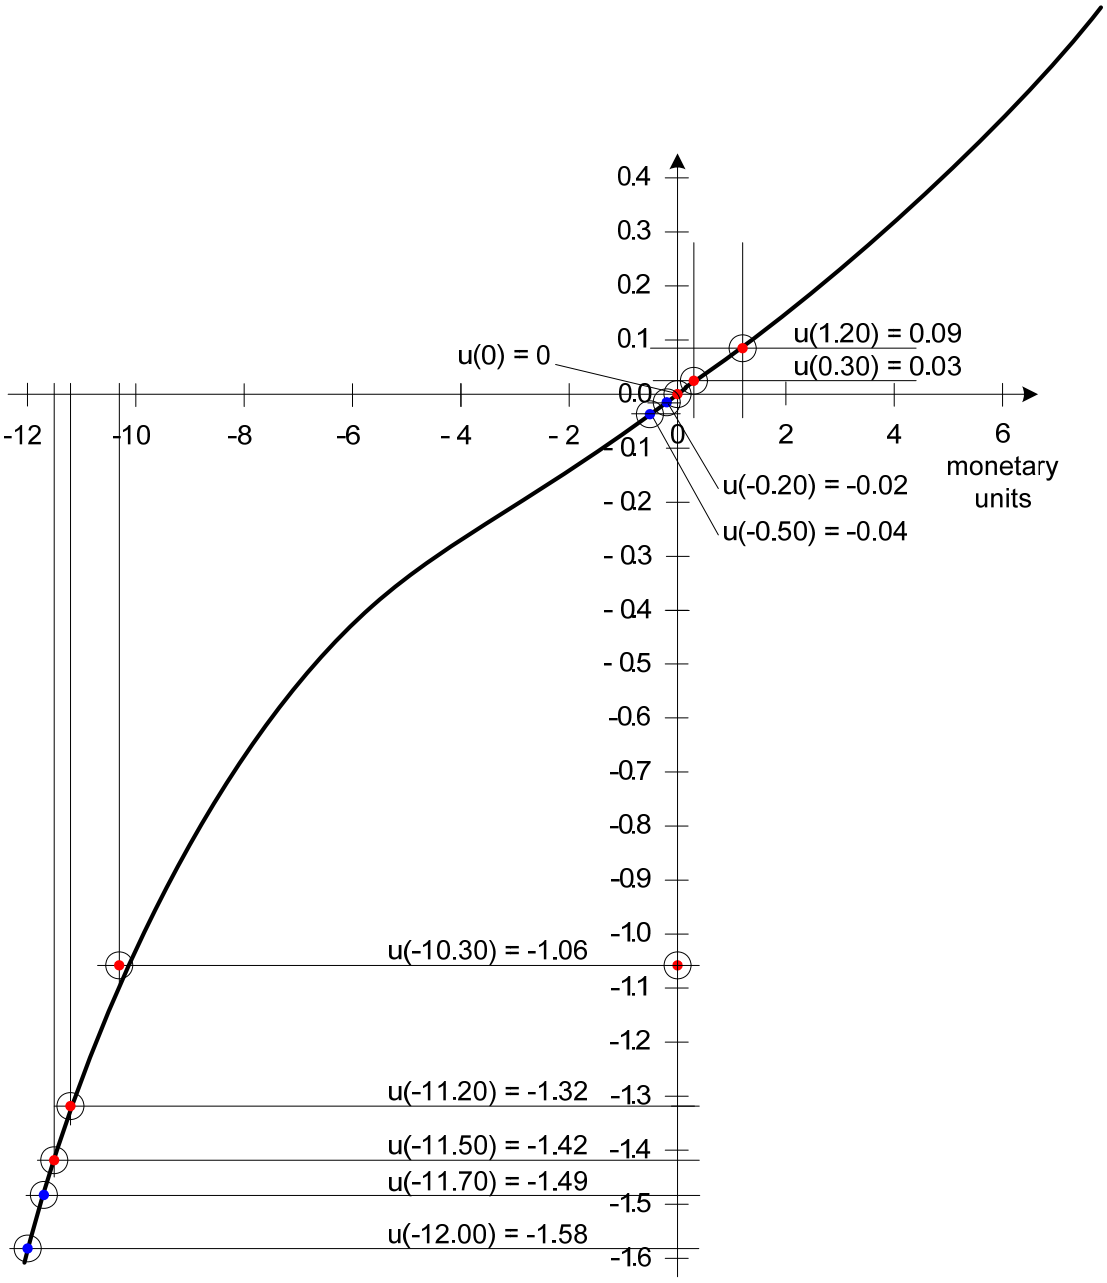

Figure 5.5: Utility Curve for "Bossert".

5.7.3 Value of Risk Analysis

In the example, security experts have debated on probabilities without reaching a conclusion. To obtain clarity, they intend to perform an additional risk analysis (costs: *0.5*). Will the decision maker mandate the additional risk analysis? To answer this, the decision tree in **Figure 5.4** is modified such to estimate the value of an additional risk analysis (see also **Figure 5.6**). For this modified decision tree the security mechanism "Weak Authentication" has been omitted as the probability of it prevailing over the threat is not acceptable. Which of the remaining alternatives is preferred by the decision maker?

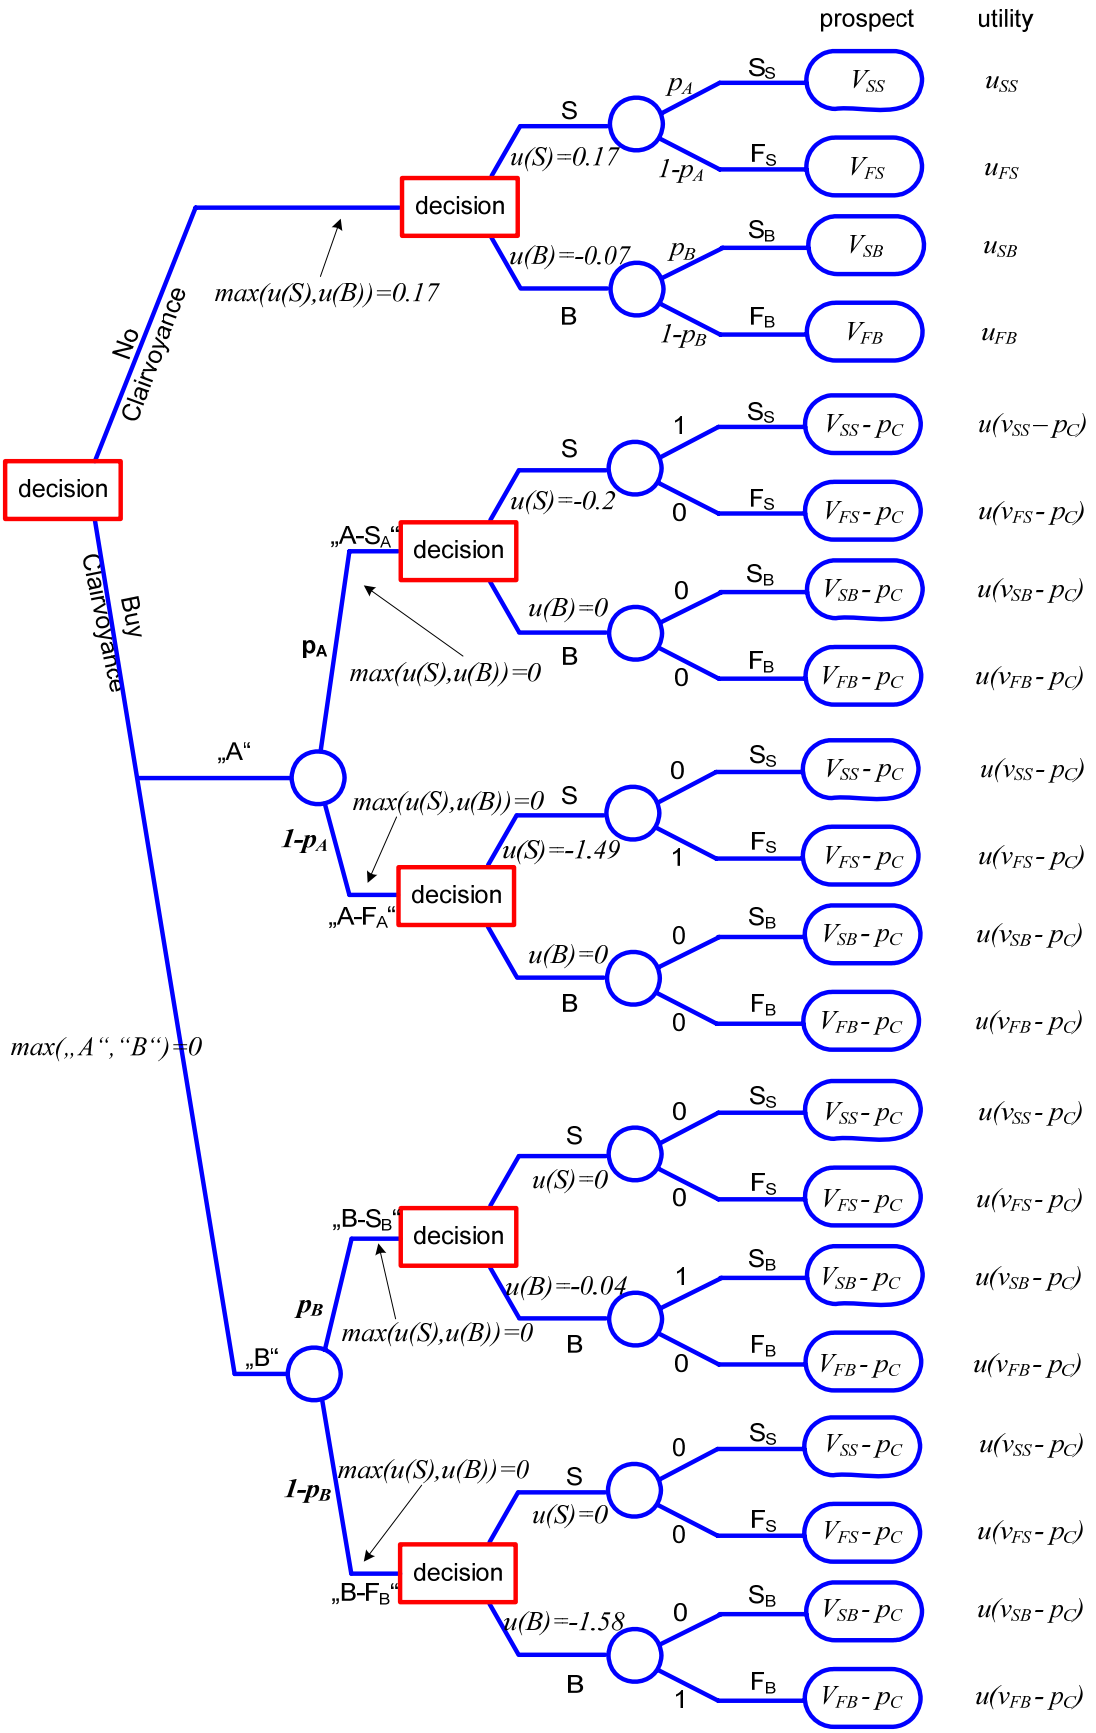

Figure 5.6: Valuing an additional Risk Analysis by Means of Clairvoyance.

5.8 Results

By using **Figure 5.5** and the consequences reported in **Figure 5.4** the following utilities for "Bossert" have been assessed:

Alternative	Utility
Biometric Authentication	***u(0.00, -11.50)*** *= 0.95 * u(0.00) + 0.05 * u(-11.50) = 0.95 * 0.00 + 0.05 * -1.42 =* ***-0.07***
Weak Authentication	***u(1.20, -10.30)*** *= 0.68 * u(1.20) + 0.32 * u(-10.30) = 0.68 * 0.09 + 0.32 * -1.06 =* ***-0.28***
Strong Authentication	***u(0.30, -11.20)*** *= 0.99 * u(0.30) + 0.01 * u(-11.20) = 0.99 * 0.03 + 0.01 * -1.32 =* ***0.17***

Table 5.1: Utility Figures for "Bossert".

The security mechanism "Strong Authentication" yields the highest utility of the three while "Weak Authentication" yields the lowest utility. Accordingly, "Bossert" should support the security mechanism "Strong Authentication".

The tree structure for the additional risk analysis is shown in **Figure 5.6**. Given a price for of *0.5*, the utility of an additional risk analysis is *0* while the utility for not performing an additional risk analysis remains at *0.17*. Consequently, the decision maker is better off not to mandate an additional risk analysis. The above is summarized in **Table 5.2**:

Alternative	Utility
No Risk Analysis	***max(u(S), u(B))*** *= max(0.17, -0.07) =* ***0.17***
Additional Risk Analysis	***max("A", "B")*** *= max(0, 0) =* ***0***

Table 5.2: Utility Figures for Additional Risk Analysis.

6. Overall Model

In Chapter 6, the Four Modules are welded together to the overall model which solves four key problems of today's IS risk management: the Ambiguity, Likelihood, Influence and Decision Problems. Each problem is addressed by one of the four modules, i.e. the Process, Function, Influence and Decision Modules respectively.

In contrast to practice where decisions in IS risk are given marginal importance (from a methodological point of view), in this work decision making is regarded as the core of IS risk management. Accordingly, the Decision Module is treated first and the other modules are regarded as data suppliers and follow in reverse order in Chapters 6.2, 6.3 and 6.4; Chapter 6.5 presents a graph of the overall model.

6.1 Decision Module

To facilitate a decision, the Decision Module adopts *RISK PREFERENCES* of individual decision makers. As such it presupposes the decision maker's utility curve, which is obtained by assessing his or her certain equivalents to business consequences. These losses or gains are induced by security mechanisms failing or successfully withstanding threats.

The business *CONSEQUENCE* is assumed in the Process Module (Chapter 3) along with the threat events, which put at risk the assets of interest. Although consequence was defined in terms of financial, reputation, and legal values, for the rest of the thesis it will be interpreted as financial losses or gains only.

Consequently, the *PROBABILITIES* of threats succeeding in overcoming the company security mechanisms are needed. To estimate them, a general approach in the Function Module has been introduced (Chapter 3). Moreover, as the occurrence of perfect vulnerabilities and controls in information systems cannot be ruled out, the *FREQUENCY* of occurrence of threats has been introduced as a measure to indicate a general protection strategy. In contrast to probabilities, no specific approach to determine frequencies of threats has been provided other than by observation and information gathering.

Security mechanisms are operated within a context where they are maintained by security processes. By evidencing the *INFLUENCE* of security processes on security mechanisms the decision maker is given the opportunity to take into account the context

of a security mechanism. As outlined in the governance example in Chapter 4, we often face the situation of not being in a position to directly shape the security mechanism but must do this via security processes. Consequently, knowing which security processes influence a security mechanism is important for decision making in practice. Eventually, along with the implementation and maintenance costs of the security mechanisms, the respective figures for the significant security processes are transferred to the Decision Module.

The above is displayed in **Table 6.1**:

Input to DM	from Module	Output from DM
Utility curve of decision maker (reflects risk preferences)	Decision	Among alternative security mechanisms and their security processes, the Decision Module indicates those security mechanisms and processes to choose, which offer the highest expected utility to the decision maker.
Consequence (in terms of financials only)	Process	
Probability of threat overcoming its opposing security mechanism	Function	
Frequency of occurrence of threat (sets general protection strategy)	Function	
(Implementation and maintenance) costs of security mechanism and security processes	Influence	

Table 6.1: Input to and Deliverables of the Decision Module.

6.2 Influence Module

In the Influence Module, security processes are evaluated in terms of their implementation quality[62], which is then related to the probability density of an event. This means that it must be verified whether the implementation quality of security processes relates either to the probability density of a threat or to the probability density of "its" security mechanism (or both). In essence, the Influence Module searches for security processes,

[62] The term "implementation quality" is not defined further as it must allow for some degree of vagueness to reflect heterogeneity in large companies. Consequently, further refinement is left at the discretion of the risk analyst of the individual company.

which are *SIGNIFICANT* for the set up, maintenance or decommissioning of threats or security mechanisms. If this search is successful, then the cost of the significant security process(es) is taken into account when evaluating, e.g., a specific security mechanism. The aim of this exercise is to support the decision set prioritized by its influence on success probabilities on attacks. This information is considered in the Decision Module.

The above is displayed in **Table 6.2**:

Input to IM	from Module	Output from IM
Implementation quality of security processes	Influence	The maintenance costs of significant security processes are forwarded to the Decision Module.
Probability density of threat	Function	
Probability density of security mechanism	Function	

Table 6.2: Input to and Deliverables of the Influence Module.

6.3 Function Module

The Function Module determines the probabilities of threats overcoming security mechanisms. Threats as well as security mechanisms are represented by temporal frequency curves; they are then fitted with probability density functions and next they are convoluted to obtain a resulting curve, the probability density of the event. The probability is represented by the cumulative distribution function of the threat event. At the basis of this computation is the insight that a threat is successful if it attacks a security mechanism before it has been updated to oppose the threat.

Probability is not the only important piece of information a decision maker requires for the selection of security mechanisms and the justification of their costs. In fact, as outlined in Chapter 3, the frequency of occurrence of a threat also needs to be known to determine a *GENERAL PROTECTION STRATEGY*.

The above is displayed in **Table 6.3**:

Input to FM	from Module	Output from FM
Frequency of occurrence of threat	Function	General protection strategy, probability density functions of threat & security mechanism; probability of threat event.
(frequency curve of) threat	Process	
(frequency curve of) security mechanism	Process	

Table 6.3: Input to and Deliverables of the Function Module.

The threat events to be analyzed in the Function Module are selected in the Process Module.

6.4 Process Module

The Process Module describes scenarios both in the information system and business contexts. Scenarios reproduce events and assets in the information system context that threaten assets in the business context. They are chosen in pivotal fashion according to the information needs of the decision maker. In the Process Module, the possibility of gaining a limited insight into one possible future by evolving scenarios along two axes of a chart has been hinted at. In the scenario chart, various scenarios are put into perspective. They are an effective tool to display potential future developments along the two axes. However, they fail if the axes of choice are not the relevant ones, i.e. if future developments suddenly occur along other axes. However, the Process Module does not claim completeness in selecting every event, which possibly threatens a specific business asset. In fact, it is rather intended as a tool for gaining a better understanding of today's and tomorrow's threat landscape related to a business asset. As such it supports communication to senior executives.

The above is displayed in **Table:6.4**

Input to PM	from	Output from PM
Selected threat events and business assets	Decision Maker	Scenarios describe impacts in the IS context and consequences in the business context. Scenario charts show the completeness of today's understanding and potential future developments.
Axes of scenario charts	Functional Experts with Decision Maker	

Table 6.4: Input to and Deliverables of the Process Module.

6.5 Overall Model

In **Figure 6.1**, the interaction of the Four Modules is visualized and is signified by the four large boxes. Each module solves one problem of today's IS risk management and the approach chosen has been hinted at, i.e. state machines for the Process Module, the converse convolution of the threat strength vs. the resistance of security mechanism for the Function Module and pattern recognition by means of Rough Sets data tables for the Influence Module. Finally, risk preferences have been expressed by means of a utility function in the Decision Module.

The four smaller boxes signify incoming information from outside of the model:

- blue stands for the information provided by the decision maker
- red stands for empirical data, which is gained through measurement
- black stands for (intermediate) results of a module.

Information Flow for the Process Module: The incoming blue arrows to its left signify the initial information provided by the decision maker. Ideally, it contains:

1. the events and the associated assets for both the information system and the business contexts

❷ the information needs of the decision maker, which became the axes of a scenario chart(s)

The Process Module produces a list of scenarios:

❸ individual scenarios and their relative position to each other in a scenario chart

❹ initial and end states of the assets in the business and IS contexts; including the related threat events

❺ consequences related to the business asset(s) being investigated.

Information Flow for the Function Module: The incoming red arrows in the lower part signify empirical data gained through measurement. Ideally, it contains the:

❻ frequency curves of threats and a description (e.g., of the threat agent)

❼ frequency curves of security mechanisms in place at the company and a description (e.g., of the resilience properties).

The Function Module produces:

❽ estimates on the frequency of a threat (through observation) and the probability of an event (through calculation)

❾ probability densities of the threat and the company security mechanisms.

Information Flow for the Influence Module: The incoming arrows in the lower part signify:

❾ probability densities of threats and security mechanisms

❿ implementation quality of security processes (empirical data).

The Influence Module produces a:

⓫ set of significant security processes influencing either the threat or the security mechanism or both.

Information Flow for the Decision Module: The incoming arrows signify:

❽ estimates on the frequency of a threat (through observation) and the probability of an event (through calculation)

⓫ setup, maintenance and decommission costs of security processes

⓬ setup, maintenance and decommission costs of security mechanisms

⓭ a set of security processes ranked according to influence

⓮ risk preferences of the decision maker.

The Decision Module produces a:

⓯ list of scenarios ranked by their utility for the decision maker.

6.6 Ten Steps to Applying the Four Modules

Each module stands on its own and no specific order is prescribed with which to apply the overall model. However, for more complex decisions the risk analyst may apply the following steps:

Step I: Assess the risk preferences of the decision maker.

Step II: According to the value of information for the decision maker, determine the maximum amount of money to be spent for a risk analysis. Consider this maximum price for all consequent steps.

Step III: With the decision maker, select the scenarios of interest for the decision maker, i.e. the threat events.

Step IV: Indentify the threats and their opposing security mechanisms. Measure their individual frequency curves and calculate probabilities.

Step V: Select the processes whose influence on probabilities is of interest. Implement the processes (if needed) and execute the measurement.

Step VI: In order to obtain a second iteration cycle (e.g., for assessments where a change over time is followed up), repeat steps 4 and 5 after a period of time then proceed with the next step.

Step VII: Compare frequency curves of first with second iteration cycle.

Step VIII: Calculate the influence of security processes on frequency curves (both iterations).

Step IX: Complete the decision tree with information gathered so far. Calculate utilities of individual security mechanisms.

Step X: Select the alternative with the highest utility.

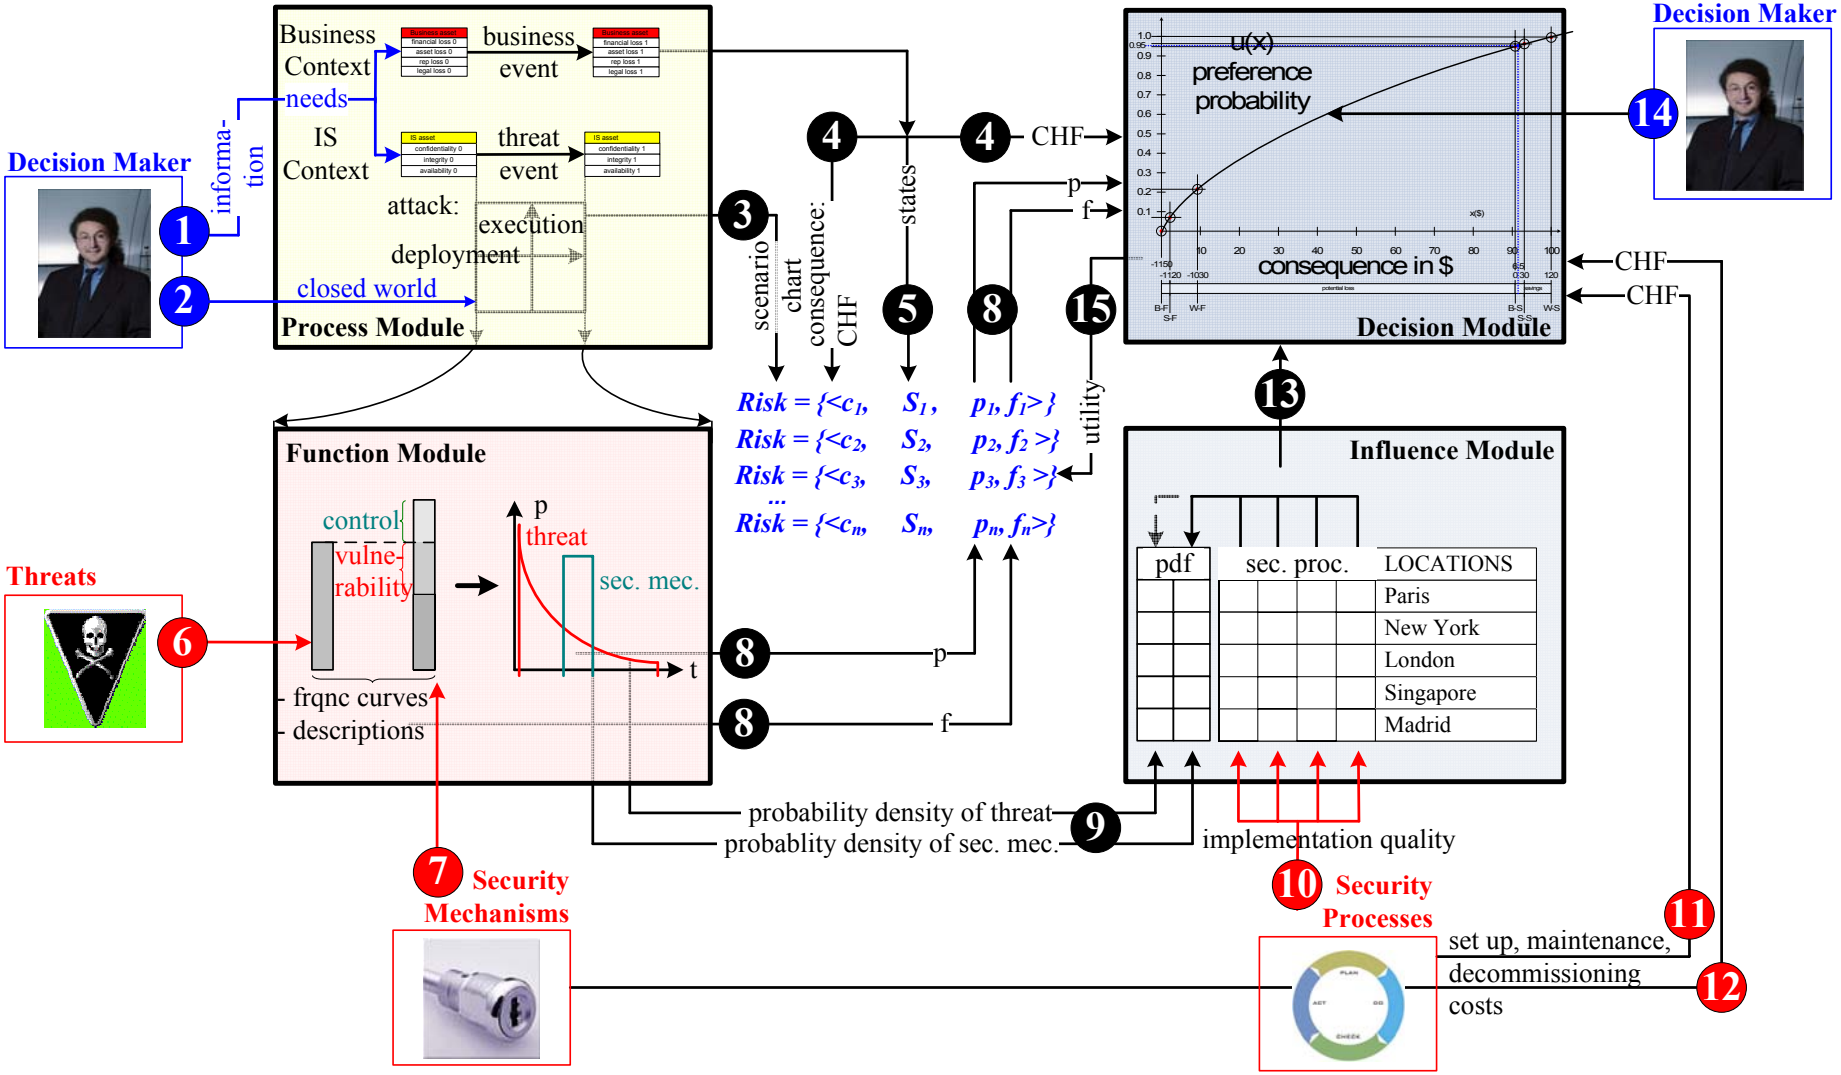

Figure 6.1: Overall Model.

7. Case Study on Phishing

This case study was mandated by a global financial institution. For its assembly, the author had the opportunity to shadow experienced practitioners from banks, one government agency and specialized security consultants. The Four Modules are applied to elaborate on phishing attacks mainly from a company perspective in the IS context of the banking industry.

Following this introduction, the state-of-the-art of phishing attacks is outlined. In Chapter 7.2, two phishing scenarios are described using the graphical notation of the Process Module. In Chapter 7.3, the success probabilities of phishing attacks are calculated by means of the Function Module. In Chapter 7.4, the Influence Module is applied to determine the influence of security awareness training onto phishing attacks. In Chapter 7.5, the Decision Module suggests specific security awareness training to counteract phishing.

7.1 State-of-the-Art of Phishing Attacks

Divulging criminal energy via the Internet is not inhibited by the physical presence of a victim; phishing attacks profit from this anonymity. They are executed to steal personal information for subsequent abuse and illegitimate monetary gain. Phishing attacks involve three parties: organized Internet criminals (the attacker), the company, and the customer of the company (the victim). The *corpus delicti* ranges from electronic banking authentication details, credit card details, social security numbers, customer loyalty card numbers, passport numbers, etc. To this end, phishers either scam their victims, spy on them, or both.

For phishing attacks, Cain [127] distinguishes six phases:

1. In the planning phase, the question is “what to steal how from whom?”
2. In the set up phase, attack material and machinery is created and contact information of potential victims is obtained.
3. In the attack phase, contact with prospective victims is made.

4. In the collection phase, the user credentials are stolen.
5. In the fraud phase, the attacker sells the stolen credentials or abuses them to illegitimately transfer money for personal enrichment. The money transfer is usually followed by money laundering.
6. Finally, in the post-attack phase, the attacker is concerned with deactivating the phishing machinery, covering his tracks, assessing his success and monitoring responses to the attack.

Phishing is not ignored by companies. They are involved in an arms race with organized Internet Crime which requires ongoing vigilance. In response to phishing, many interest groups have emerged such as the:

- Antiphishing Working Group, see antiphishing.org
- International Identity Theft Technology Council, see cyber.st.dhs.gov/ittc.html
- Special Interest Group (SIG) on Identity Theft, see projectliberty.org.

In this struggle, which is eroding trust in the Internet channel, companies have chosen to approach their customers, e.g., the Bank of New York [128] informs customers via informational leaflets. However, many companies remain silent when it comes to the consequences caused by phishing. In the financial industry, monetary losses suffered by victims are usually refunded in-turn for signing a non-disclosure agreement. Consequently, phishing attacks find moderate attention by mass media, although they have been ongoing for a couple of years.

There are many forms of phishing and some commonly found variants of it shall be explored:

Description "Classic Phishing": *CLASSIC PHISHING* starts with randomly sent e-mails. These e-mails lie to or pressure to victims in order to obtain their Internet banking login credentials. As many users do not comply with the fraudulent request, a great number of potential victims are targeted. For example, in April 2004, Warner [129] reported that 3.1 billion phishing e-mails were sent out worldwide.

In response to the e-mail, the victims may surrender their login credentials. Then, the attacker impersonates a victim and commits some type of electronic fraud; usually an illegitimate money transfer from the victim's Internet banking account to an accomplice who

immediately cashes out the money and forwards it to another accomplice. This forwarding uses channels other than Internet banking to blur the trail of the money. At the end of a chain of accomplices the attacker awaits the money.

The accomplice is often unsuspecting of these money laundering activities he or she is involved in as s/he was lured into forwarding the money to the attacker by some false pretext. To emphasize this ignorance the accomplice is also called the "money mule".

Description "Spear Phishing": *SPEAR PHISHING* is a colloquial term used to describe personalized attacks targeting a specific group or individuals. Spear phishers use available information on specific groups or individuals to make e-mails appear genuine. Because of this, the victims comply with the request in the e-mail to, e.g., download and execute malicious software which spies on login credentials. This variant of phishing is often used for industrial espionage involving, e.g., stealing information on salaries, figures of the company's financial statements, etc.

Description "Vishing"[63]: Phishers use sometimes (Internet) telephony to trick a victim into revealing their login information. The victim is invited to call a phone number provided by the attacker and believes they are talking to the bank's call center when, in fact, the victim is talking to the attacker. Because of this, the victim surrenders Internet banking credentials upon request to the supposed call center. This kind of attack is called *VISHING* – a term used to denote "voice" and "phishing" and leverages IP-based voice messaging together with social engineering techniques.

Description "SMiShing": Phishing attacks are not always deployed by e-mail and do not always aim at Internet banking login credentials. The following transcription, for example, shows parts of an early SMS phishing attack reported by F-Secure [131] in April 2007: "*We're confirming you have won a prize. Please contact the following number to collect your prize …*" When calling the number, the victims were prompted to surrender the security codes for access to automated teller machines in order to transfer the prize. This form of phishing is called *SMiSHING* to emphasize the use of SMS.

[63] One of the first papers on "Vishing" was presented by Ollmann [130].

Description "Phishing with Malicious Software": According to a joint report of the US Department of Homeland Security, the International Identity Theft Technology Council and the Anti-Phishing Working Group [132], attackers increasingly apply sophisticated technical tricks. In principle, they attempt to intercept the connection line between the victim and the Internet banking computer by tricking the user into downloading some kind of malicious software on his or her computer prior to connecting to the Internet banking site. This so-called *PHISHING WITH MALICIOUS SOFTWARE* activates when the victim accesses the Internet banking to modify the beneficiary account number and the amount in financial transactions.

Description "Drive-by Infection": In another variant of phishing with malicious software, an attacker modifies a website to automatically exploit potential vulnerabilities in the browser of the victim. If successful, malicious software is inwardly installed on the computer of the victim to execute some kind of electronic fraud. The Swiss Federal Police [133] calls this a *"DRIVE-BY" INFECTION*.

Description "Pharming": In yet another variant, the attacker redirects a user to, e.g., a false Internet banking page instead of the genuine website. There, the unsuspecting victim surrenders his or her Internet banking credentials. Such attacks are also referred to as *PHARMING*[64].

Phishing works because it combines technical subterfuge with social engineering[65] which aims at deceiving an unsuspecting victim into performing actions he or she would otherwise not do. To counteract social engineering, security awareness training has been proposed. For example, Sheng et al [138] have developed "Anti-Phishing Phil", an online game that teaches users to recognize phishing e-mails based on a variety of cues. Kuramaguru et al [139] proposed user training to be embedded in the e-mail system of the user. However, according to Dhamija, Tygar and Hearst [140] participants of a study on strategies for deceiving general users *proved vulnerable across the board to phishing attacks ... neither education, age, sex, previous experience, nor hours of computer use showed statistically significant correlation with vulnerability to phishing*.

---

[64] For further information on "Pharming" refer, e.g., to Srivastava [134].

[65] For an overview on social engineering techniques see, e.g., Drake et al [135]. For examples on social engineering approaches refer to Jagatic [136] or Jakobsson [137].

The phishing community has specialized in complementary tasks: resources are allocated in specialized markets and Abad [141] has demonstrated that they are professionally applied in organized crime networks where phishers often use so called botnets[66], a virtual army of compromised computers, to send out spam e-mails, deploy technical subterfuge or hide the tracks of the attacker.

According to a study by Ramzan and Wüest [144] there are fluctuations in the number of attacks on weekends and during major events. For example, activities related to maintaining a botnet, decline considerably over weekends while they rise sharply during, e.g., Christmas, New Year, etc. Finally, the financial industry is targeted most (with some attempts on online merchandise sales amongst others) and the majority of the victims are English-speaking, elderly and students.

[66] For information on botnets refer to Dittrich [142] and for an overview on recruiting and operating botnets refer to Cole et al [143].

7.2 Process Module: Scenarios in an Information System Context

Assumption 7.1 "Attack Deployment and Execution": Phishing attacks are determined by the techniques used for their *DEPLOYMENT* and subsequent *EXECUTION*.

Accordingly, the "Number of Targeted Victims" (deployment) and the "Victim/Attacker Interaction" (execution) are designated to be the axes of choice for the scenario chart.

7.2.1 Number of Targeted Victims

Typically, techniques adopted for deploying phishing attacks delimit the maximum number of victims, which can be reached by the attacker. **Table 7.1** shows some examples of deployment techniques:

Deployment Technique	Description
via Internet worms	Internet worms automatically exploit security mechanisms in commonly installed software such as operating systems, browsers, etc.
via e-mails	The e-mail, usually sent in great numbers, is used as a carrier, which may contain malicious software, e.g., a virus, or simply delivers instructions to the recipient.
via website (software download)	Malicious software is injected into a website and is subsequently downloaded by an unaware victim.
via website (via drive by infections)	Drive-by infections occur on websites where the browser of the victim is checked for vulnerabilities, which are then automatically exploited. This technique reaches many more users than "deployment via download".
via affiliate marketing	Some affiliate marketing programs offer financial rewards to operators for installing malicious software on computers visiting their website (USD

programs	0.08 to USD 0.50 per computer according to Emigh [132]).
via exploits (hacking)	The attacker deploys the phishing attack manually and "hacks" specific target environments.
via personal phone calls	The attacker engages the victim in personal phone calls.

Table 7.1: Deployment Techniques for Phishing Attacks.

In practice, the above deployment techniques are often combined. They determine the number of victims: single victims, groups or mass victims **(Figure 7.1)**:

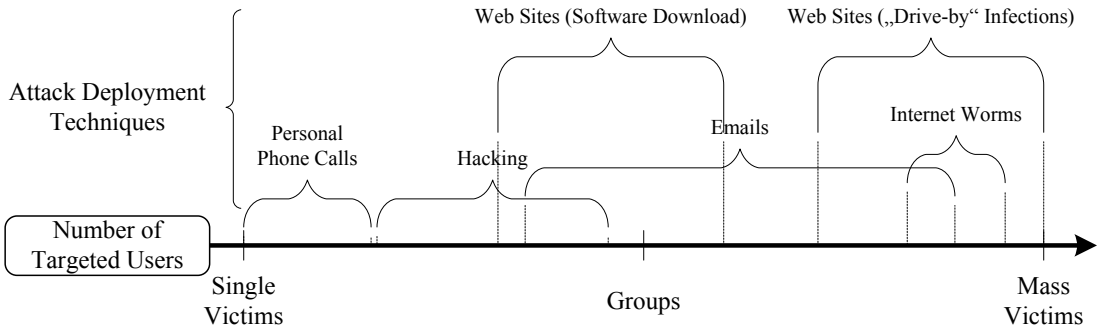

Figure 7.1: Deployment Techniques on the Axis "Number of Targeted Victims".

7.2.2 Victim / Attacker Interaction

Assumption 7.2 "Attacker/Victim Interaction": The technique adopted for executing the attack determines the interaction between the attacker and the victim.

Table 7.2 shows some techniques for executing phishing attacks once they have been deployed (next page):

Execution Technique	Description
via *SOCIAL ENGINEERING*	An attacker may • misrepresent him/herself as a large banking organization, • act to cause a technical problem, which troubles the victim and then offers himself to solve that problem the victim experiences, • personalize an e-mail with information found, e.g., in public data bases Alternatively, fear, greed, sex, etc. may be exploited.
via *KEYLOGGERS* and *SCREENLOGGERS*	A keylogger monitors sensitive data on a computer and forwards it to a phishing server. Screenloggers, in addition, monitor the graphical display of the computer to counteract potential on-screen security mechanisms. In May 2006, Emigh [132] reports *215* unique signatures of keyloggers.
via *E-MAIL* and *INSTANT MESSAGING REDIRECTORS*	Redirectors intercept outgoing e-mails and send a copy to an attacker. Likewise, instant messaging redirectors transmit transcripts of on-line messages to an attacker.
via *BROWSER REDIRECTORS* (hostname lookup attacks)	Browser redirectors modify the *HOSTS FILE* of the victim's computer. The hosts file matches IP addresses with domain/host names and is looked up by the browser before accessing the Internet. If the IP address of a requested domain/host name is in the hosts file then it is used. Conversely, if the desired IP address is not shown in the hosts file, a query to an external domain name server (DNS) is started. Redirectors change the entries in the hosts file to redirect the browser to the phishers website.
via *BROWSER REDIRECTORS* (DNS cache poisoning)	A browser queries the DNS located in its domain to resolve a URL (e.g., ethz.ch) into an IP address. If the domain of the requested URL does not lie within the same domain for which the DNS is authoritative, it will forward the resolution request to the DNS of that domain. Its response is cached into the memory space of the first DNS, which then sends it to the requesting browser. As DNS' usually don't authenticate each other, an attacker can insert a malicious DNS into the network, which "poisons"

	the caches of the other DNS with bogus information. This attack is also known as "zone transfer" or "pharming"[67].
via *SESSION HIJACKING*	Session hijackers typically reside in the browser of the user and take over the connection (session) on behalf of the attacker. Since the connection to, e.g., the bank was initiated by the user, the session hijacker can now impersonate the victim.
via *WEB TROJANS*	Web trojans pop up over login screens to collect credentials. The user is tricked into believing that he or she is entering information on a legitimate website while, in fact, the information is being collected for an attacker.
via *TRANSACTION GENERATORS*	Unlike the types of malicious software previously discussed, a transaction generator does not aim at gaining personal information from a victim. It resides within a computer at the transaction processing center of, e.g., a credit card company. There it generates fraudulent transactions.

Table 7.2: Execution Techniques for Phishing Attacks.

Execution techniques are often combined and determine the interaction between the attacker and the victim, i.e. "NIL", "click-to-install" or "actively surrender" information:

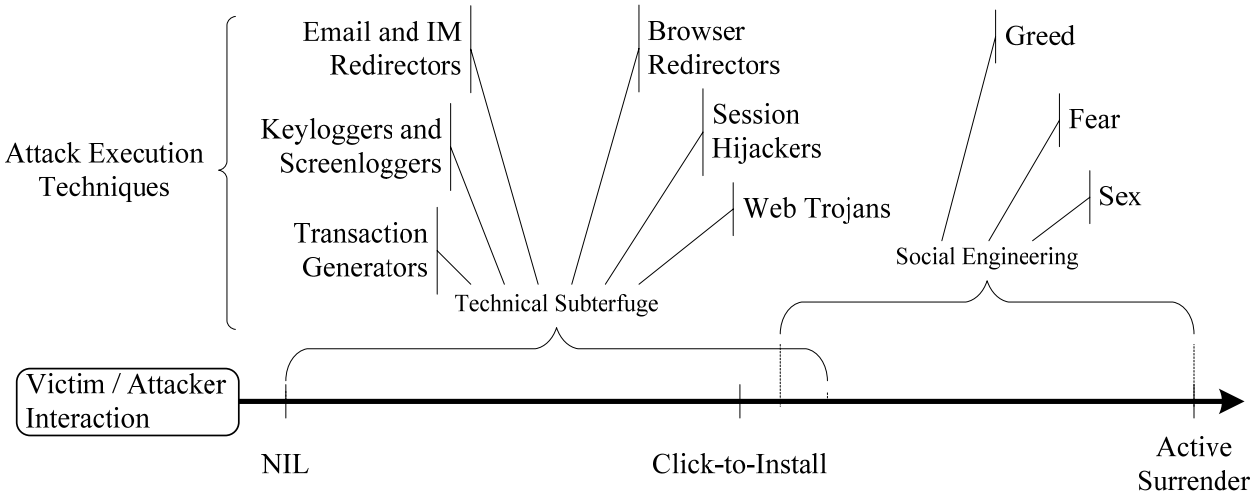

Figure 7.2: Attack Execution Techniques on the Axis "Victim / Attacker Interaction".

[67] Refer to Olzak [145] for further reading.

7.2.3 Phishing Chart

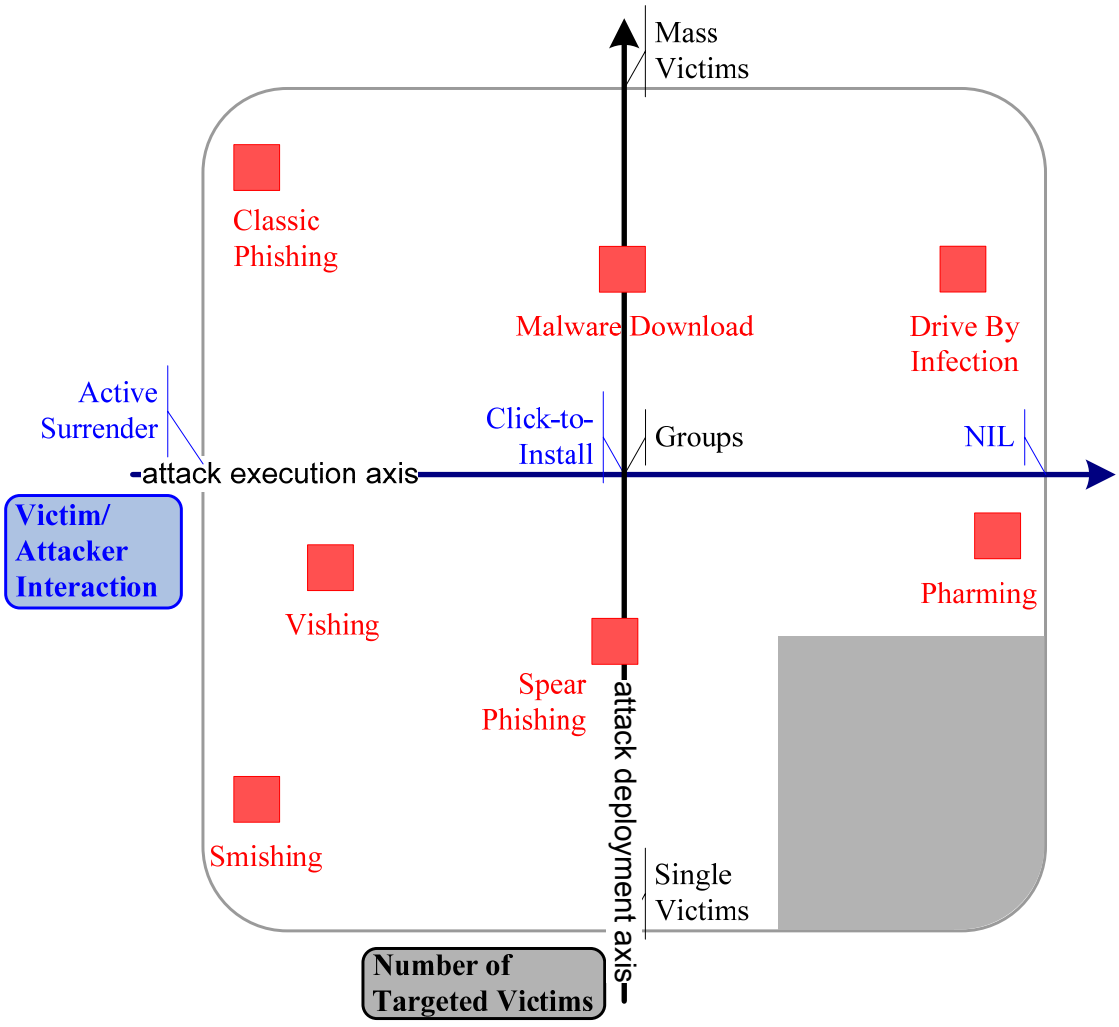

Figure 7.3: Scenario Chart Displaying Seven Phishing Attacks.

Figure 7.3 shows the positioning of the aforementioned seven phishing attacks in an IS context as well as relative to the axes "Number of Targeted Victims" and "Victim/Attacker Interaction". The placement of the attacks has been verified by various secu-

rity experts (e.g., by representatives of MELANI[68]). The grey colour evidences "attack-free" areas of the chart.

Remark: Instead of displaying the attack dynamics, the scenario chart may show axes evidencing security mechanisms. However, this does reflect the special interest of the remitter of the case study and is not followed up.

7.2.4 Applying the Process Module Notation

In **Figure 7.4**, the seven phishing scenarios are depicted in terms of the Process Module**.**

[68] MELANI is the Reporting and Analysis Centre for Information Assurance of the Swiss Federal Police, for more information see melani.admin.ch.

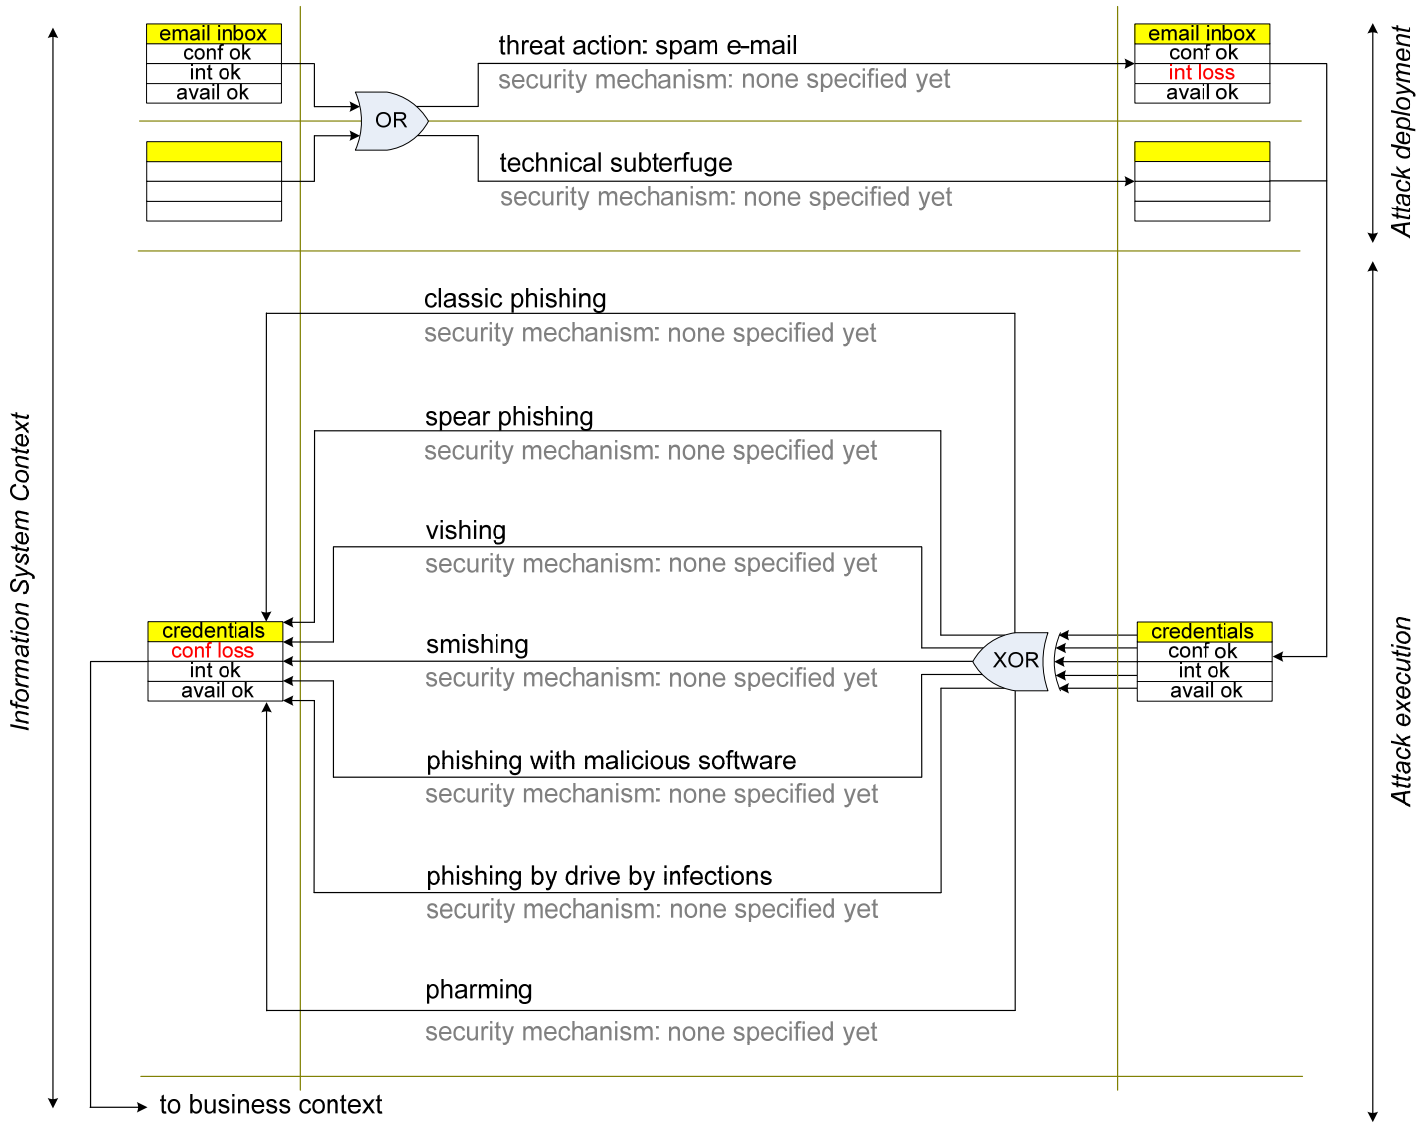

Figure 7.4: Phishing Scenarios in the IS Context.

The upper part of **Figure 7.4** shows that phishing attacks are deployed either by spam e-mails, technical subterfuge, or both. Yellow boxes on the left and right hand side of the drawing denote an asset in the IS context, in which the assets "user e-mail inbox" or the "user credentials" are of importance. Empty boxes indicate assets that have not been yet specified. The arrows connecting two assets symbolize a threat event which changes the state of the assets in terms of confidentiality, integrity and availability. Threat actions are displayed on the upper side of the arrows. The security mechanisms are displayed on the lower side of the arrows.

For our remitter, two of the seven scenarios are important: classic phishing and malware phishing.

CLASSIC PHISHING has been chosen because it is the most common. In a classic phishing scenario (shown in **Figure 7.5**), the **attacker acquires the user e-mail address** whose investigation is **out of scope** for this case study. Then a **spam e-mail** is sent to the user which is counteracted by the company **spam filter**. However, some spam e-mails may pass the filter and are received by the user. On the one hand, if the user chooses to act upon the spam e-mail, then he or she responds with his/her Internet banking credentials (**user response**). On the other hand, if the company recognizes the ongoing phishing attack then it attempts to shut down the Internet site of the attacker (**company response**) to prevent him or her from receiving more user responses. A comprehensive description of this scenario including its IS context is found in **Appendix F**.

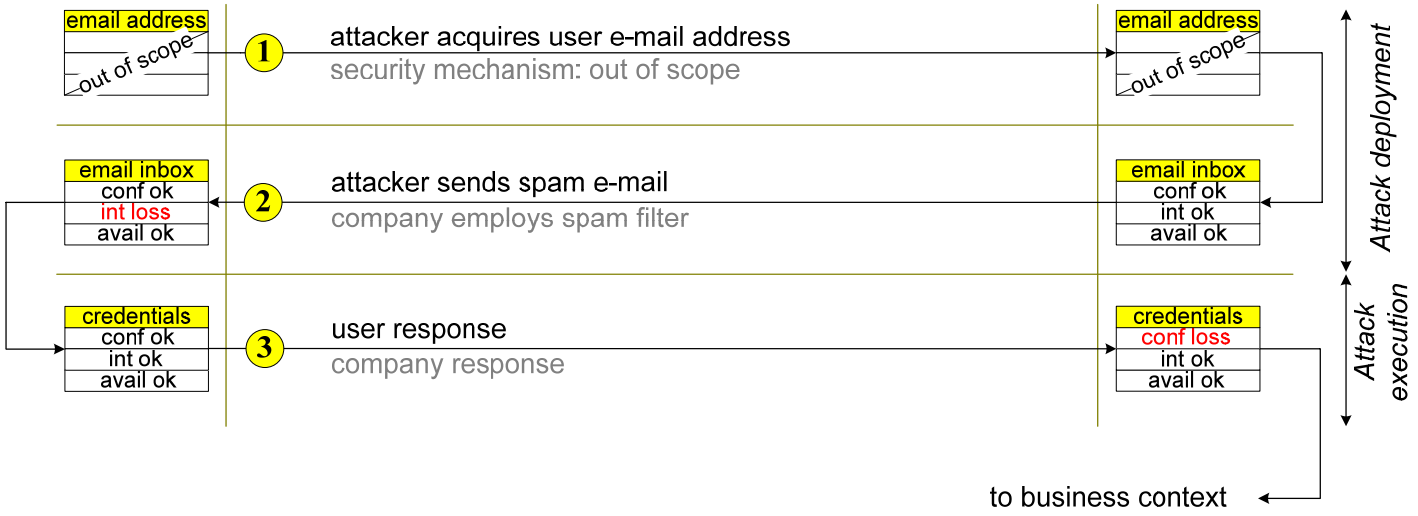

Figure 7.5: Classic Phishing Scenario.

PHISHING WITH MALWARE DOWNLOAD has been chosen because it is employed for industrial espionage (see Cachin et al [15], the European Union [14] and GAO [146]). In a spear phishing scenario (shown **Figure 7.6**), the attacker **acquires the e-mail address** and **sends a spam e-mail** to the user, which is counteracted by the company **spam filter**. Upon receiving the e-mail, the user **downloads the malicious software** (malware) which is counteracted by **anti-virus** software. Upon successful download, the malware captures, e.g., user identities (IDs), passwords and screen-shots and **communicates them back to the attacker**. This outgoing communication is counteracted by the company **firewall**. A comprehensive description of this scenario including its IS context is found in **Appendix G**.

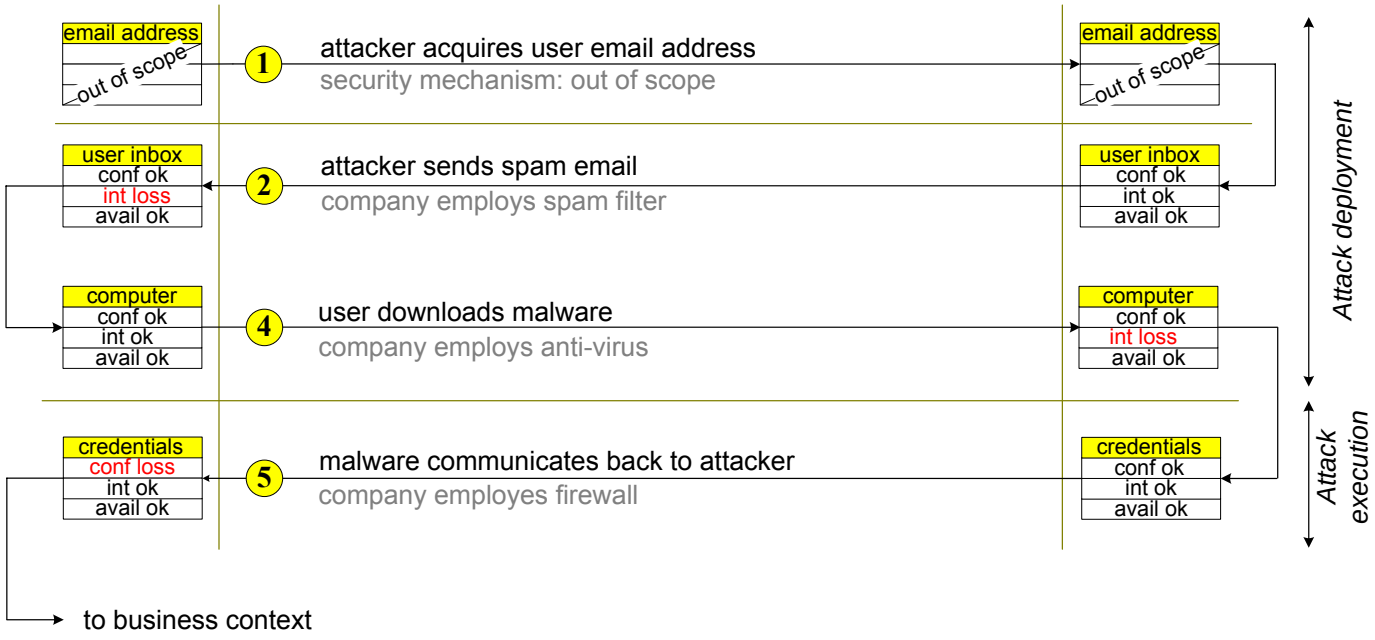

Figure 7.6: Phishing with Malicious Software (Malware).

7.2.5 Potential Future Developments

Clamping victims into executing phishing attacks is laborious. From an attacker point of view, the more interaction an attack requires with a victim the fewer users can be victimized with the same resources. Therefore, an increasing degree of technical subterfuge can be hypothesized for future phishing attacks as this diminishes the interaction with the victim.

In the scenario chart in **Figure 7.3**, the above argument signifies future attacks to shift to the right to emphasize "drive-by" infections and "phishing with malware download". And indeed, although technical subterfuge requires highly specialized programming skills, Emigh [132] has noted that it is on the rise. Moreover, the adoption of technical subterfuge is favored by the fact that:

- it is not possible to promptly issue and install security patches all over the country and
- distribution channels for deploying malware such as YouTube and Facebook are readily available.

Accordingly, missing security mechanisms such as automated security patches or anti-virus software make it attractive for organized Internet crime to deploy and execute

phishing attacks. In fact, Moll [147] notes the disconcerting emergence of spyware. In the first quarter of 2006, scan data reveals that *87* percent of consumer PCs in the U.S. are infected with an average of *34* pieces of spyware. This tendency has been confirmed by representatives of MELANI [148] in October 2007.

Finally, the assumption that the quality of spam e-mails will be improved for future attacks signifies phishing attacks to shift upwards in the scenario chart in **Figure 7.3**. Overall this yields a worst case scenario where masses of victims fall for attacks with little or no interaction from their side.

7.2.6 Cash out Scenarios in a Business Context (Financial Consequence)

Financial figures on the consequence of classic phishing attacks at our remitter as well as scenario charts in the business context have been distorted or omitted upon request. They are available upon written authorization by the remitter of the case study. Accordingly, the following statements on financial losses suffered from phishing attacks are fictitious. They comprise:

- refunds paid to customers which have been victimized by phishing attacks. According to Menotti and Walder [149] of the electronic banking department of the remitter, the refunds to customers represent only a small portion of the overall transaction volume
- time invested for investigating fraudulent cases, checking short lists of customer transactions, setting up and task forces to enhance internal control processes.
- efforts invested in training the internal work force or to deploy security awareness among customers.
- projects to counteract phishing attacks
- co-ordination efforts among institutions in the financial industry or MELANI, the Reporting and Analysis Centre for Information Assurance of the Swiss Federal Police.

7.3 Function Module: Frequencies and Probabilities

The probability is determined by the density functions showing employees answering spam e-mails and employees reporting the ongoing phishing attack to an internal authority. The form of these curves is hypothesized next.

7.3.1 Hypotheses

Hypothesis 7.1 "Lognormal User Response": In analogy to Aitchinson [150] where the behaviour of many systems in biology *arises from a theory of elementary errors* (actions) *combined by a multiplicative process*, the temporal behaviour of victims responding to phishing e-mails is interpreted as the result of many independent elementary actions. Therefore, a lognormal distributed density function can be used to describe it.

Such elementary actions could be the phone ringing (and thus distracting the employee from reading the phishing e-mail), a colleague dropping by, etc. Moreover, lognormal distributions are useful in describing high-variability phenomena, which span over a prolonged period of time.

Hypothesis 7.2 "Lognormal Company Response": Analogously, for the curve displaying the temporal response of the company a lognormal distribution is also assumed.

Hypothesis 7.3 "Characteristic User Response": The form of the curve "employees answering spam e-mails" is the same regardless whether the user is asked to surrender personal information or "click-to-install" malicious software.

If **Hypothesis 7.3** holds, then the insights gained from employees answering spam e-mails may be transferred to related actions such as users responding to internal circular e-mails and vice versa.

7.3.2 Data Collection

For **Hypothesis 7.1**, the raw data[69] represents an ethical phishing attempt commissioned in November 2006. In particular, it reflects the temporal behavior of *1159* employees responding to phishing e-mails with malware download.

The data is displayed in the below frequency curve and is called the *USER RESPONSE*, F_{1159}. On the *x-axis*, **Figure 7.7** shows the first 60 minutes after the phishing attack; the number of respondents per time unit is shown on the *y-axis*:

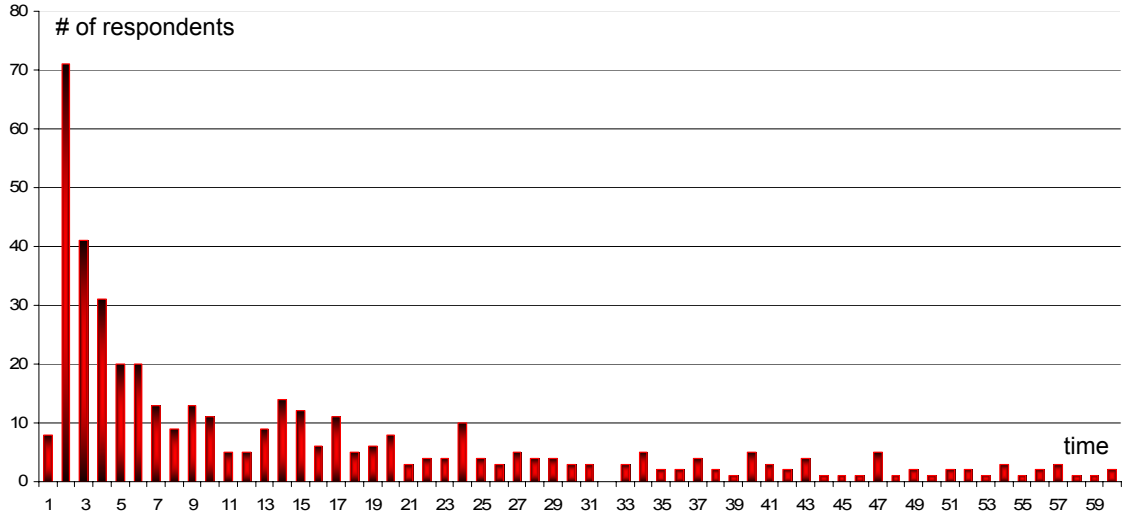

Figure 7.7: Temporal Behavior of Victims (F_{1159}).

The following information is supplied for the above data set:

Threat	Phishing with Malware Download[70]
Attack Deployment	Via e-mail on a working day in November 2006; from 9:17 to 10:15
Attack Execution	E-mail contains a link to download malware (Web Trojan)
Targets	*1159* victims (employees)

[69] The data has been made available courtesy of scip, a consulting company in the area of information security, www.scip.ch.

[70] The same data set will also be used to represent the user response curve to classic phishing.

Overall Duration	*274* minutes (from which the first *60* minutes are displayed)
Respondents	• *562* respondents over a period of *274* minutes, around *48%* • *426* respondents over a period of *60* minutes, around *37%*

For **Hypothesis 7.2**, the raw data[71] reflects the point in time when MELANI received confirmation from a hosting provider that a suspicious Internet site had been brought down upon previously requesting it. The data consists of *34* data points, which have been gathered from December 2006 through September 2007. This frequency curve is called the *COMPANY RESPONSE*, F_{34} **(Figure 7.8)**.

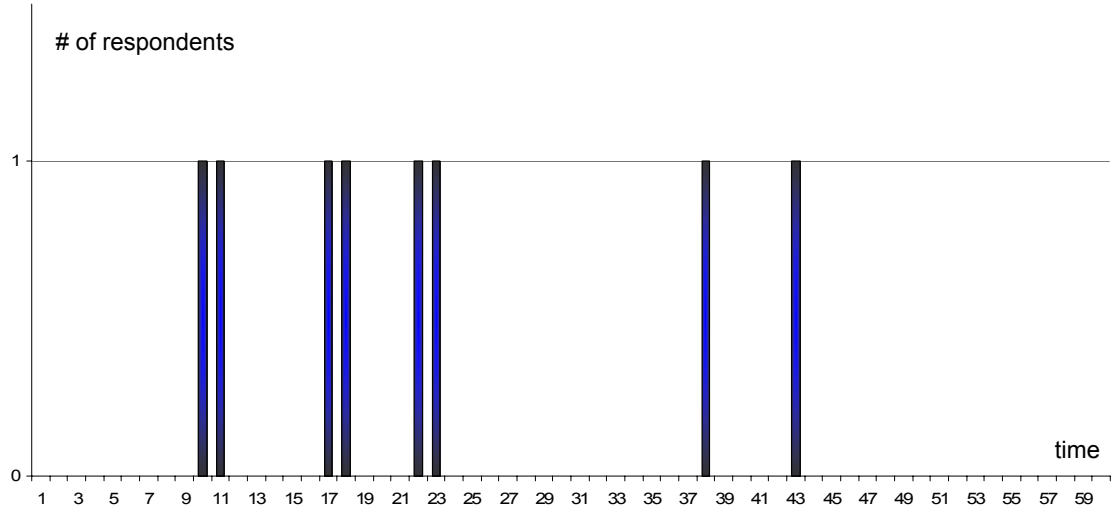

Figure 7.8: Temporal Response of Hosting Providers (F_{34}).

The following information is supplied for the above data set:

Security Mechanism	Shut down fraudulent Internet site
Attack Deployment	No impairment of attack deployment

[71] The data has been made available by MELANI of the Swiss Federal Police, www.melani.admin.ch.

Attack Execution	Impairment of attack execution by obstructing user credentials from reaching the attacker Shutting down the attacker's site was executed on working days at various times
Targets	*34* fraudulent Internet sites of *34* sites in total
Overall Duration	*6932* minutes (from which the first *60* minutes are displayed)
Respondents (ISP[72])	• *32* shut downs within a period of *6932* minutes, around *94%* • *8* shut downs within a period of *60* minutes, around *25%*.

For **Hypothesis 7.3**, the raw data[73] reflects the temporal response of users to a company-internal circular e-mail. It was gathered in December 2007. This frequency curve is called the *USER RESPONSE*, F_{20667}. **Figure 7.9** shows the number of victims per time unit answering the company-internal circular e-mail.

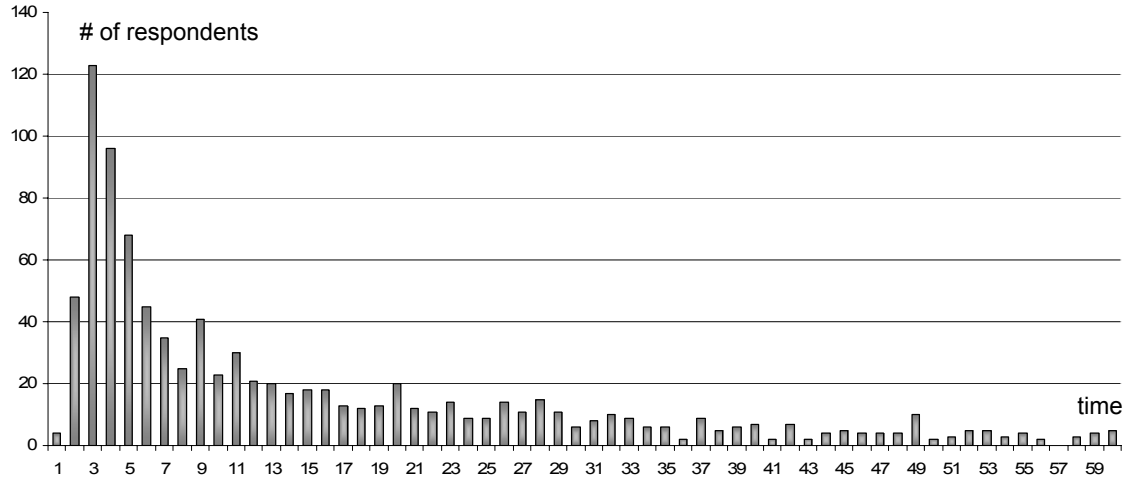

Figure 7.9: Response to Internal E-mail (F_{20667}).

The following information is supplied for the above data set:

[72] Internet Service Provider

[73] The data has been made available courtesy of Credit Suisse, a global bank, www.credit-suisse.com.

"Threat"	Response to company-internal circular e-mail
Deployment	Via e-mail on a working day in November 2007; from 14:44 to 14:55
Execution	E-mail contains a link, which takes the user to a company intranet site
Targets	*20667* employees
Overall Duration	ca. *5* days (from which the first *60* minutes are displayed)
Respondents	• *1426* respondents over a period of *5* days, around *6.8%* • *951* respondents over a period of *60* minutes, around *4.6%*.

7.3.3 Data Biases

As with any real world experiment, there are biases, which offset the three data sets:

- the temporal behavior of victims of an ethical phishing attempt (user response, F_{1159})
- the reaction time of MELANI and the hosting provider (company response, F_{34})
- the reaction time of users to circular, company-internal e-mails (user response, F_{20667}).

Data Biases for the Temporal Behavior of Victims (F_{1159})

Bias	Description	Comment
A	The data point at *t=32* has the value *zero* (**Figure 7.7**). As the data is to be approximated by a lognormal distribution, this value needs to be treated (logarithm of *zero* returns an infinite result).	According to Aitchinson and Brown [150], to treat *zero* observations in lognormal distributions a positive constant can be added to all sample values or the zero values are replaced with a constant. In this work it is chosen to add *1* to the *zero* value at *t=32* because it signifies the least overall change in the data.

Table 7.3: Data Biases for the Temporal Behavior of Victims.

Data Biases for the Reaction Time of Hosting Provider (F_{34})

Bias	Description	Comment
❶	The data reflects the points in time when hosting providers are requested to bring down a suspicious Internet site by MELANI. It covers only a small proportion of the communication chain. In particular, it does not include users recognizing a phishing attack and notifying it to company-internal authorities until this information reaches MELANI. Refer to **Figure 7.10** for the complete notification chain.	The missing data is not investigated as it would not add significant value to verifying the overall applicability of the model.
❷	Once the Internet site of the attacker is disabled, it can still be reached for a small amount of time because the access data is still present in the browser caches of users and ISP.	The missing data is not investigated (reason: see bias 1).
❸	There are only *32* data points representing *34* requests to shut down Internet sites. This number is too low to be statistically relevant.	The missing data is constructed using a lognormal distribution.
❹	F_{34} also contains *7* requests to shut down websites advertising jobs for financial agents. Shutting down websites of attacker looking for money mules is usually not time critical.	The superfluous data is not extracted (reason: see bias 1).
❺	F_{34} reflects the points in time when the hosting provider has been requested by MELANI to shut down a fraudulent site. It is possible that the hosting provider does not notify MELANI immediately after shutting it down.	The missing data is not investigated (reason: see bias 1).

Table 7.4: Data Biases for the Reaction Time of Hosting Provider.

Data Biases for Circular E-mails (F_{20667})

Bias	Description	Comment
I	The data point at *t=57* has the value *zero* (**Figure 7.9**). As the data is to be approximated by a log-normal distribution, this value needs to be treated (logarithm of *zero* returns an infinite result).	The missing data point is treated as in comment 1 of 7.4.3.1.

Table 7.5: Data Biases for Circular E-mails.

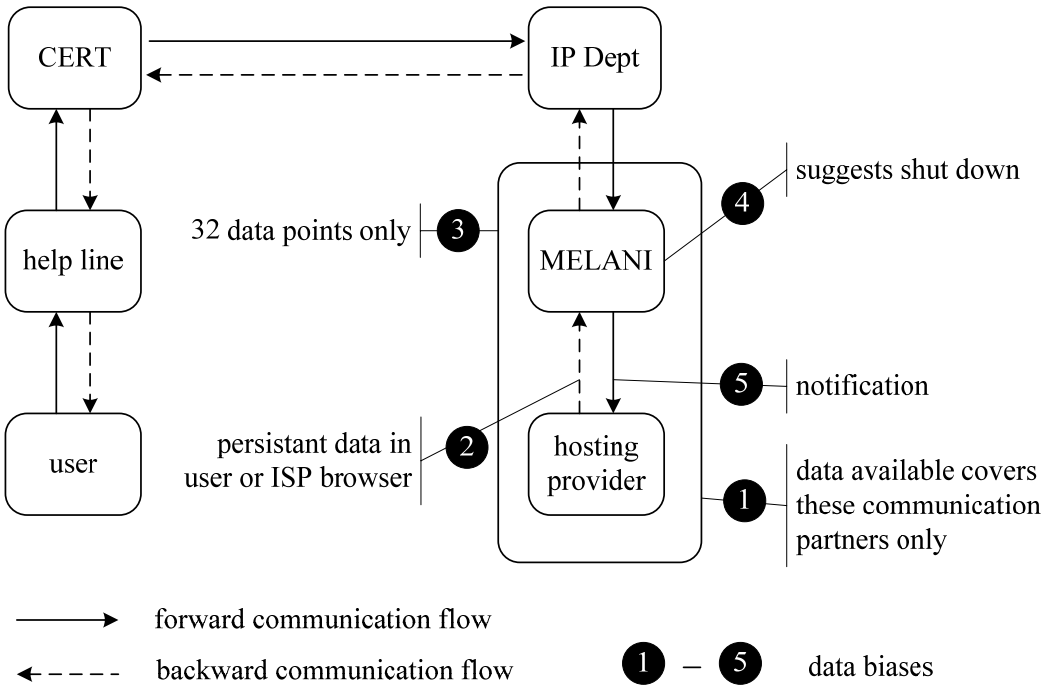

Figure 7.10: Data Biases in Notification Flow.

Figure 7.10 shows a complete notification flow, evidencing the biases for F_{34} (grey area) which have been validated by security experts (Salvati [149]). The flow is triggered by users recognizing the ongoing phishing attack and reporting it to the company help line, which reports it to the company computer emergency response team (CERT), which notifies the company department trusted with the registration and deregistration of IP addresses (IP Dept), which notifies MELANI, which notifies the hosting provider.

7.3.4 Curve Fitting: Lognormal Distribution

As argued before in **Hypothesis 7.1** and **7.2**, the temporal response of users and companies to phishing attacks follows a lognormal distribution. This two-parameter curve is unimodal, skewed to the left and it is defined by:

$$\Lambda(\mu,\sigma\ ) = f_{u,\sigma}\ (t) = \begin{cases} \dfrac{1}{t\sigma\sqrt{2\pi}} \cdot e^{-\frac{(\ln(t)-\mu)^2}{2\sigma^2}} & (t > 0) \\ 0 & (t \le 0) \end{cases}, \tag{7.1}$$

where $\Lambda(.)$ is the lognormal probability density
μ is the mean
σ is the standard deviation
t is time.

To fit two continuous lognormal curves, C_{1159} and C_{20667}, based on the related frequency curves, the statistical calculator by Wessa [151] has been used and a Maximum Likelihood fit for the first *60* minutes of the available data was performed, see **Table 7.6**.

Curves	Parameter	Estimated Value	Standard Deviation, σ
C_{1159}	μ_{1159}	*1.38 ± 2·σ (95% interval)*	*0.13*
	σ_{1159}	*1.00± 2·σ (95% interval)*	*0.09*
	Parameter	**Estimated Value**	**Standard Deviation, σ**
C_{20667}	μ_{20667}	*2.20 ± 2·σ (95% interval)*	*0.13*
	σ_{20667}	*1.01 ± 2·σ (95% interval)*	*0.09*

Table 7.6: Curve Fitting for C_{1159} and C_{20667} (Maximum Likelihood).

The above results show a reasonably low standard deviation for the individual μ and σ-values. In particular, the coefficient of variation, σ / μ, ranges between *0.06* and *0.09* for both curves. Moreover, the QQ-plots depict a good overall fit (refer to **Appendix H**). It is noted that the standard deviation σ has the same value for C_{1159} and C_{20667}. This is attributed to the similar actions, which must be performed by the victim prior to answering either a phishing e-mail or a company-internal circular e-mail.

Given the low amount of data, for F_{34} it was not possible to produce a good fit based on Maximum Likelihood (refer to **Appendix J**). However, it is advocated to approximate F_{34} with the lognormal distributed C_{34} because:

- lognormal distributions express a multiplicative concatenation of individual probability densities of many small and independent actions [150]. For F_{34} it is argued that it equally represents a lineup of independent small actions:
  (1) Upon verifying that a suspicious website is indeed related to a phishing attack, MELANI needs to identify and contact the hosting provider
  (2) The contact details of the hosting provider are usually readily available; however, it is possible for the information to be out-of-date, which delays the subsequent actions
  (3) Once the correct contact information is available, the hosting provider is informed by MELANI about the phishing attack, usually via e-mail
  (4) This e-mail again starts a cycle of independent small actions for the hosting provider (reading the e-mail, verifying its contents, reading the phishing e-mail, verifying its contents, checking the suspicious website, possibly obtain a second opinion, shut down the attacker's website).
- If we agree that C_{34} follows a lognormal distribution then there is a need to assess μ_{34} and σ_{34}. Let $\sigma_{34} \approx 1.00$; the same as it was for σ_{1159} and σ_{20667} because the actions performed for C_{1159}, C_{20667} and C_{34} are similar.
- Finally, for C_{34}, a value μ_{34} is estimated. This is done by *GUESSING*, i.e. an optical fit for the first *60* minutes of F_{34} is performed and $\mu_{34} \approx 4.50$ is obtained (**Table 7.4**).

Curve	Parameter	Estimated Value	Rationale
C_{34}	μ_{34}	*4.50*	Graphical fit
	σ_{34}	*1.00*	Similar actions for C_{1159} or C_{20667}

Table 7.7: Curve Fitting for C_{34} (Educated Guess).

7.3.5 Calculation of Success Probability

There are three possibilities to approach the calculation of probability by applying *MATHEMATICAL ANALYSIS*, numeric integration and Monte Carlo techniques. Analytically

speaking, the black curve in **Figure 7.11,** C_{sim}, signifies an approximation to the converse convolution of C_{1159} with C_{34}. Integrating C_{sim} for $0 < t < \infty$ yields the desired probability of phishing attacks being successful while integrating C_{sim} for $-\infty < t \leq 0$ yields the probability of phishing attacks being unsuccessful.

Unfortunately, an analytical solution to the sum (or difference) of independent log-normal random variables (C_{1159} and C_{34}) is unavailable because the characteristic function[74] of lognormal random variables is not known. Consequently, no closed-form expression is found. Approximations to the analytical solutions; however, do exist, e.g., by Schwartz and Yeh [152]. According to Joshua Lam and Le-Ngoc [153] the accuracy of the approximations relies *highly on the method, the region of the resulting distribution and the individual lognormal parameters being examined*.

Next, *NUMERIC INTEGRATION* simulation is difficult due to the tails of lognormal probability distribution functions. In fact, Beaulieu et al [154] attribute this to *slowly decaying envelopes*, which are very demanding on computational power.

Finally, we turn to *MONTE CARLO* techniques and start by simulating $\boldsymbol{C_{1159}}$. This red curve in **Figure 7.11** has been generated by *10,000* lognormal distributed draws in MATLAB[75] and represents the curve *"user response"* with $\mu_{1159} = 1.38$ and $\sigma_{1159} = 1.00$. Aside from the red curve, a blue curve representing the *"company response"*, $\boldsymbol{C_{34}}$, is generated with $\mu_{34} = 4.50$ and $\sigma_{34} = 1.00$. For the red and the blue curves, the *x-axis* signifies the time t of ongoing phishing attacks (the attacks start at $t = 0$) while the *y-axis* signifies the number of user responses (red curve) and the number of company responses (blue curve) at the time t.

In general, phishing attacks are successful if the user answers a phishing e-mail before the Internet site of the attacker is shut down. Accordingly, $\boldsymbol{C_{sim}}$ (black curve) has been generated, by subtracting all *10,000* random draws for C_{1159} from the *10,000* draws for C_{34}. If the time difference for the individual draw is smaller than zero, i.e. $t_{diff} < 0$ then the user response *does not* reach the attacker (the attack was not successful). Conversely, if $t_{diff} \geq 0$, then the attack is successful. For C_{sim}, the *x-axis* is interpreted as t_{diff} while the

[74] In probability theory, characteristic functions completely define the probability distributions of any random variable.

[75] MATLAB offers an interactive environment for performing computationally intensive scientific calculations http://www.mathworks.com/products/matlab/.

y-axis reflects the number of attacks that were unsuccessful (in case of $t_{diff} < 0$) and successful (in case of $t_{diff} \geq 0$).

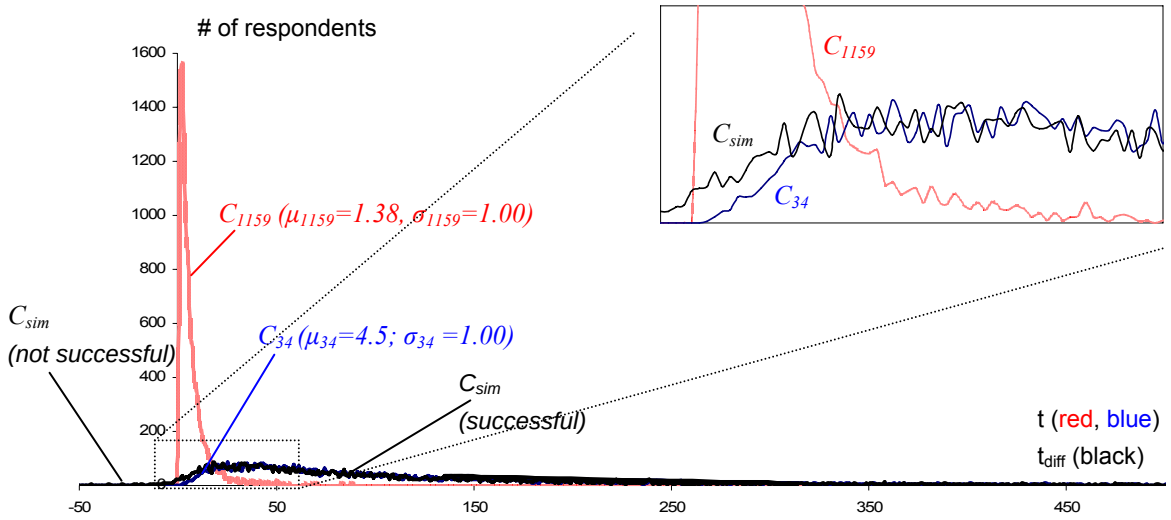

Figure 7.11: Subtracting C_{1159} from C_{34} yields C_{sim}.

To obtain the probability curve shown in **Figure 7.12**, the three curves (C_{1159}, C_{34}, C_{sim}) were normalized by dividing each of their occurrences along the time axis by the overall number of responses for the respective curve (*10,000* in this case). In order to calculate the probability the area under the black curve in **Figure 7.11** left and right of *0* of the time axis is of interest. The probability has been computed by summing up the number of occurrences, which were ascertained for each time difference. The below graph shows the probability that the random variable pair "user response" and "company response" is successful in terms of the threat overcoming the security mechanism displayed by the time difference. On the *x-axis* of the below graph, the time difference t_{diff} is shown. On the *y-axis*, the numbers of pairs of user and company responses (normalized to *1*) have been plotted, yielding the time difference t_{diff}.

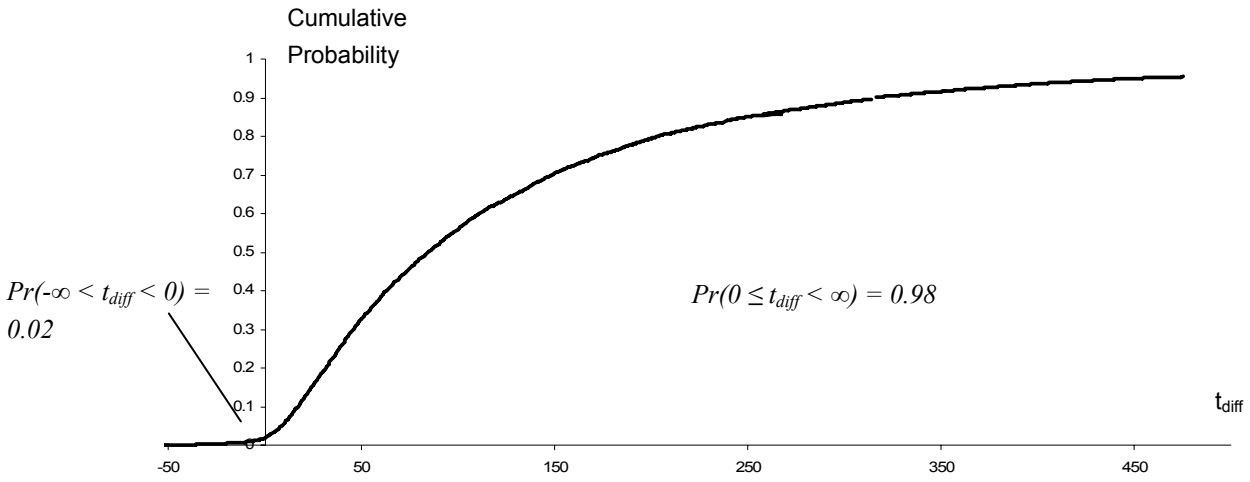

Figure 7.12: Probability of Successful Phishing Attacks (for curves C_{1159} and C_{34}).

The security mechanism related to the company response displays a *PERFECT VULNERABILITY* towards the user response as $Pr(t_{diff} \leq 0) = 0.02$. The probabilities resulting from integrating the convolution of C_{1159} with C_{240}, C_{20667} with C_{34} and C_{20667} with C_{240} are reported in **Appendix J**. These results will be needed in the Decision Module.

7.3.6 Discussion of Results (Function Module)

With respect to **Hypothesis 7.1** (users responding to phishing e-mails), F_{1159} was fitted with the lognormal C_{1159}. For its parameters μ_{1159} *and* σ_{1159}, the coefficient of variation varies between acceptable limits (*0.06* and *0.09*).

With respect to **Hypothesis 7.2** (company responding to phishing e-mails), F_{34} could not be positively built in into the lognormal C_{34} due to the low amount of data points. Nevertheless, a lognormal distribution was applied to describe F_{34} as it is conceived – like F_{1159} – as a multiplicative concatenation of many small and independent actions.

With respect to **Hypothesis 7.3** (temporal response of employees to a company-internal, circular e-mail) F_{20667} was fitted into the lognormal probability density C_{20667}. The virtually identical σ-values for C_{1159} and C_{20667} ($\sigma_{1159} = \sigma_{20667} = 1.00$) indicate that the

skew of a lognormal curve reflecting the temporal behaviour of employees responding to (phishing) e-mails is not affected by the:

- quality and type of request in the fraudulent e-mail,
- number of employees and size of the company,
- industry where the employees work in,
- security awareness of the employees.

Accordingly, only μ remains to describe the influence of the above points on probability densities. As the contents of the e-mails which yielded the user response curves C_{1159} and C_{20667} were very different (one e-mail requested the user to install a piece of malicious software while another requested the nomination of candidates for an internal contest) it should be possible to designate probability densities to phishing attacks from sources other than the attacks themselves, thus confirming **Hypothesis 7.3**. This is an important result for practice, as this allows the use of more readily available data.

Colloquially speaking, companies display almost *PERFECT VULNERABILITIES* in their defence system as the success probability of phishing attacks are:

- *0.202* within the first hour,
- *0.424* within the first two hours and
- *0.687* within the first four hours.

In more precise terms, the above result expresses the following:

- C_{1159} represents users responding to one instance of a phishing attack. In the case study it has been implicitly assumed that other instances of phishing attacks will cause the users to respond in the same or at least in an analogous way.
- C_{34} denotes various hosting providers shutting down a malicious Internet site on behalf of MELANI. The data represents the reaction to *34* different attacks and was called the company response.

As it has been assumed that the user response – *ceteris paribus* – will look the same for every phishing attack, C_{1159} has been convoluted with C_{34}. A discrete converse convolution "compares" every single point of C_{1159} with every single point of C_{34} and determines whether the attack succeeds or not in the following way:

- For one phishing attack, the passwords of all users will have been disclosed to the fraudulent Internet site of an attacker up to the point where the hosting provider shuts it down. All other passwords of user responding will not reach the attacker. A percentage of disclosed passwords is obtained.
- The discrete converse convolution compares all lognormal distributed phishing attacks with all lognormal distributed site shutdowns and a percentage of disclosed passwords over all phishing attacks is obtained.
- **Figure 7.12** shows the percentage of password disclosure dependent on the time difference between passwords being sent to the attacker and the point in time where the fraudulent Internet site of the attacker is shut down.

Remark: The approach to calculating the success probabilities where employees report the ongoing phishing attack to a company-internal authority rather than to MELANI is shown in **Appendix L**.

Remark: The analysis of additional sets of data[76] related to three phishing attacks on small and medium-sized companies indicates that the time of sending the e-mail has an impact on the general shape of the user response curve in the sense of a shift to the right along the *x-axis*, if sent at noon. Surprisingly, the time employed to discover the attack does not primarily seem to depend on the number of targeted users. In fact, by analyzing the data available for the three companies (with a number of respondents well below *130*), the time needed to discover the ongoing attack ranged between *45* and *55* minutes. Based on the above, it can be hypothesized that phishing attacks ought to be performed when the majority of the employee force is working. This might be an explanation in support of Ramzan and Wüests [144] finding that the attacks usually take place during working days rather than on weekends.

[76] The data has been made avaialable courtesy of Infoguard, a consulting company in the area of information security, www.infoguard.com.

7.4 Influence Module: Influence of the Context on Security Mechanisms

This module aims at evidencing the varying influence of different types of security awareness training on the frequency curves "user response" and "company response". Three hypotheses regarding these influences are formulated next.

7.4.1 Hypotheses

Hypothesis 7.4 "Training on User Response": Applying appropriate security training the number of respondents to phishing attacks diminishes. In addition, they tend to respond later, see sketch in **Figure 7.13**.

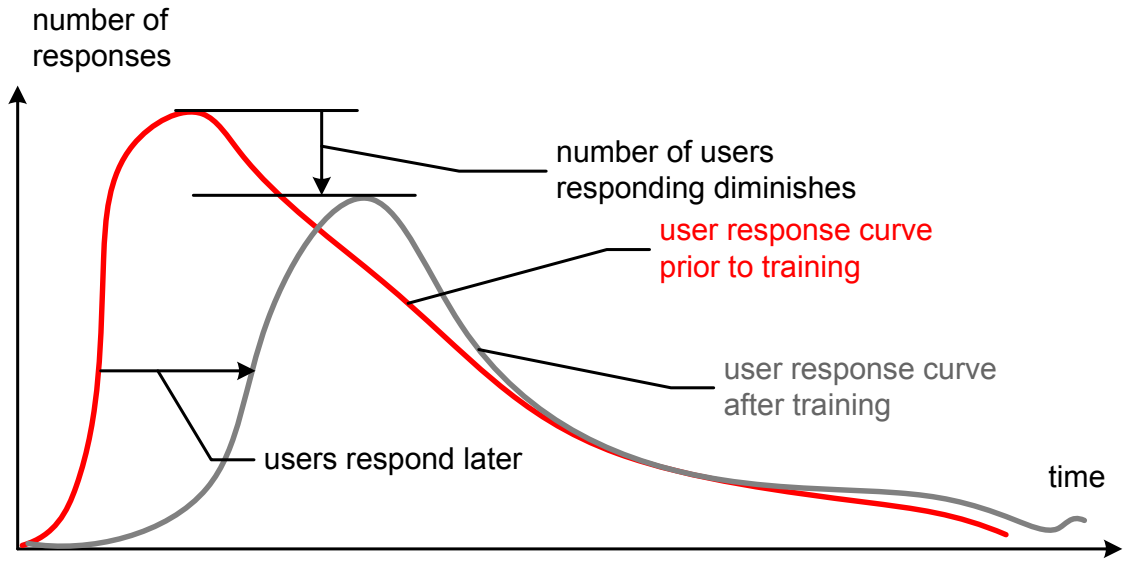

Figure 7.13: Less Users Respond Later to a Phishing Attack.

Hypothesis 7.5 "Training on Company Response": Applying the appropriate security awareness training causes the number of users reporting ongoing phishing attacks to their help line to rise. In addition, they are reported earlier, see sketch in **Figure 7.14**.

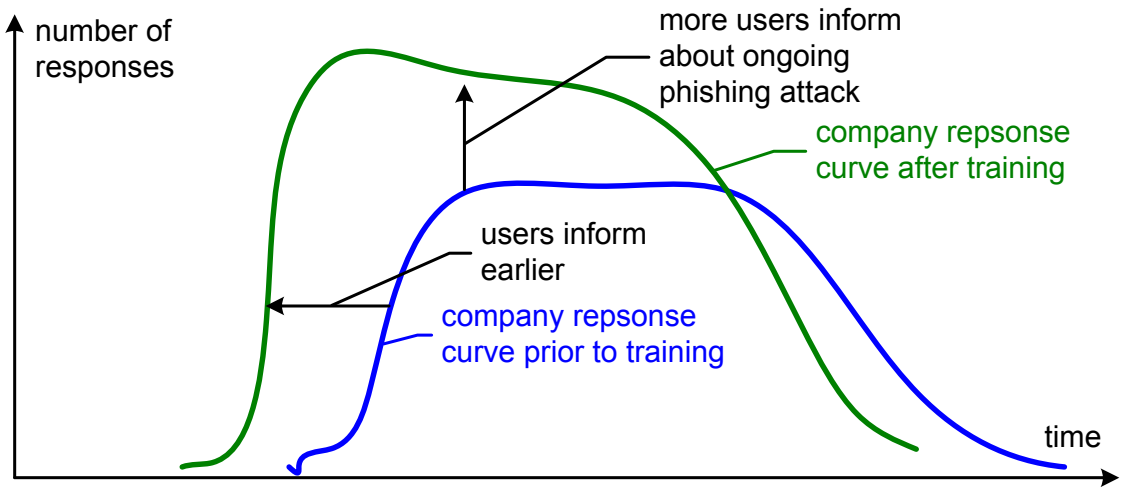

Figure 7.14: More Users Report a Phishing Attack Earlier.

Hypothesis 7.6 "Optimal Training": The contents conveyed in the security awareness training influences the "user response", the "company response" or both. By selecting an appropriate message it is possible to favorably influence the associated curves, preferably both.

7.4.2 Data Collection and Evaluation Methodology

The security awareness training was originally intended for a body of *4800* (*6400*) employees, which were to be split into three (four) groups of *1600* individuals each:

- the first group was intended for training in *RECOGNIZING* phishing attacks
- the second group was intended for training on how to *PREVENT* becoming a potential victim
- the third group was intended for training on how to *REACT* to an ongoing phishing attack
- (the fourth group was taken as a control group and was to receive no training)

The intention was to measure the influence of security awareness training by means of two company-commissioned phishing attacks. One phishing attack was to be performed prior to the training and the other was to be performed after the training. The data was to be evaluated in terms of the Influence Module.

Due to a shortage in funding the scope of the case study for the Influence Module was restricted to verifying the applicability of Rough Sets Theory. Therefore, to reflect the security awareness of users, three data sets have been generated each reflecting the security awareness of one group. For the generation of data the following was assumed:

- training the users of group 1 in recognizing phishing attacks influences their user response curve
- training the users of group 2 in preventing phishing attacks neither influences the user response curve nor the company response curve
- training the users of group 3 in reacting to phishing attacks influences the company response curve

Finally, the security awareness of each group was depicted into three Rough Sets data tables schematically shown in **Table 7.8** where the orange colour signifies that the respective group has received training as outlined above:

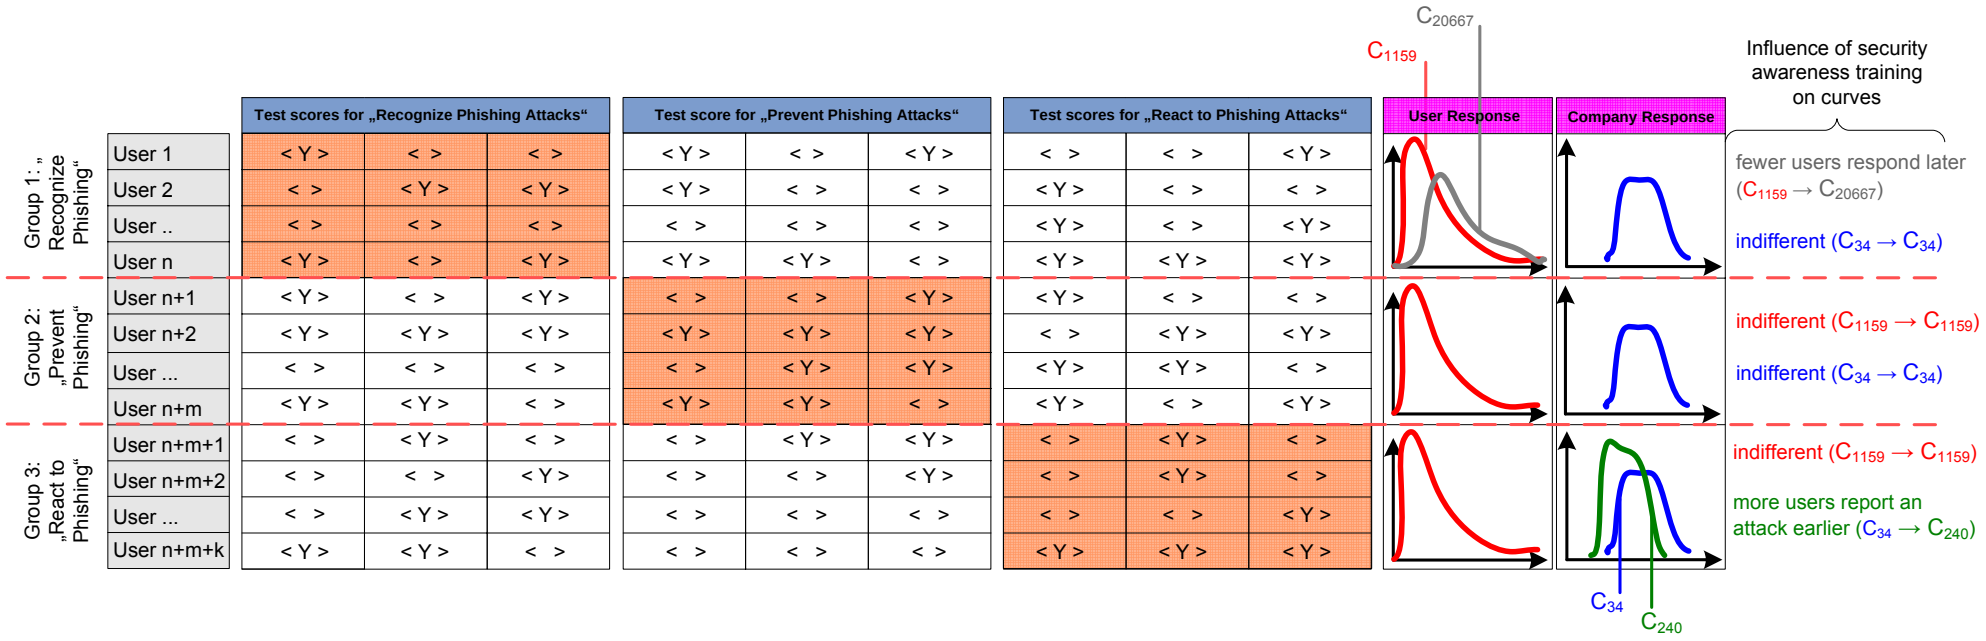

Table 7.8: Influence of Awareness Training on User and Company Responses.

To describe the *USER RESPONSE* of users who have received no security awareness training in recognizing phishing attacks C_{1159} was adopted. Users who received this training were represented by C_{20667}. The training is very successful because the first *25%* of user responses (lower quartile[77]) are only reached after *7.77* minutes (for C_{20667}) while the equivalent figure for C_{1159} is *2.03* minutes. Furthermore, *75%* of user responses (upper quartile) are reached only after *17.64* minutes for C_{20667} as opposed to *4.62* minutes for C_{1159}.

To describe the *COMPANY RESPONSE* of security unaware users C_{34} was employed. For users who have been trained in reacting to phishing attacks, C_{34} was replaced by a fictitious C_{240} ($\mu_{240} = 3.70$, $\sigma_{240} = 1.00$). Switching from C_{34} to C_{240} indicates very successful security awareness training as the first *25%* of the company responses are reached already after *14.88* minutes (for C_{240}) while the equivalent figure for C_{34} is only at *46.06* minutes. Furthermore, *75%* of the company responses take place after *79.04* minutes for C_{240} as opposed to *175.91* minutes for C_{34}.

In order to obtain a symbolic representation of C_{1159}, C_{20667}, C_{34}, and C_{240} suited for display in a data table, the respective line sections[78] on the x-axis have been coded as follows:

line section on time axis (x-axis)	discretization (symbolic)
zero to mode	D++
mode to lower quartile	D+
lower quartile to median	C
median to mean	D-
mean to upper quartile	D--

, where

[77] For mathematical description see **Appendix K**.

[78] Line sections represent segments on the *x-axis* of a lognormal curve, delineated by well known points such as the “mode”, “median”, “mean”, as well as “lower quartile” and “upper quartile”, for mathematical description see **Appendix K**.

point on time axis	for C_1159 at (minutes)	for C_20667 at (minutes)	for C_32 at (minutes)	for C_240 at (minutes)
zero	0	0	0	0
mode	1.46	3.32	33.16	14.88
lower quartile	2.03	4.62	46.06	20.7
median	3.97	9.03	90.02	40.45
mean	6.55	14.88	148.41	66.69
upper quartile	7.77	17.64	175.91	79.04

Table 7.9: Discretization of the Lognormal User and Company Response Curves.

For example, for C_{1159}, D++ reflects users who employed between *0* minutes (*zero* point) and 3.32 minutes (modal point) to answer a phishing e-mail. The number of users classified by D++ corresponds to *16%* of the overall number of respondents.

Taking all of the above into consideration, data was generated for three Rough Sets data tables, which are displayed next in **Tables 7.10** through **7.12**. Moreover, the graphical depiction of the individual curves has been replaced by a symbolic representation as shown in **Table 7.9**.

In the above table, the column to the extreme left shows their categorization into three groups of users according to the type of training received (recognition, prevention or reaction to phishing attacks). The columns under the blue headings exemplify the security awareness of the users where < Y > indicates that a user is aware of a certain security topic while < > shows that the user is not. The orange areas signify the topics where training has been provided. For example, the orange area in the upper left indicates those security topics that have been provided by the training “recognizing phishing attacks”. The two columns on the right of **Table 7.9** contain a graphical depiction of C_{1159}, C_{20667}, C_{34} and the fictitious C_{240}. The grey columns signify that no change in the user or company response has taken place while the white columns indicate change. All other data in the table was generated by a uniform probability distribution.

Rough Sets data table for Group 1 "Recognizing Phishing Attacks":

	CONDITIONAL ATTRIBUTES															DECISION ATTRIBUTES	
	Recognize					Prevent					React						
User	What is phishing?	What are money mules?	Groups (depl. techniques)	Phones (depl. techniques)	Malware download	Read privacy policy	Install spam filter	Regularly update software	Use restricted profile	Anti malware software	Do	Do not answer unknown emails	Do not surrender credentials via phone	Do report unexpected errors	Do report job offers	user response (C_20667)	company response (C_32)
1		Y				Y	Y	Y	Y		Y	Y	Y	Y		D--	D--
2			Y		Y	Y		Y	Y	Y	Y			Y	Y	D++	D++
3	Y			Y				Y			Y	Y		Y		D-	D-
4					Y	Y	Y		Y		Y	Y	Y	Y	Y	C	C
5					Y	Y			Y	Y	Y	Y	Y			C	C
6			Y			Y				Y		Y				D+	D+
7					Y	Y		Y	Y		Y	Y	Y	Y	Y	C	C
8			Y					Y			Y	Y		Y	Y	D+	D+
9		Y					Y	Y		Y	Y	Y			Y	D--	D--
10					Y				Y	Y		Y	Y		Y	C	C
11	Y			Y			Y	Y	Y			Y	Y	Y	Y	D-	D-
12					Y			Y	Y		Y	Y		Y		C	C
13					Y		Y	Y	Y		Y	Y				C	C
14			Y			Y		Y				Y		Y	Y	D+	D+
15			Y		Y	Y	Y	Y	Y	Y						C	C
16			Y		Y			Y	Y				Y			D++	D++
17	Y			Y		Y	Y	Y		Y			Y	Y		D-	D-
18					Y				Y	Y			Y	Y	Y	C	C
19	Y			Y			Y	Y		Y					Y	D-	D-
20					Y						Y	Y	Y	Y	Y	C	C
21					Y			Y		Y		Y	Y			C	C
22			Y			Y		Y	Y				Y	Y		D+	D+
23					Y	Y			Y			Y	Y			C	C
24			Y		Y		Y	Y				Y	Y			D++	D++
25	Y			Y			Y	Y	Y	Y		Y				D-	D-
26					Y			Y	Y	Y		Y	Y			C	C
27					Y		Y			Y	Y	Y	Y			C	C
28			Y					Y		Y	Y	Y			Y	D+	D+
29					Y		Y		Y		Y	Y		Y		C	C
30			Y							Y		Y	Y		Y	D+	D+
31					Y						Y		Y	Y	Y	C	C
32			Y		Y	Y	Y						Y	Y	Y	D++	D++

Table 7.10: Influence of Training "Recognizing" on User and Company Response.

Rough Sets data table for Group 2 "Preventing Phishing Attacks":

	CONDITIONAL ATTRIBUTES															DECISION ATTRIBUTES	
	Recognize					Prevent					React						
	What is phishing?	What are money mules?	Groups (depl. techniques)	Phones (depl. techniques)	Malware download	Read privacy policy	Install spam filter	Regular software updates	Restricted user profile	Anti malware software	Answer 1	Do not answer unknown emails	Do not surrender credential via phone	Do report unexpected error	Do report job offers	user response (C_1159)	company response (C_32)
User																	
101		Y	Y		Y	Y				Y		Y		Y		D--	D--
102	Y				Y	Y		Y	Y	Y	Y		Y	Y		D++	D++
103	Y		Y	Y	Y		Y		Y		Y					D-	D-
104	Y	Y		Y	Y	Y								Y	Y	C	C
105		Y	Y	Y	Y	Y	Y	Y	Y			Y	Y		Y	C	C
106	Y				Y	Y		Y		Y	Y	Y	Y			D+	D+
107	Y	Y	Y	Y		Y					Y	Y				C	C
108			Y	Y		Y			Y		Y	Y	Y		Y	D+	D+
109	Y	Y			Y			Y		Y	Y	Y	Y			D--	D--
110	Y		Y	Y						Y		Y		Y		C	C
111	Y	Y	Y	Y			Y	Y	Y		Y	Y			Y	D-	D-
112			Y			Y	Y	Y	Y		Y	Y	Y	Y	Y	C	C
113				Y	Y			Y		Y				Y		C	C
114	Y		Y	Y	Y	Y	Y	Y			Y	Y	Y	Y	Y	D+	D+
115	Y			Y	Y						Y			Y		C	C
116	Y	Y	Y	Y					Y			Y	Y	Y		D++	D++
117	Y			Y			Y		Y		Y	Y				D-	D-
118			Y				Y	Y	Y	Y	Y	Y		Y		C	C
119	Y		Y		Y			Y		Y		Y	Y	Y		D-	D-
120	Y		Y	Y		Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	C	C
121				Y	Y	Y		Y		Y		Y	Y	Y		C	C
122			Y	Y				Y	Y			Y			Y	D+	D+
123		Y				Y	Y	Y	Y		Y					C	C
124			Y				Y	Y	Y	Y					Y	D++	D++
125		Y			Y	Y	Y			Y			Y			D-	D-
126			Y	Y	Y		Y		Y	Y		Y	Y		Y	C	C
127	Y				Y		Y	Y			Y				Y	C	C
128					Y		Y	Y		Y					Y	D+	D+
129	Y		Y	Y	Y		Y		Y	Y				Y	Y	C	C
130	Y			Y		Y	Y	Y		Y			Y		Y	D+	D+
131			Y		Y		Y	Y						Y	Y	C	C
132	Y	Y			Y	Y					Y	Y	Y		Y	D++	D++

Table 7.11: Influence of Training "Preventing" on User and Company Response.

Rough Sets data table for Group 3 "Reacting to Phishing Attacks":

	CONDITIONAL ATTRIBUTES															DECISION ATTRIBUTES	
	Recognize					Prevent					React						
	What is phishing?	What are money mules?	Groups (depl. techniques)	Phones (depl. techniques)	Malware download	Read privacy policy	Install spam filter	Regular software updates	Restricted user profile	Anti malware software	Do report xy7	Do not answer unknown emails	Do not surrender credential via phone	Do report unexpected error	Do report job offers	user response (C_1159)	company response (C_240)
User																	
201			Y	Y									Y			D--	D--
202	Y	Y		Y			Y	Y	Y	Y	Y			Y	Y	D++	D++
203			Y	Y	Y	Y		Y	Y			Y				D-	D-
204	Y	Y	Y		Y		Y	Y	Y	Y	Y				Y	C	C
205		Y	Y	Y	Y						Y				Y	C	C
206			Y	Y		Y	Y				Y	Y		Y		D+	D+
207			Y		Y	Y	Y	Y	Y	Y	Y				Y	C	C
208		Y				Y			Y	Y	Y			Y		D+	D+
209		Y	Y			Y			Y			Y	Y			D--	D--
210		Y		Y	Y	Y	Y	Y	Y		Y	Y			Y	C	C
211	Y		Y		Y	Y		Y					Y			D-	D-
212	Y	Y	Y		Y	Y	Y	Y	Y		Y	Y	Y		Y	C	C
213	Y	Y	Y				Y			Y	Y	Y			Y	C	C
214			Y		Y		Y	Y		Y	Y	Y	Y	Y		D+	D+
215		Y	Y		Y	Y			Y		Y		Y		Y	C	C
216					Y	Y	Y	Y	Y		Y		Y	Y	Y	D++	D++
217	Y			Y	Y		Y		Y	Y		Y	Y			D-	D-
218	Y	Y									Y	Y			Y	C	C
219					Y		Y		Y	Y						D-	D-
220	Y			Y	Y	Y					Y	Y	Y		Y	C	C
221		Y		Y	Y			Y		Y	Y	Y	Y		Y	C	C
222				Y		Y		Y			Y			Y		D+	D+
223				Y	Y	Y	Y		Y	Y	Y		Y		Y	C	C
224		Y	Y		Y	Y	Y				Y	Y	Y	Y	Y	D++	D++
225		Y	Y	Y		Y	Y			Y		Y				D-	D-
226	Y	Y	Y	Y			Y	Y			Y		Y		Y	C	C
227			Y					Y	Y	Y	Y	Y	Y		Y	C	C
228		Y	Y	Y		Y		Y		Y	Y	Y		Y		D+	D+
229		Y		Y	Y	Y		Y	Y	Y	Y				Y	C	C
230									Y		Y			Y		D+	D+
231	Y	Y					Y	Y		Y	Y		Y		Y	C	C
232		Y			Y	Y	Y	Y		Y	Y		Y	Y	Y	D++	D++

Table 7.12: Influence of Training "Reacting" on User and Company Response.

7.4.3 Results (Influence Module)

Analysis for Group 1, "Recognizing Phishing Attacks": Upon executing the training on recognizing phishing attacks, the first group shows the desired change (**Hypothesis 7.5**). In fact, the curve representing the user response has shifted from C_{1159} (asserted in the first fictitious phishing attack) to C_{20667} (asserted in the second fictitious phishing attack). The "company response" curve remains as is as the training appears to have no influence on C_{34}.

Consequently, the following exercise focuses on the user response C_{20667}. By applying ROSE2 [117], for Group 1, the following *CORE ANSWERS* were found.

Response Curve	Set of Security Questions	Quality of Classification	Core Answers	Quality of Classification
C_{20667}	*A, B, C, D, E*	*1.000*	*C, E*	***0.781***
C_{20667}	P_1, P_2, P_3, P_4, P_5	*0.344*	P_1, P_2, P_3, P_5	*0.344*
C_{20667}	R_1, R_2, R_3, R_4, R_5	*0.563*	R_1, R_2, R_3, R_4, R_5	*0.563*

Table 7.13: Core Answers to each Subset of Group "Recognizing".

For Group 1, the core (*C, E*) was identified, which shows a clear relationship with C_{20667}. This or a similar relationship is not established by the core answers (P_1, P_2, P_3, P_5) and (R_1, R_2, R_3, R_4, R_5) as the quality of classification is low (*0.344* and *0.563* respectively) and the number of attributes in the cores is high.

Next, for group 1, two sets of questions are analyzed conjunctively:

Response Curve	Set of Security Questions	Quality of Classification	Core Answers	Quality of Classification
C_{20667}	*A, B, C, D, E, P_1, P_2, P_3, P_4, P_5*	*1.000*	*C, E*	***0.781***
C_{20667}	*A, B, C, D, E, R_1, R_2, R_3, R_4, R_5*	*1.000*	*C*	*0.000*
C_{20667}	*P_1, P_2, P_3, P_4, P_5, R_1, R_2, R_3, R_4, R_5*	*1.000*	*P_1, R_5*	*0.000*

Table 7.14: Core Answers for two Subsets of Group "Recognizing".

For the security questions (*A, B, C, D, E, P_1, P_2, P_3, P_4, P_5*), the core (*C, E*) remains the same as before. For the set of security questions (*A, B, C, D, E, R_1, R_2, R_3, R_4, R_5*) and (*P_1, P_2, P_3, P_4, P_5, R_1, R_2, R_3, R_4, R_5*), the cores (*C*) and (*P_1, R_5*) are observed respectively. However, their classification quality is NIL.

Next, for Group 1, all three sets of security answers are taken into account:

Response Curve	Set of Security Questions	Quality of Classification	Core Answers	Quality of Classification
C_{20667}	*A, B, C, D, E, P_1, P_2, P_3, P_4, P_5, R_1, R_2, R_3, R_4, R_5*	*1.000*	*Empty*	*0.000*

Table 7.15: Core Answers for the three Subsets of Group "Recognizing".

For the selected questions (*A, B, C, D, E, P_1, P_2, P_3, P_4, P_5, R_1, R_2, R_3, R_4, R_5*), the core is empty and its classification quality is *0.000*.

Next, whether or not the quality of approximation can be increased by adding individual security questions/answers to the core is verified:

Response Curve	Core Answers to (A, B, C, D, E)	Quality of Classification	With Additional Answer to Core	Quality of Classification
C_{20667}	C, E	0.781	A, C, E	0.781
C_{20667}	C, E	0.781	B, C, E	**1.000**
C_{20667}	C, E	0.781	C, D, E	**1.000**

Table 7.16: Quality of Classification for Group "Recognizing".

In **Table 7.16** the triplets *(B, C, E)* and *(C, D, E)* are of interest as they yield a higher *QUALITY OF APPROXIMATION* than the core alone. This results is satisfactory as the classification quality cannot be better than *1.000*.

Next, the *ERROR OF CLASSIFICATION*, which occurs if one security process is removed from the triplets *(B, C, E)* and *(C, D, E)* is investigated.

Selected Answers	Quality of Classification	Answer to Omit	Quality of Approximation after Omission	Error of Classification
B, C, E	*1.000*	*B*	*0.781*	*0.219*
B, C, E	*1.000*	*C*	*0.062*	*0.938*
B, C, E	*1.000*	*E*	*0.062*	*0.938*
C, D, E	*1.000*	*C*	*0.156*	***0.844***
C, D, E	*1.000*	*D*	*0.781*	***0.219***
C, D, E	*1.000*	*E*	*0.156*	***0.844***

Table 7.17: Error of Classification for Omitted Attributes (Group "Recognizing").

Based on the above, the triplet of security questions *(C, D, E)* is interpreted as the most suited for establishing (describing) a relationship between answers to security questions and C_{20667}. Accordingly, (C, D, E) is chosen over (B, C, E) because omitting one of the core attributes causes a slightly smaller error of classification.

Analysis for Group 2, "Preventing Phishing Attacks": From the data in the Rough Set table there is neither change in the "user response" curve, C_{1159}, nor in the "company response" curve, C_{34}. This group is not followed up.

Analysis for Group 3, "Reacting to Phishing Attacks": The laboratory data indicates that the "company response" curve has changed. In fact, C_{34} has shifted to become C_{240} while the "user response" curve, C_{1159}, has remained unchanged. Accordingly, the data mining exercise focuses on C_{240}. By applying ROSE2 [117] the following core of answers for each set of security questions are found:

Response Curve	Set of Security Questions	Quality of Classification	Core Answers	Quality of Classification
C_{240}	*A, B, C, D, E*	*0.594*	*A, B, C, D, E*	*0.594*
C_{240}	P_1, P_2, P_3, P_4, P_5	*0.531*	P_1, P_2, P_3, P_4, P_5	*0.531*
C_{240}	R_1, R_2, R_3, R_4, R_5	*0.875*	R_3, R_4, R_5	***0.875***

Table 7.18: Core Answers to each Subset of Security Questions (Group "Reacting").

Despite a classification quality of *0.594* and *0.531*, the core related to the security questions (*A, B, C, D, E*) and (P_1, P_2, P_3, P_4, P_5) has not reduced in number. The core answers *($R_3, R_4,$* and *R5)* related to the security questions (R_1, R_2, R_3, R_4, R_5) show a reduction in size with a classification quality of *0.875*. Consequently, the core answers *($R_3, R_4,$* and *R5)* are followed up.

Next, it is verified whether the quality of approximation can be increased by adding additional answers to the core. We start bottom up and calculate the quality of approximation for four security processes. For this we concentrate on the third set of answers:

Response Curve	Answers to Core of (R_1, R_2, R_3, R_4, R_5)	Quality of Classification	Additional Answers to Core	Quality of Classification
C_{240}	R_3, R_4, R_5	*0.875*	R_1, R_3, R_4, R_5	***0.875***
C_{240}	R_3, R_4, R_5	*0.875*	R_2, R_3, R_4, R_5	***0.875***

Table 7.19: Quality of Classification for Group "Recognizing".

The above results indicate that the quality of approximation induced by the core cannot be enhanced any further. Accordingly, the answers *(R_3, R_4, R_5)* describe best the shape of C_{240}. For the sake of completeness, the error of classification induced by removing one core answer is looked at:

Selected Answers	Quality of Classification	Answers to Omit	Quality of Classification	Error of Classification
R_3, R_4, R_5	*0.875*	R_3	*0.781*	***0.107***
R_3, R_4, R_5	*0.875*	R_4	*0.094*	***0.893***
R_3, R_4, R_5	*0.875*	R_5	*0.094*	***0.893***

Table 7.20: Error of Classification for Core Answers of Group "Reacting".

7.4.5 Discussion

The core security questions centred on attack deployment techniques concerning groups (question *C*) and malware downloads (question *E*) appear to be the most suited piece of information to be conveyed in a security training on the *RECOGNITION OF*

PHISHING ATTACKS. They describe the "user response" curve, C_{20667}, with a classification quality of *0.781*. By adding one security question (centred on attack deployment techniques via phones) the "user response" curve can be described with a classification quality of *1.000*. As expected, the other security questions related to the prevention and reaction to phishing attacks do not increase the quality of classification. Finally, the small error of classification confirms the security questions *(C, D, E)* to be a good choice for delivering awareness training on the recognition of phishing attacks.

For the *PREVENTION OF PHISHING ATTACKS*, no calculations have been performed as their curves have remained unchanged by security awareness training.

For the *REACTION TO PHISHING ATTACKS*, analogous results have been obtained as for group 1: the triplet (R_3, R_4, R_5) has been found to be the core answers with a classification quality of *0.875* in describing C_{240}.

Finally, optimized security training can be assembled by using the questions (*C, D, E,* R_3, R_4, R_5) which delivers C_{20667}, C_{240} in one go.

7.5 Decision Module: Selection of Security Mechanisms

In this module a decision in terms of risk preferences is taken on which security training to adopt to counteract phishing.

7.5.1 Data Collection and Evaluation Methodology

The risk preferences of two decision makers in the online banking department of the remitter of the case study are displayed next. Seven points of their individual utility curves have been ascertained[79] by means of a questionnaire which is partly shown in

[79] Utility curves are based on the subjective perception of individuals and vary in time. Therefore, its general shape suffices for our purposes (rather than a precise mathematical description of it). The above utility curves were assessed in November 2007. A follow up assessment by Salvati [155] in March 2008 placed special emphasis on the "savings" portion of the chart. It showed:

Appendix M. By connecting the seven points the following utility curves have been obtained (**Figure 7.15**).

The online banking department is organized as a cost center rather than a profit center. Therefore, **Figure 7.15** displays "savings" and "losses" rather than "earnings" and "losses".

- The "savings" portion of the curve in March 2008 indicates a tendency to linearity (risk neutral) rather than risk preferring as suggested in November 2007.
- The general shape in the area "losses" could be confirmed.

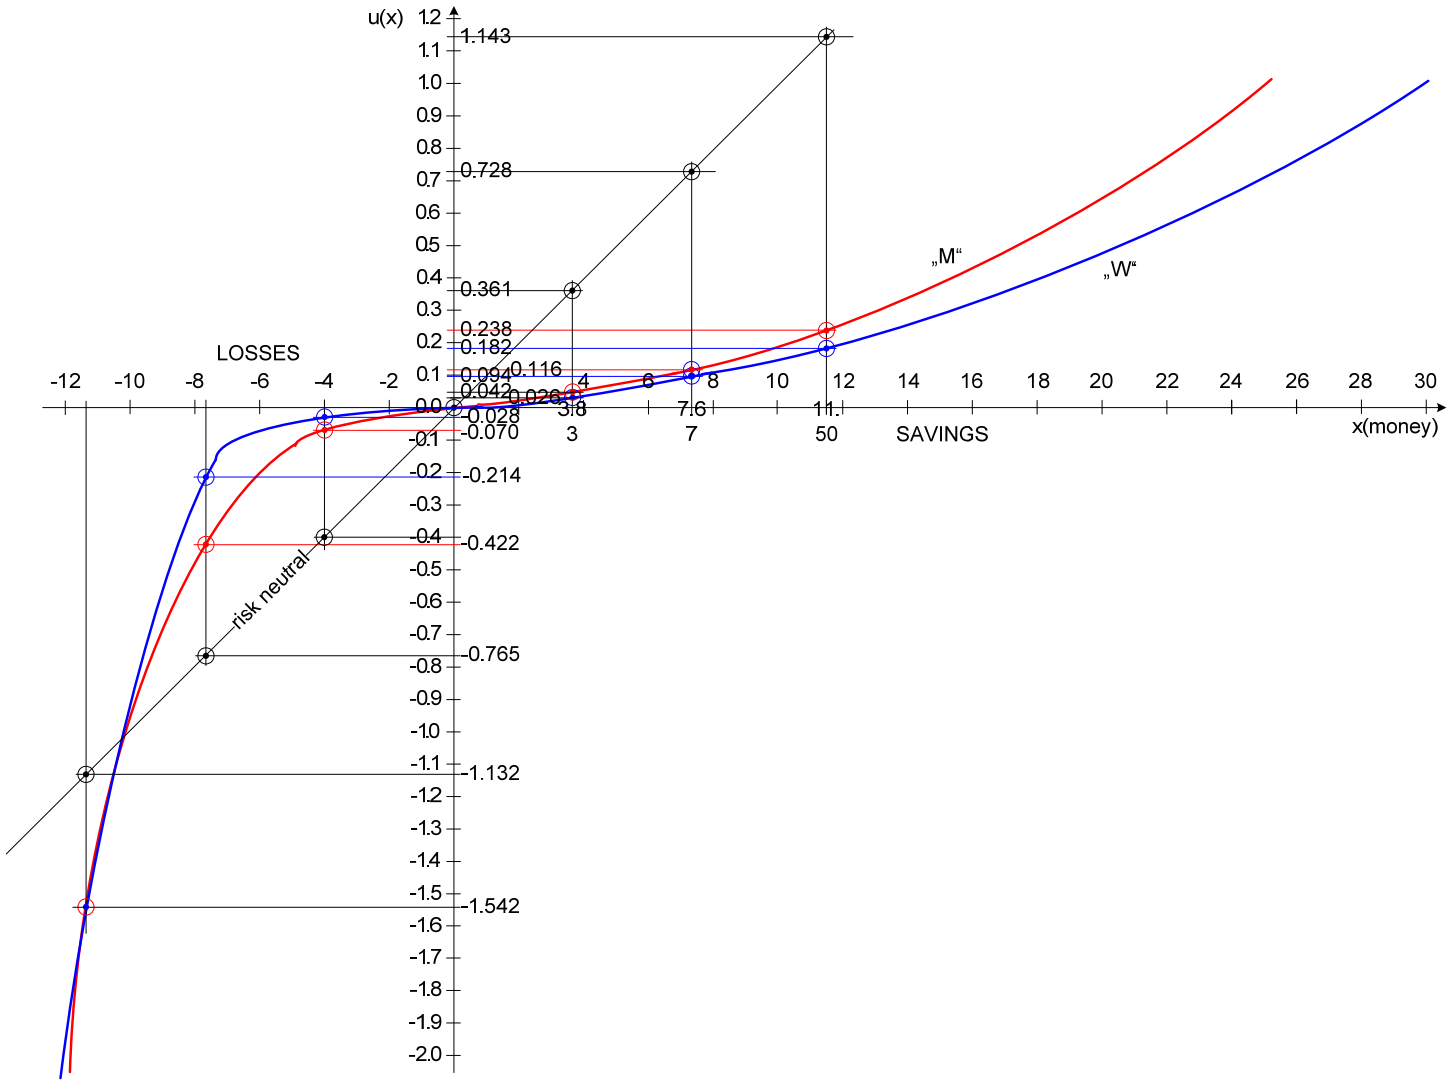

Figure 7.15: Utility Curves for Decision Makers “M” / “W”.

7.5.2 Decision Tree for Classic Phishing

The previous results from all other modules are merged in the Decision Module:

Data from the Process Module: A description of the phishing scenarios is obtained from the Process Module including data on financial losses and gains. The fictitious loss data is reported in **Figure 7.16**.

Data from the Function Module: From the Function Module the frequencies and probabilities of phishing attacks being successful are obtained. The probability calculation of the Function Module for all curves (C_{1159}, C_{20667}, C_{34}, and $C240$) shows:

Combination of Curves	Probability of Unsuccessful Attack	Probability of Successful Attack
C_{1159} with C_{34}	*0.016*	*0.984*
C_{1159} with C_{240}	*0.050*	*0.950*
C_{20667} with C_{34}	*0.052*	*0.948*
C_{20667} with C_{240}	*0.134*	*0.861*

Table 7.21: Probability of Successful Attack for C_{1159}, C_{20667}, C_{34}, C_{240}.

Data from the Influence Module: The relevant security questions for successful awareness training are obtained from the Influence Module. As three types of training for security awareness are distinguished, there are eight ($= 2^3$) possibilities to combine them. However, the security training "Prevention" does neither support the curves "user response" from decreasing nor the "company response" from increasing. Consequently, it is not considered it any further. Therefore, the number of possibilities of assembling security training reduces from eight ($= 2^3$) to four ($= 2^2$).

All of the above information is displayed in the following decision tree (**Figure 7.16**):

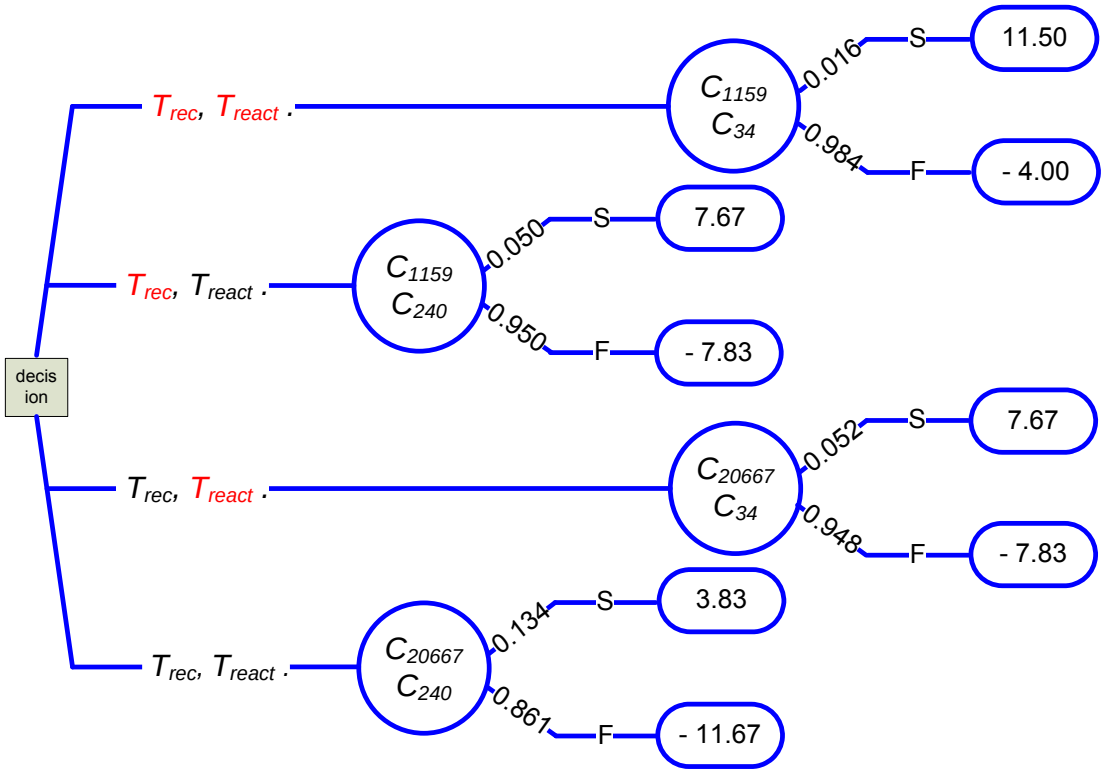

Figure 7.16: Decision Tree for T_{rec} and T_{react}.

Figure 7.16 shows a reduced decision set where T_{rec} signifies training designed to recognize phishing attacks and T_{react} signifies training that educates people on how to react to them. Letters **in red** (e.g., T_{rec}) mean that the training has not been executed while lettering **in bold black letters** (e.g., $\boldsymbol{T_{react}}$) means that training has been performed. The circles display the response curves (e.g., C_{1159} with C_{34}) induced by the training and the probabilities are taken from **Table 7.21** where phishing attacks are successful (S) or unsuccessful (F = failure). The figures in the boxes signify the savings or losses given a combination of security training.

7.5.3 Results (Decision Module)

Table 7.22 shows the utilities for the individual decision makers. They indicate that the decision makers are almost indifferent in choosing between **no** security training and suffering a financial loss. Moreover, a situation where security training for the recognition of and reaction to phishing attacks applies is the least desirable as it yields the lowest utility numbers (*-1.322* for "M" and *-1.324* for "W" respectively).

Given the above negative utility numbers, the value of a risk analysis is not calculated.

7.5.4 Discussion

The results show that the perfect vulnerability displayed by the company response cannot be significantly reduced by security awareness training. To prevent *EMPLOYEES* from becoming victims of classic phishing attacks, measures related to the temporary prevention of e-mails leaving the company in case of a phishing attack are much more promising rather than shutting down the opponents Internet site. Another approach would be to train selected employees in recognizing and reacting to phishing attacks. This approach may reduce costs and allow for a better reaction time.

The strategy adopted by financial companies to refund their *CUSTOMERS* in case they become victims of a phishing attack has proven successful. The above utility numbers indicate that decision makers are better off refunding their customers rather than training employees. Fortunately, this practice also positively influences media reports in Switzerland.

Because of the vulnerability shown by employees in falling for phishing e-mails, in order to counteract *ATTACKERS* the continuous evolution of security measures designed to achieve and maintain a higher security level than competitors appears appropriate. Alternatively, the role and number of money mules in the execution of phishing attacks could be investigated further as they may represent a limiting factor for the attacker.

Alternative	Utility for "M"	Utility for "W"	Utility for Risk Neutral
T_{rec}, T_{react}	***u(11.50, -4.00)*** = *0.016·u(11.50)+0.984·u(-4.00)= 0.016 · 0.238 + 0.984 · -0.070 =* ***-0.065***	***u(11.50, -4.00)*** = *0.016·u(11.50)+0.984·u(-4.00) = 0.016 · 0.182 + 0.984· -0.028 =* ***-0.025***	***u(11.50, -4.00)*** = *0.016·u(11.50)+0.984·u(-4.00)= 0.016 · 1.143 + 0.984 · -0.422 =* ***-0.396***
T_{rec}, T_{react}	***u(7.67, -7.83)*** = *0.050·u(7.67)+0.950·u(-7.83) = 0.050 · 0.116 + 0.950· -0.422 =* ***-0.395***	***u(7.67, -7.83)*** = *0.050·u(11.50)+0.950·u(-4.00) = 0.050 · 0.094 + 0.950· -0.214 =* ***-0.199***	***u(7.67, -7.83)*** = *0.050·u(7.67)+0.950·u(-7.83) = 0.050 · 0.728 + 0.950· -0.765 =* ***-0.690***
T_{rec}, T_{react}	***u(7.67, -7.83)*** = *0.052·u(7.670)+0.948·u(-7.83) = 0.052 · 0.116+ 0.948 · -0.422 =* ***-0.394***	***u(7.67, -7.83)*** = *0.052·u(7.67)+0.948·u(-7.83) = 0.052 · 0.094+ 0.948 · -0.214 =* ***-0.198***	***u(7.67, -7.83)*** = *0.052·u(7.67)+0.948·u(-7.83) = 0.052 · 0.728+ 0.948 · -0.765 =* ***-0.685***
T_{rec}, T_{react}	***u(3.83, -11.67)*** = *0.134·u(3.83)+0.861·u(-11.67) = 0.134 · 0.042 + 0.861 · -1.542 =* ***-1.322***	***u(3.83, -11.67)*** = *0.134·u(3.83)+0.861·u(-11.67) = 0.134 ·0.026 + 0.861 · -1.542 =* ***-1.324***	***u(3.83, -11.67)*** = *0.134·u(3.83)+0.861·u(-11.67) = 0.134 · 0.361 + 0.861 · -1.132 =* ***-0.926***

Table 7.22: Utility Numbers for the U-Curves for Decision Makers "M"/"W".

8. Conclusion and Outlook

In this chapter, the benefits of and limitations to the Four Modules are discussed and their practical implementation is addressed. In Section 8.3, further research topics are proposed. Finally, concluding remarks are provided.

8.1 Benefits of and Limitations to the Four Modules

In today's increasingly regulated business world, compliance management is spreading into fields like information security. This trend is disconcerting as it entails a bureaucratic and inefficient approach to the selection, implementation and maintenance of security mechanisms and related processes. Consequently, decisions concerning IS risks are driven by checklists rather than entrepreneurial spirit. In addition, as today's underlying concepts to the management of IS risks resembles a patchwork of concepts rather than a coherent framework, the Ambiguity, Likelihood, and Influence Problems prevail. To solve these problems the epistemological nature of IS risk has been revisited. As a result, this thesis provides a revised terminology and concepts, establishes a general approach for determining probabilities, incorporates the influence of the business and engineering contexts on these probabilities and enables decisions according to risk preferences of decision makers.

8.1.1 Benefits of and Limitations to the Process Module

The *AMBIGUITY PROBLEM* has been mitigated by considering the information system and business contexts when describing security related events in scenarios. Accordingly, the motivation and resources of an attacker executing a threat, the security processes which maintain security mechanisms, or the market a company operates in is valuable and indispensable information.

As many risks lie in the eye of an observer, the scenarios in the Process Module are chosen by the decision maker according to information needs. Scenarios "tell a story" by concatenating individual events. Taking into account the information needs of decision makers, a simple but powerful graphical notation has been developed which draws from state-based event trees.

Benefits: Operating a risk analysis in the business context provides a better understanding for the protection requirements of assets compared to the same risk analysis operated in the information system context only. Moreover, scenario charts allow for a limited insight in to the future of threats and security mechanisms and provide hints on the completeness of a risk analysis.

During the case study, the graphical notation of the Process Module supported communication between security specialists and senior executives by bridging the engineering and the business worlds.

Limitations: The information needs of a decision maker and the use of scenario charts are no guarantee for a comprehensive risk analysis. Consequently, scenarios may not sufficiently reflect the actual risk landscape.

8.1.2 Benefits of and Limitations to the Function Module

Solving the *LIKELIHOOD PROBLEM* required distinguishing between frequencies of a threat occurring and the probability of that threat overcoming security mechanisms. On one hand, frequency describes how many times in a period of time a threat is present and attempts to overcome a security mechanism. This number enables the decision maker to select a general protection strategy, i.e. to set up process-based security for recurring threats or to follow up on measures designed to ensure business continuity for nonrecurring threats. On the other hand, probability is determined by the interplay between threat and security mechanism. In this interpretation, probability reflects the notion that security mechanisms are vulnerable to threats if they are not ready to face them within a given time window, e.g., anti-virus software is not up-to date during the time when confronted with a new virus. Accordingly, probability "counts" how many times a threat is faster in reaching the security mechanism before it is ready to face it (which means that the threat prevails) or how many times the security mechanism has been updated on time before the threat reaches it (which means that the security mechanism prevails). A security mechanism which is always up-to-date when confronted with a threat is called perfect control; conversely, it is called perfect vulnerability.

Benefits: The previously described approach for solving the Likelihood Problem offers a general method for calculating probabilities based on time differences. This is key to the calculation of many probabilities and the data is readily available. For example, as shown in Chapter 3, the probabilities of cracking passwords given some characteristics like length, number of characters or computational resources of the attacker can actually be derived. As such, the complaint voiced many times by risk analysts that there is not enough data to measure probabilities can be put into perspective.

In practice, given the fact that probability figures can be obtained by measurements and calculation, the largest benefit lies in the avoidance of tedious discussions.

Limitations: The Function Module is limited by three basic problems. The first one is typical to any frequentist approach: although the measurements are simple in nature and can be performed by placing sensors in "neuralgic" areas of a company (e.g., the help line, e-mail gateways, servers used for maintenance), a sufficiently large number of data points is needed to determine probabilities which may not always be available.

The second problem is of a legal nature. Although processing probability in terms of the Function Module is feasible with the privacy requirements of many countries, possible breaches to data privacy of employees must be asserted diligently and acted upon to preserve it. This is especially the case when measuring, e.g., the responses of company personnel to specific Internet attacks as has been done in the case study.

The third problem relates to ambiguity and potentially arises if the risk analysis does not disclose the scope of measurement and calculation of probabilities. For example, the question whether all occurrences of a specific threat were considered or were just the specific instances that were directed against the company under investigation is relevant as it directly influences the results. If the scope is not agreed on prior to commencing the measurements and calculations there is room for misinterpretation of the results along the way.

8.1.3 Benefits of and Limitations to the Influence Module

The Influence Module explores relationships between security processes and success probabilities (i.e. the probability density functions of either the threat, the security mechanism, or both). The *INFLUENCE PROBLEM* and the *GOVERNANCE PROBLEM* are solved by devising the concepts of dependency, dispensability and significance of Rough Sets Theory. Dependency is reflected by the accuracy with which one or more security processes describes the probability density functions of a threat and a security mechanism or, in terms of RST, by the equivalence class structure induced by the attributes of one or more security processes approximating the equivalence class structure induced by the attributes of a threat and/or a security mechanism.

To ascertain the dispensability of security processes the concept of reducts has been introduced. A reduct is a minimal subset of a security processes which preserves the equivalence class structure induced by the original set of security processes. There may be more than one reduct common to a set of security processes and the security processes which are common to every reduct are called the core security processes; the others are called dispensable. The core security processes are regarded as an indispensable minimal set which influences the interaction of the pair threat/security mechanism. In terms of the

Governance Problem, the core security processes represent the indispensable minimal set with which head offices control the level of security at their branches.

Finally, the significance of security processes has been interpreted as the reduction in error of classification if one or more security processes were added to the core.

Benefits: Solving the Governance Problem offers the possibility for global companies to measurably steer risks by requiring their branches to implement specific security processes. For example, out of a given set of security processes, the Influence Module identifies how pronounced the influence of each one is on success probabilities of threats at a specific branch. By devising the core security processes, head offices find the optimal balance in choosing between fully centralized governance and a *laissez-faire* approach. The concept of significance is useful if head offices are willing to exercise their governance with more than the core security processes.

In the case study, the notions of dependent, dispensable and significant security processes were also applied to determine the relationship between security awareness training and the behaviour of users responding to a phishing attack. Consequently, just like the Function Module represents a general approach to calculate probabilities, the Influence Module represents a structured approach to identifying relationships between security processes and the interaction threat/security mechanisms. The Influence Module points to those security processes which must be taken into account in the Decision Module with vague, inconsistent, and a small amount of data.

Limitations: RST does not explain relationships among security processes but merely describes them. For example, which of two security process is to be omitted from further consideration if they are fully dependent on each other? Moreover, is this exclusion permanent or should it be reevaluated after a specific amount of time? To answer these questions, the Influence Module needs to be applied by an expert on a regular basis.

Another problem is the small amount of data which may entail putting the use of RST as a tool for statistical inference into perspective as future measurements may entirely change the interpretation of today's results.

An additional limitation to the force of expression of the Influence Module lies within the activities for gathering data on security processes. Such assessments may span weeks or months, which may contrast the time limit within which a decision must be reached.

8.1.4 Benefits of and Limitations to the Decision Module

Given a list of alternative security processes, security mechanisms and the risk preferences of a decision maker, solving the *DECISION PROBLEM* indicates an ideal *RISK BASED* approach to decision making. The decision maker decides on whether or not to adopt the security mechanisms (including security processes) which may fail or succeed in counter-

acting threats. Should a security mechanism fail then the money allocated for it is lost and a consequence of this failure becomes noticeable in the business context. Prior to selecting security mechanisms the decision maker faces an ill-specified decision problem which can be corrected by executing a risk analysis. However, a decision maker will execute a risk analysis only if it is worth the effort, i.e. the utility in having the risk analysis performed is higher than not having it performed. In this thesis, emphasis was put on keeping the Decision Module simple and to elaborate on the key aspects only. Additional features of "Howard's Model" such as a sensitivity analysis and time preferences were not treated.

Benefits: The risk preferences of decision makers can be applied to other decisions as long as they remain within the same monetary boundaries of the decisions for which they were originally assessed. Accordingly, the decision maker can delegate specific decisions by anchoring risk preferences in company risk policies. This also means that the decision maker is not required to have specialist knowledge in the field of IS security for the selection of security mechanisms. However, the aforementioned needs to be put into perspective as according to Apostolakis [156], purely *RISK BASED* decision making as it is presented in this work does not exist in practice and decision are *RISK INFORMED* instead. Thus, the value of the Decision Module is to provide decision support rather than strict decision automation.

The approach presented allows the estimation of the maximum price to a risk analysis to reduce uncertainty. This feature is welcomed for planning and budgeting purposes.

Finally, the strength of this approach lies in its simplicity and its "sense for reality". For example, it presupposes an IS to be a cost driver rather than a profit centre. Consequently, implementation and maintenance costs of security mechanisms are at the centre of attention as well as the financial consequences should they fail. In contrast, other decision approaches like the Return on Security Investment (ROSI) regard an IS to be an investment. This approach is far fetched from reality as it is impossible to attribute a ROSI to an IS.

Limitations: The concept of expected utility may not always deliver consistent expected results when comparing them with actual observed choices. Such an inconsistency was voiced in 1953 by Allais [157]. The Allais Paradox[80] states that under certain circumstances two independent lotteries cannot be merged into one lottery as the decomposability axiom would suggest.

Another inconsistency is the Ellsberg Paradox [158] in which observed choices violate the expected utility hypothesis. This is generally the case when decision makers are confronted with known risks versus unknown risks they will tend to decide to go with the known risk. This behaviour describes an attitude of preference for known risks over

[80] Refer to **Appendix N** for an example of the Allais Paradox.

unknown risks and is called *UNCERTAINTY AVERSION*. Uncertainty aversion contrasts risk aversion in the sense that risk aversion is determined with a growing consequence while uncertainty aversion is determined by ignorance of the probabilities.

The Decision Module emphasizes a subjectivist approach rather than a frequentist view on probability estimates. This presupposes a high level of professionalism both from decision makers and experts when assigning figures to prospect probabilities, preference probabilities, certain equivalents, potential business losses, etc. because their estimates may be biased[81], e.g., by ignoring base rates or overweighing data.

Finally, the Decision Module considers the risk preferences of one single decision maker. Although it could be argued that the risk preferences of *one* decision maker are sufficient to set the entrepreneurial course of a company, a decision mechanism may be envisaged where preferences of *multiple* decision makers are accommodated. However, such extensions face Arrows' Impossibility Theorem[82] [160] which states that a body cannot reach a conclusion that satisfies all decision makers when taking into account requirements of a fair voting mechanism such as non-dictatorship of single voters, accountability of all individual preferences, independence from irrelevant or unknown decision alternatives, positive association of social and individual values, achievability of every possible group preference.

[81] For an overview on such biases refer to see Kleindorfer et al [159].

[82] In social welfare, Arrows Impossibility Theorem states that no voting system can extract a decision from a multitude of voters' preferences to form an aggregated ranking accepted by the entire community.

8.2 Practical Implementation

Ideally, prior to a risk analysis, the basic scenario elements are preassembled from different sources and stored in a repository. Frequency curves of threats may be obtained from Internet companies offering such information services. Frequency curves describing the response of security mechanisms to attacks are measured in the companies themselves. Examples of such company-internal measurement points are the:

- help line (for measuring the time employees take to notify an ongoing attack)
- teams trusted with sending of circular e-mails (for measurements on the user responsiveness)
- teams trusted with patching servers (for measuring patching activities).

To transfer the methodology into a specific company a software tool offers valuable support. The following components are important:

Process Module: A graphical interface depicts the "state-event-state" diagrams.

Function Module: Central repositories contain temporal response curves (for threats and security mechanisms, see also above); cross-company databases which are used within the same or related industries considerably ease data gathering activities.

Influence Module: A central repository contains data on the implementation quality of security processes (e.g., at branches but also in head offices of global companies). This repository also includes measurements on the security awareness of employees.

Decision Module: A central repository contains utility curves of decision makers.

8.3 Further Work

Not all potential applications of this work have been fully explored. The following additional studies and further developments are proposed:

Process Module: It is proposed to investigate approaches from economic theory for estimating losses in reputation and to evaluate a potential introduction into IS risk management. For example, VaR techniques[83], based on historical data, could deliver valuable results with respect to financial business consequences if applied to an abundant series of

[83] For an excellent introduction refer to Linsmeier and Pearson [118].

losses arising from a specific attack. However, difficulties are expected with the amount of loss data for the calculation of a long term VaR[84].

Function Module: An empirical study is proposed on the probability distribution of user passwords to refine the uniform distribution of Chapter 3.8.

Influence Module: An empirical study is proposed **(1)** to select a base set of security processes for the Governance Problem, **(2)** to measure their implementation quality at companies and **(3)** to measure the security awareness of employees

Decision Module: In response to the Allais and Ellsberg Paradoxa, the Cumulative Prospect Theory of Kahnemann and Tversky [161, 162] as well as the Reference Dependent Preferences by Köszegi and Rabin [163] are proposed for further consideration. Moreover, a comparison of risk preferences of decision makers in cost centres as opposed to preferences in profit centres is suggested. Further research is also proposed for determining the maximum budget and size of a company security department[85] by the means laid out in Chapter 5.6. Finally, the Decision Module could be explored for decision processes involving the risk appetite of various divisions (business models) in large organizations.

[84] For an approach for calculating a long term VaR refer to Dowd et al [121].

[85] For the Economic Risk Capital for Operational Risk as required in the Basle II Accord [89].

8.4 Concluding Remarks

The successful application of the Four Modules in the Case Study indicates a considerable leap forward for the management of IS risks:

- the ambiguity in describing risk for classic and spear phishing has been greatly reduced
- the measurement and subsequent calculation of the probability of an attack being successful which relates to the company is much more accepted by decision makers rather than best "guesstimates"
- the display of the influence of specific security processes (in the Case Study: security questions) on probabilities is reassuring for decision makers because it allows for a reduction of the decision set
- the decision support based on risk preferences has been pointed out by senior executives as the most interesting area of application.

The presented approach comprehensively covers the management of IS risks and, in particular, cyber attacks. Examples are brute force, dictionary or phishing attacks, viruses and anti-viruses, hacking exploits, etc.

A final word of caution: although during his practical field work the author has occasionally been confronted with statements on *risk being the authoritative paradigm for decision making in all situations* it is important to emphasize that it is not. In particular, it should not be used instead of jurisdiction and laws. Is this what Bernstein [51] meant when he stated that *the mathematically driven apparatus of modern risk management contains the seeds of a dehumanizing and self-destructive technology*?

Bibliography

1. ISO/IEC-Guide73, *Risk Management —Vocabulary — Guidelines for Use in Standards*. 2002, ISO/IEO: Geneva.

2. ISO/IEC-TR-13335-1, *Information Technology - Guidelines for the Management of IT Security - Part 1: Concepts and Models for IT Security*. 1996: Geneva.

3. CSE-MG2. *A Guide to Security Risk Management for Information Technology Systems*. 1996. Ottawa: Government of Canada, P.O. Box 9703, Terminal, Ottawa, Ontario, Canada, K1G 3Z4.

4. Stoneburner, G., A. Goguen, and A. Feringa, *Risk Management Guide for Information Technology Systems*. 2002, National Institute of Standards and Technology (NIST), Technology Administration, U.S. Department of Commerce: Gaithersburg, MD.

5. Leiner, B., et al. *A Brief History of the Internet*. [Internet] January 1999 [cited March 2008; Available from: http://arxiv.org/abs/cs.NI/9901011.

6. Dunn, M. and I. Wigert, *International CIIP Handbook 2004, An Inventory and Analysis of Protection Policies in Fourteen Countries*, ed. A. Wenger and J. Metzger. 2004: Swiss Federal Institute of Technology Zurich.

7. Weijnen, M. *Critical Infrastructures at Risk - The Need for Innovation*. [Powerpoint presentation] 2005 [cited December 2005]; Available from: http://www.lsa.ethz.ch/news/Zurich_ETH_220605.pdf.

8. Freeh, L.J. *Statement for the Record of Louis J. Freeh, Director Federal Bureau of Investigation on Cybercrime Before the Senate Committee on Judiciary Subcommittee for the Technology, Terrorism, and Government Information Washington, D.C.* [Internet] March 28th, 2000 [cited May 29th, 2007]; Available from: http://permanent.access.gpo.gov/lps10084/www.cybercrime.gov/freeh328.htm.

9. GAO-07-705, *Cybercrime - Public and Private Entities Face Challenges in Addressing Cyber Threats*. 2007, United States Government Accountability Office (GAO).

10. Kröger, W., et al., *White Paper on Managing and Reducing Social Vulnerabilities from Coupled Critical Infrastructures*. 2006, International Risk Governance Council, IRGC: Geneva. p. 68.

11. Skoudis, E. and SANS, *Computer and Network Hacker Exploits*. 2003.

12. Schneier, B., *Secret and Lies - Digital Security in a Networked World*. 2000, New York: John Wiley & Sons.

13. Schwarz, M., *Universe of Discourse für IT Risk Management in global tätigen Unternehmen*, in *Fachbereich Wirtschaftswissenschaften*. 2005, Hochschule Liechtenstein: Vaduz.

14. Schmid, G., *On the Existence of a Global System for the Interception of Private and Commercial Communications (ECHELON Interception System)*. 2001, European Parliament: Brussels.

15. Cachin, C., et al. *Malicious- and Accidental-Fault Tolerance for Internet Applications (MAFTIA) - Reference Model and Use Cases*. August 2000 [cited March 2008; Available from: http://www.maftia.org/deliverables/D1.pdf.

16. Kunz, A., *Reich der Unsitte*, in *Weltwoche 25/08*. 2008: Zürich.

17. von Lucius, R., *Am Boden, in der Luft und im Cyberspace*, in *Frankfurter Allgemeine (FAZ)*. August 14th, 2008: Frankfurt.

18. *State of Internet Security (Q107)*. [Internet] 2007 [cited March 2008; Available from: http://www.webroot.com/.

19. Fyffe, G., *Addressing the Insider Threat.* Network Security, 2008(3): p. 11-14.

20. Shinder, D.L., *Scene of the Cybercrime: Computer Forensics Handbook*. 2002, Rockland, MA 02370, USA: Syngress Publishing. Inc.

21. Serdiouk, V. *Technologies for Protection Against Insider Attacks on Computer Systems*. in *MMM-ACNS 2007*. 2007. St. Petersburg, Russia: Springer-Verlag Berlin Heidelberg.

22. Chen, T. and C. Davis. *An Overview of Electronic Attacks*. 2006 [cited July 22nd, 2007]; Available from: http://engr.smu.edu/~tchen/papers/dig-forensics06.pdf.

23. Sarbanes, P.S. and M. Oxley, *Sarbanes-Oxley Act*. 2002, U.S Government Printing Office. p. 66.

24. *Gramm-Leach-Bliley Act*. [Internet] US Congress, November 11th, 1999 [cited July 24th, 2007]; Available from: http://banking.senate.gov/conf/confrpt.htm.

25. *Directive 95/46/EC of the European Parliament and of the Council (The Data Protection Directive)*. [Internet] October 1995 [cited July 24th, 2007]; Available from: http://ec.europa.eu/justice_home/fsj/privacy/law/index_en.htm.

26. *Bundesdatenschutzgesetz*. [Internet] 1990. Bundesrepublik Deutschland, latest update: January 14th, 2003 [cited July 24th, 2007]; Available from: http://www.bmi.bund.de/.

27. *Tutela delle persone e di altri soggetti rispetto al trattamento dei dati personali* [Internet] 1996. Italian Parliament, latest update: December 28th, 2001 [cited July 24th, 2007]; Available from: http://www.garanteprivacy.it/garante/doc.jsp?ID=28335.

28. *Health Insurance Portability and Accountability Act*. [Internet] 1996. US Congress [cited July 24th, 2007]; Available from: http://www.cms.hhs.gov/HIPAAGenInfo/Downloads/HIPAALaw.pdf.

29. ISO/IEC-17799, *Information Technology - Security Techniques- Code of Practice for Information Security Management*. 2005: Geneva. p. 1-128.

30. *COBIT Management Guidelines*, ed. C.S. Committee. 2000, Rolling Meadows, IL: Information System Audit and Control Foundation, IT Governance Institute.

31. BSI, *IT-Grundschutzhandbuch*. 2004, Bundesamt für Sicherheit in der Informationstechnik: Bonn.

32. Geissler, C. *Was ist ... Compliance Management?* [Internet] 2004 [cited May 27th, 2007]; Available from: http://www.harvardbusinessmanager.de/img/cat/HBMO/compliancemanagement-Was.pdf.

33. Dummer, S. and C. Locher. *Compliance durch Standards für Informationssicherheit - Untersuchung von gesetzlichen Anforderungen an das Management der Informationssicherheit und deren Erfüllung durch ISO27001/17799*. [Internet] 2006 [cited March 29th, 2007]; Available from: http://www.vda.de/de/service/bestellung/downloads/VDA_Compliance_durch_Standards_fuer_Informationssicherheit.pdf.

34. ISO/IEC-17799, *Information Technology - Code of Practice for Information Security Management*. 2000, ISO/IEO: Geneva. p. 1-84.

35. ISF, *Threat and Vulnerability Assessment*. 2005, Information Security Forum: London. p. 49.

36. Team, C.P., *CMMI for Development*. 2006, Software Engineering Institute, Carnegie Mellon University: Pittsburgh. p. 587.

37. Salvati, D., *Maturity of IT Risk Management at Credit Suisse (Internal Paper)*. 2007: Zurich. p. 6.

38. Burrell, G. and G. Morgan, *Sociological Paradigms and Organisational Analysis*. 1979, London: Heinmann.

39. Althaus, C.E., *A Disciplinary Perspective on the Epistemological Status of Risk.* Risk Analysis, 2005. **25**(3): p. 567-587.

40. ISF, *ROSI - Return on Security Investment*. 2005, Information Security Forum: London.

41. Altorfer, P. *Präsentation der Umfrage-Ergebnisse zu ROSI*. Security Zone, 2006 [cited March 2008]; Available from: http://www.isss.ch/fileadmin/publ/agrosi/ROSI_Survey-Results.pdf.

42. Siponen, M. and R. Willison. *A Critical Assessment of IS Security Research Between 1990-2004*. in *15th European Conference on Information Systems (ECIS)*. 2007. St. Gallen (Switzerland).

43. *Oxford English Dictionary*. [Internet] 2006 [cited March 13th, 2006]; Available from: http://dictionary.oed.com.

44. NBS, *Federal Information Processing Standards Publication (FIBS PUB)*. 1974, U.S. Department of Commerce, National Bureau of Standards: Washington, D.C.

45. Salvati, D. and M. Diergardt. *Towards a Scenario Based Risk Model for Information Systems*. [Internet] 2007 [cited March 2008]; Available from: http://www.lsa.ethz.ch/people/phd/salvatid/DS-Scenario-Based-Risk-Model.pdf.

46. Marquis de Laplace, P.S. *A Philosophical Essay on Probabilities*. [Internet] 1902 [cited March 2008]; Available from: http://www.archive.org/details/philosophicaless00lapliala.

47. Carnap, R., *Logical Foundations of Probability*. 1950, Chicago: University of Chicago Press.

48. Popper, K., *The Propensity Interpretation of the Calculus of Probability and the Quantum Theory.* The Colston Papers, 1957. **9**: p. 65-70

49. Popper, K., *The Propensity Interpretation of Probability.* The British Journal for the Philosophy of Science, 1959. **10**(37): p. 25-42.

50. Hayek, A. *Interpretations of Probability*. [Internet] July 2007 [cited March 2008]; Available from: http://plato.stanford.edu/entries/probability-interpret/.

51. Bernstein, P.L., *Against the Gods - The Remarkable Story of Risk*. 2nd ed. 1998, New York: John Wiley & Sons, Inc. 383p.

52. Fisher, R., *Statistical Methods for Research Workers*. 1925, Edinburgh: Oliver and Boyd. 239 pp.

53. Neyman, J. and E. Pearson, *On the Problem of the Most Efficient Tests of Statistical Hypotheses.* Philosophical Transactions of the Royal Society of London. Series A, Mathematical and Physical Sciences 1933. **231**: p. 289-337.

54. Neyman, J., *Outline of a Theory of Statistical Estimation Based on the Classical Theory of Probability.* Philosophical Transactions of the Royal Society of London. Series A, Mathematical and Physical Sciences 1937. **236** (767): p. 333-380.

55. Cooke, R.M., *Experts in Uncertainty - Opinion and Subjective Probability in Science*. Environmental Ethics and Science Policy. 1991, New York: Oxford University Press.

56. Morgan, M.G., M. Henrion, and M. Small, *Uncertainty - A Guide to Dealing with Uncertainty in Quantitative Risk and Policy Analysis*. 1990, New York: Cambridge University Press.

57. *Northern Prairie Wildlife Research Center, The Insignificance of Statistical Significance Testing*. [Internet] U.S. Department of the Interior, August 3rd, 2007 [cited August 15th, 2007]; Available from: http://www.npwrc.usgs.gov/resource/methods/statsig/stathyp.htm.

58. Kain, Z., *The Legend of the P Value.* International Anesthesia Research Society, 2005. **101**(05): p. 1454-1456.

59. Howson, C., *Theories of Probability.* The British Journal for the Philosophy of Science, 1995. **46**(1): p. 1-32.

60. Diergardt, M., *Modeling Complex Scenarios in Computer Based Information Systems for Risk Analysis*. 2006, Laboratory for Safety Analysis: ETH Zurich.

61. Wharton, F., *Risk Management: Basic Concepts and General Principles*, in *Risk: Analysis, Assessment and Management*, F. Wharton, Editor. 1992, John Wiley & Sons.

62. *Encyclopedia Britannica Online*. [Internet] 2007 [cited April 14th, 2007]; Available from: http://info.britannica.co.uk.

63. Kolmogoroff, A.N., *Grundbegriffe der Wahrscheinlichkeitsrechnung, Ergebnisse der Mathematik und ihrer Grenzgebiete*. 1933, Berlin: Springer-Verlag.

64. Kaplan, S. and B.J. Garrick, *On the Quantitative Definition of Risk.* Risk Analysis, 1981. **1**(No. 1): p. 11-27.

65. Stamatelatos, M., et al., *Probabilistic Risk Assessment Techniques for NASA Managers and Practitioners*, in *NASA Headquarters of Safety and Mission Assurance*. 2002: Washington, DC.

66. Kröger, W. and R. Mock, *Methoden der Risikoanalyse und des Risikomanagements - Lecture Notes*. 2004: Zürich.

67. Kyas, O. and M. Campo, *IT Crackdown. Sicherheit im Internet*. 2000, Bonn: MITP.

68. Eckert, C., *IT-Sicherheit: Konzepte – Verfahren – Protokolle*. 2003, München/Wien: Oldenbourg.

69. SANS. *WhatWorks*. 2006 [cited January 30th, 2007]; Available from: http://www.sans.org/whatworks/poster.pdf.

70. *Merriam-Webster OnLine English Dictionary*. 2007 [cited February 4th, 2007]; Available from: http://www.m-w.com/.

71. Jakobsson, M. *Modeling and Preventing Phishing Attacks*. [Internet] 2005 [cited May 2007]; Available from: http://www.informatics.indiana.edu/markus/papers/phishing_jakobsson.pdf.

72. Schneier, B. *Attack Trees*. 1999 [cited 2007]; Available from: http://www.schneier.com/paper-attacktrees-ddj-ft.html#rf8.

73. Moore, A., R. Ellison, and R. Linger. *Attack Modeling for Information Security and Survivability - Survivable Systems*. [Internet] March 2001 [cited January 2007]; Available from: http://www.sei.cmu.edu/pub/documents/01.reports/pdf/01tn001.pdf.

74. Lee, W., D. Grosh, and C. Tillman, *Fault tree analysis, methods, and applications - a review.* IEEE Transactions, 1985: p. 194-203.

75. Roberts, N.H., et al., *Fault Tree Handbook (NUREG-0492)*. 1981, Washington, D.C.: U.S. Nuclear Regulatory Commission. 1-209.

76. Schubert, M., *FMEA - Fehlermöglichkeits- und Einflussanalyse*. 1. ed. DGQ-Schrift 11-13. 1993, Frankfurt/Main: Deutsche Gesellschaft für Qualität e.V. 1-48.

77. Lahres, H., *Einführung in die diskreten Markoff-Prozesse und ihre Anwendungen*. Die Wissenschaft. Vol. 120. 1964, Braunschweig: Friedrich Vieweg & Sohn.

78. VDI-4008-Blatt3, *Markoff-Zustandsänderungsmodelle mit endlich vielen Zuständen*. 1999, Beuth Verlag GmbH: Berlin. p. 17.

79. Petri, C.A., *Kommunikation mit Automaten*, in *Schriften des Rheinisch-Westfälischen Institutes für Instrumentelle Mathematik an der Universität Bonn*. 1962, University of Bonn: Bonn. p. 1-128.

80. Garrett, C. and G.E. Apostolakis, *Context in the Risk Assessment of Digital Systems*. Risk Analysis, 1999. **19**(1).

81. Weisstein, E. *Wolfram Mathworld*. [Internet] May 11th, 2008 [cited May 12th, 2008]; Available from: http://mathworld.wolfram.com/Likelihood.html.

82. Goel, A. and R.J. Graves, *Electronic system reliability: Collating prediction models*. IEEE Transactions on Device and Materials Reliability, 2006. **6**(2): p. 258-265.

83. Pan, J. *Software Reliability*. [Internet] 1999 [cited June 3rd, 2007]; Available from: http://www.ece.cmu.edu/~koopman/des_s99/sw_reliability/.

84. Kröger, W. and R. Mock, *Risiko und Sicherheit - Lecture Notes at ETH Zurich*. 2005: Zürich.

85. Pham, H.E., *Handbook of Reliability Engineering*, ed. H. Pham. 2003, London: Springer-Verlag. 664 p.

86. ISO/IEC-TR-13335-4, *Information Technology - Guidelines for the Management of IT Security - Part 4: Selection of Safeguards*. 2000: Geneva.

87. Verhoef, C., *Quantifying the Value of IT-Investments.* Science of Computer Programming, 2005. **56**: p. 315-342.

88. Canal, A.V., *ISM3 1.2: Information Security Management Maturity Model*, E. Stansfeld, Editor. 2006, Institute for Security and Open Methodology (ISECOM). p. 1-81.

89. Supervision, B.C.o.B., *International Convergence of Capital Measurement and Capital Standards*. 2006, Basel: Bank for International Settlements. 333.

90. Saaty, T.L., *The Analytic Hierarchy Process*. 1980, New York: McGraw-Hill.

91. Edwards, W., *How to use multiattribute utility measurement for social decision making.* IEEE Transactions on Systems, Man and Cybernetics, 1977. **SMC- 7**: p. 326-340.

92. Olson, D.L. and J.F. Courtney, *Chapter 10: Modeling Selection Decisions*, in *Decision Support Models and Expert Systems*, Houst, Editor. 1998, Dame Publications: Houston. p. 140-154.

93. Goto, J.H., M.E. Lewis, and M.L. Puterman, *Coffee, Tea, or ...?: A Markov Decision Process Model for Airline Meal Provisioning.* informs, 2004. **38**(1): p. 107-118.

94. Figueira, J., S. Greco, and M.E. Ehrgott, *Multiple Criteria Decision Analysis: State of the Art Surveys*. 1st edition ed. 2004, New York: Springer. pp. 1045.

95. Gheorghe, A. and R. Mock. *Employing Fuzzy Logic into Regional Risk Assessment and Safety Management*. in *Probabilistic Safety Assessment and Management, -ESREL 96-PSAM-III*. 1996. Crete: Springer-Verlag.

96. Pawlak, Z. and R. Slowinski, *Rough Set Approach to Multi-attribute Decision Analysis.* European Journal of Operational Research, 1994. **72**(3): p. 443-459.

97. Yang, J.-B. and M.G. Singh, *An Evidential Reasoning Approach for Multiple-Attribute Decision Making with Uncertainty.* IEEE Transactions on Systems, Man, and Cybernetics, 1994. **24**(1).

98. Huynh, V.N. and Y. Nakamori, *Multiple Attribute Decision Making Under Uncertainty: The Evidential Reasoning Approach Revisited.* IEEE Transactions on Systems, Man and Cybernetics, 2006. **36**(4): p. 804-822.

99. Zgurovskii, M.Z. and N.D. Pankratova, *An Information Platform for Scenario Analysis in Technology Foresight Problems.* Cybernetics and Systems Analysis, 2003. **39**(4): p. 564 - 575.

100. Plous, S., *The psychology of judgment and decision making*. Series in Social Psychology, ed. C. Rogers and J. Belser. 1993: McGraw-Hill.

101. Bradfield, R., et al. (2005) *The Origins and Evolution of Scenario Techniques in Long Range Business Planning*. Futures, 37(8): **Volume**, 795-812

102. *Learning from the Future*, ed. L. Fahey and R. Randall. 1998: John Wiley and Sons.

103. Salvati, D., *Telephone Interview with Max Moser, member of CERT at Credit Suisse*. 2006.

104. Kaufman, C., R. Perlman, and M. Speciner, *Network Security: PRIVATE Communication in a PUBLIC World*. 2002, Upper Saddle River, NJ: Prentice-Hall.

105. Crutchfield, S. *Joy of Convolution*. [Internet] 1999 [cited September 2007]; Available from: http://www.jhu.edu/~signals/index.html.

106. Czitrom, V., *One-Factor-at-a-Time Versus Designed Experiments.* The American Statistician, 1999. **53**(2): p. 126-131.

107. Estivill-Castro, V., *Why so many clustering algorithms - A Position Paper.* ACM Special Interest Group on Knowledge Discovery and Data Mining, 2006. **4**(1): p. 65 - 75.

108. Faber, V., *Clustering and the Continuous k-Means Algorithm.* Los Alamos Science, 1994(22): p. 138 - 144.

109. Wooldridge, S. *Bayesian Belief Networks*. [Internet] 2003 [cited February, 2008]; Available from: http://www.mrcmekong.org/download/programmes/ep/Bayesian_Network_Reading2.pdf.

110. Backhaus, K., et al., *Multivariate Analysemethoden*. 11th ed. 2006, Berlin: Springer. pp. 830.

111. Pawlak, Z., *Rough Sets.* International Journal of Computer and Information Sciences, 1982. **11**(5): p. 341-356.

112. Rauszer, C., *Reducts in information systems.* Fundamenta Informaticae, 1991(15): p. 1-12.

113. Skowron, A. and C. Rauszer, *The discernibility matrices and functions in information systems*. Intelligent Decision Support. Handbook of Applications and Advances of Rough Sets Theory, ed. R. Slowinski. 1992, Dordrecht: Kluwer. 331-362.

114. Delic, D., H.-J. Lenz, and M. Neiling. *Improving the Quality of Association Rule Mining by Means of Rough Sets*. Free University of Berlin, Institute of Applied Computer Science, Garystr. 21, D-14195 Berlin, Germany [Internet] 2002 [cited December 12th, 2008; Available from: http://cis.cs.tu-berlin.de/~mneiling/publications/delic_lenz_neilingSMPS2002.pdf.

115. Zhong, N., J. Dong, and S. Oshuga, *Using Rough Sets with Heuristics for Feature Selection.* Journal of Intelligent Information Systems, 2001(16): p. 199-215.

116. Nguyen, H.S., *Discretization Problem for Rough Sets Methods*. RSCTC'98, LNAI 1424, ed. L. Polkowski and A. Skowron. 1998, Berlin Heidelberg: Springer-Verlag. 545-552.

117. ProSoft. *ROSE 2 - Rough Set Data Explorer, User`s Guide*. [Internet] 1999 [cited 2004].

118. Linsmeier, J.T. and N.D. Pearson. *Risk Measurement: An Introduction to Value at Risk*. [Internet] 1996 [cited March 2008]; Available from: http://econpapers.repec.org/.

119. Holton, G.A. *Value at Risk - Theory and Practice*. [Internet] 2003 [cited March 2008]; Available from: www.value-at-risk.net.

120. Simons, K., *Value at Risk? New Approaches to Risk Management.* New England Economic Review, 1996. **1996**(5): p. 3-13.

121. Dowd, K., D. Blake, and A. Cairns. *Long-Term Value at Risk*. [Internet] 2003 September 2003 [cited March 2008]; Available from: http://www.lse.ac.uk/ubs/pdf/dp17.pdf.

122. Artzner, P., et al. *Coherent Measures of Risk*. [Internet] 1998 [cited December 11th, 2008]; Available from: http://www.math.ethz.ch/~delbaen/ftp/preprints/CoherentMF.pdf.

123. Howard, R.A., *The Principles and Applications of Decision Analysis*. Risk Preference, ed. R.A. Howard and J.E. Matheson. 1984: Strategic Decisions Group. 629 - 663.

124. Howard, R.A., *The Foundation of Decision Analysis.* IEEE Transactions on Systems, Man, and Cybernetics, 1968. **4**(3).

125. Howard, R.A., *Decision Analysis: Practice and Promise.* Management Science, 1988. **34**(6): p. 679-695.

126. Howard, R.A., *Decision Analysis (Lecture Notes)*. 2004, University of Stanford: Stanford.

127. Cain, P., *Vocabulary of Phishing Terms*. 2004, The Financial Services Technology Consortium.

128. *Identity Theft*. 2006, Bank of New York: New York.

129. Warner, B. *Billions of Phishing E-mails Sent Monthly*. [Internet] May 6th, 2004 [cited February 4th, 2008]; Available from: http://www.ladlass.com/ice/archives/cat_phishing_identity_theft.html.

130. Ollmann, G. *The Vishing Guide*. [Internet] May 2007 [cited September 2nd, 2007]; Available from: http://www.iss.net/documents/whitepapers/IBM_ISS_vishing_guide.pdf.

131. *SMS phishing on the rise in SE Asia?* [Internet] April 2007 [cited October 2007]; Available from: http://www.f-secure.com/weblog/archives/archive-042007.html#00001173.

132. Emigh, A. *The Crimeware Landscape: Malware, Phishing, Identity Theft and Beyond*. [Internet] September 19th, 2006 [cited September 1st, 2007]; Available from: http://www.cyber.st.dhs.gov/crimeware-ittc.pdf.

133. Salvati, D., *Swiss Federal Police, Phishing Attacks in Switzerland - Notes of Interview with Stephan Glaus*. December 2007: Berne.

134. Srivastava, T. *Phishing and Pharming: The Deadly Duo*. [Internet] January 29th, 2007 [cited August 28th, 2007]; Available from: http://www.sans.org/reading_room/whitepapers/privacy/1731.php.

135. Drake, C., J. Oliver, and E. Koontz. *Anatomy of a Phishing Email*. in *CEAS 2004 - First Conference on Email and Anti-Spam*. 2004. Mountain View, California, USA.

136. Jagatic, T., et al. *Social Phishing*. [Internet] 2005 [cited September 2007]; Available from: http://www.indiana.edu/~phishing/social-network-experiment/phishing-preprint.pdf.

137. Jakobsson, M. *The Human Factor in Phishing*. [Internet] 2007 [cited September 2007]; Available from: http://www.informatics.indiana.edu/markus/papers/aci.pdf.

138. Sheng, S., et al. *Anti-Phishing Phil: The Design and Evaluation of a Game That Teaches People Not to Fall for Phish*. in *Symposium On Usable Privacy and Security (SOUPS)*. 2007. Pittsburgh, Pennsylvania, USA.

139. Kuramaguru, P., et al., *Protecting People from Phishing: The Design and Evaluation of an Embedded Training Email System*. 2006, Carnegie Mellon: Pittsburgh, Pennsylvania, USA. p. 18.

140. Dhamija, R., J. Tygar, and M. Hearst. *Why Phishing Works*. in *CHI 2006*. 2006. Quebec, Canada: ACM 1-59593-178-3/06/0004.

141. Abad, C. *The Economy of Phishing: A Survey of the Operations of the Phishing Market*. [Internet] 2006 [cited June 4th, 2007]; Available from: http://www.cloudmark.com/serviceproviders/research/messaging_security/.

142. Dittrich, D. *Attack of the Zombies and How to Respond (Webcast)*. [Internet] March 18th, 2005 [cited March 2007]; Available from: http://searchsecurity.techtarget.com/.

143. Cole, A., M. Mellor, and D. Noyes. *Botnets: The Rise of the Machines*. [Internet] 2006 [cited February 4th, 2008]; Available from: http://www.mellorsecurity.com/Botnets.pdf.

144. Ramzan, Z. and C. Wüest. *Phishing Attacks: Analyzing Trends in 2006*. in *CEAS 2007 - Fourth Conference on Email and Anti-Spam*. 2007. Mountain View, CA, US.

145. Olzak, T. *DNS Cache Poisoning: Definition and Prevention*. [Internet] March, 2006 [cited February 4th, 2008]; Available from: http://www.infosecwriters.com/text_resources/pdf/DNS_TOlzak.pdf.

146. GAO-07-1128, *U.S.-China Economic and Security Review Commission*. 2007, United States Government Accountability Office (GAO).

147. Moll, D. *State of Spyware, Q1 2006 - A Review and Analysis of the Impact of Spyware on Consumers and Corporations*. [Internet] 2006 [cited September 2007]; Available from: http://www.antyspyware.pl/state-of-spyware/2006-q1-sos.pdf.

148. Salvati, D., *Swiss Federal Police, Phishing Attacks in Switzerland - Notes of Interview with Stephan Glaus*. October 2007: Berne.

149. Salvati, D., *Credit Suisse, Phishing Attacks in Switzerland - Notes of Interviews with Michael Fossati*. September 2007: Zurich.

150. Aitchinson, J. and J. Brown, *The Lognormal Distribution*. 4th ed. Monographs by the University of Cambridge Department of Applied Economics. Vol. 5 (15). 1969, New York. pp 176.

151. Wessa, P. *Maximum-likelihood Lognormal Distribution Fitting (v1.0.1) in Free Statistics Software (v1.1.22-r4)*. [Internet] 2007 [cited December 17th, 2007]; Available from: http://www.wessa.net/rwasp_fitdistrlnorm.wasp/.

152. Schwartz, S. and S. Yeh, *On the distribution function and moments of power sums with log-normal components.* Bell System Tech. J., 1982. **61**: p. 1441 - 1462.

153. Joshua Lam, C.-L. and T. Le-Ngoc, *Estimation of Typical Sum of Lognormal Random Variables using Log Shifted Gamma Approximation.* IEEE Communication Letters, 2006. **10**(4).

154. Beaulieu, N., A. Abu-Dayya, and P. McLane, *Estimating the Distribution of a Sum of Independent Lognormal Random Variables.* IEE Transactions on Communications, 1995. **43**(12): p. 2869-2873.

155. Salvati, D., *Credit Suisse - Phishing Attacks in Switzerland, Notes of Interview with Martin Walder*. March 2008: Zurich.

156. Apostolakis, G.E., *How Useful is Quantitative Risk Assessment?* Risk Analysis, 2004. **24**(3): p. 515-520.

157. Allais, M., *Le Comportement de l`Homme Rationnel devant le Risque: Critique des Postulats et Axiomes de l`Ecole Americaine.* Econometrica, 1953. **21**(4): p. 503-546.

158. Ellsberg, D., *Risk, Ambiguity, and the Savage Axioms.* The Quarterly Journal of Economics, 1961. **75**(4): p. 643-669.

159. Kleindorfer, P.R., H.C. Kunreuther, and P.J.H. Schoenmaker, *Decision Sciences: An Integrative Approach*, in *Decision Sciences, An Integrative Perspective*. 1993, Cambridge University Press: New York.

160. Arrows, K.J., *A Difficulty in the Concept of Social Welfare.* Journal of Political Economy 1950. **58**(4): p. pp. 328–346.

161. Kahnemann, D. and A. Tversky, *Prospect Theory: An Analysis of Decision under Risk.* Econometrica, 1979. **47**: p. 313-327.

162. Kahnemann, D. and A. Tversky, *Advances in prospect theory: Cumulative representation of uncertainty.* Journal of Risk and Uncertainty, 1992. **5**: p. 297-323.

163. Köszegi, B. and M. Rabin, *A Model of Reference-Dependent Preferences.* The Quarterly Journal of Economics, 2006. **121**(4): p. 1133-1165.

164. Pawlak, Z., *Classification of Objects by Means of Attributes*. 1981, Institute for Computer Science, Polish Academy of Sciences Report 429.

165. Pawlak, Z., *Rough Sets - Basic Notions*. 1981, Institute for Computer Science, Polish Academy of Sciences Report 431.

166. Komorowski, J., L. Polkowski, and A. Skowron, *Rough Sets: a Tutorial*. 1998: Singapore.

167. Pawlak, Z., *Rough Sets*. Rough Sets and Data Mining, ed. N. Cecrone. 1997: Kluwer Academic Publishers. 3-8.

168. Pawlak, Z., *Some Issues on Rough Sets.* Transactions on Rough Sets I, 2004: p. pp. 1-58.

169. Dubois, D., et al. *An Information-Based Discussion of Vagueness*. in *IEEE International Conference on Fuzzy Systems*. 2001. Melbourne, Australia.

170. Sorensen, R. *Vagueness*. [Internet] August 21st, 2002 [cited March 28th, 2006]; Available from: http://plato.stanford.edu/entries/vagueness.

171. Keefe, R., *Theories of Vagueness*. 2000, Cambridge: Cambridge University Press.

172. Stefanowski, J. and D. Vanderpooten, *Induction of decision rules in classification and discovery-oriented perspectives.* International Journal of Intelligent Systems, 2001. **16**(1): p. 13-27.

173. Pawlak, Z., et al., *Rough Sets.* Communication of the ACM, 1995. **38**(11).

174. Grzymala-Busse, J., *Rule Induction*, in *Data Mining and Knowledge Discovery Handbook*, O. Maimon and L. Rokach, Editors. 2005, Springer US. p. 277-294.

175. Stefanowski, J., *On rough set based approaches to induction of decision rules*, in *Rough Sets in Knowledge Discovery*, L. Polkowski and A. Skowron, Editors. 1998, Physica Verlag: Heidelberg, Germany. p. 501 - 529.

176. Stefanowski, J., *On Rough Set Based Approaches to Induction of Decision Rules*. Rough Sets in Data Mining and Knowledge Discovery, ed. L. Polkowski and A. Skowron. Vol. 1. 1998: Physica-Verlag. p. 500-529.

177. Bazan, J., *A comparison of dynamic non-dynamic rough set methods for extracting laws from decision tables*, in *Rough Sets in Knowledge Discovery vol. 1 and 2.*, A. Skowron and L. Polkowski, Editors. 1998, Physica Verlag: Heidelberg. p. 321 - 365.

178. Grzymala-Busse, J., *Three Strategies to Rule Induction from Data with Numerical Attributes.* Transactions on Rough Sets II, 2004. **LNCS 3135**: p. 54–62.

179. Ziarko, W. and N. Shan. *Discovering Attribute Relationships, Dependencies and Rules by Using Rough Sets*. in *Hawaii International Conference on System Sciences (HICSS)*. 1995. Hawaii, USA: IEEE.

180. Cox, R., *Probability, frequency, and reasonable expectation.* American Journal of Physics, 1946. **14**: p. 1-13.

181. Salvati, D., *Credit Suisse - Phishing Attacks in Switzerland, Notes of Interviews with Menotti / Walder*. November 2007: Zurich.

182. Salvati, D., *Credit Suisse - Phishing Attacks in Switzerland, Notes of Interview with Gutermann*. September 2007: Zurich.

183. APWG. *Phishing Activity Trends Report for the Month of December, 2007*. [Internet] December 2007 [cited March 2008]; Available from: http://www.apwg.com/phishReportsArchive.html.

Appendix A: Threat Modeling

A.1 Jakobsson's Model

In Jakobsson's model [71] *nodes correspond to knowledge or access rights, and (directed) edges correspond to means of obtaining information or access rights from already possessed information or access rights; edges may also be associated with probabilities, costs, or other measures of the hardness of traversing a graph*. He argues that quantifying the effort of traversing a graph from a starting node to a target node can be used for quantification of the risks by economic analysis. Jakobsson puts the associated costs to traversing the graph in the forefront to compute success probabilities. If it is sufficiently high, then the attack is considered feasible and is potentially threatening.

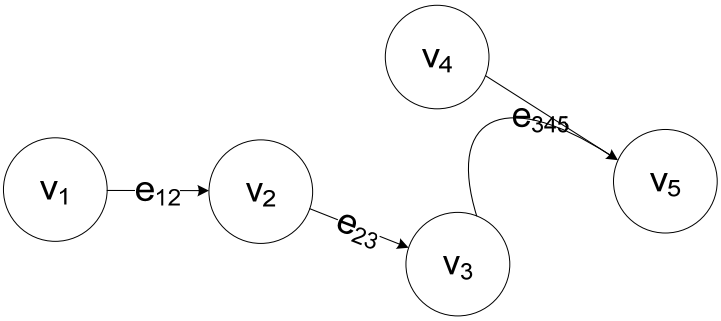

Figure A.1: Graph-based model by Jakobsson.

Figure A.1 shows a simplified graphical representation of the graph-based model. A detailed representation would label edges with the effort, probability, and other costs to execute threat actions.

A.2 Schneier's Model

The attacks propagate from top (root node) to bottom (leaf nodes) and a control flow is introduced by AND and OR gates. The basic approach assigns any Boolean or continuous value to the leaf nodes and then propagates them up the tree structure to obtain,

e.g., the cost for performing the attack described in the node or its probability of succeeding. The main goal is to increase the survivability of IS by documenting attack information in a structured and reusable form and, consequently, improve their designs. **Figure A.2** displays an example of such a tree structure.

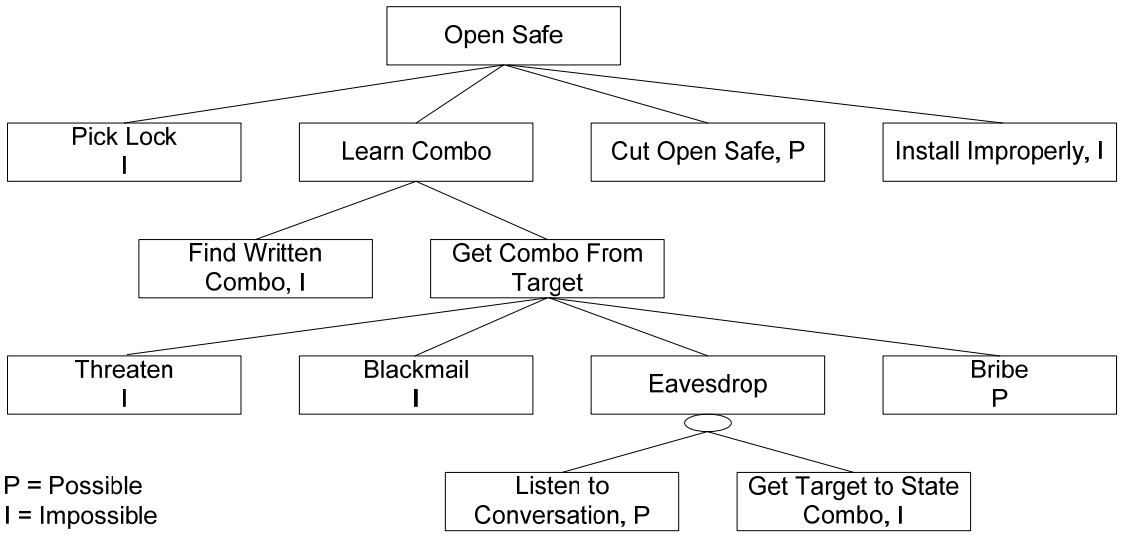

Figure A.2: Attack Nodes.

A.3 Bathtub Curve

Pan [83] assumes a distribution along the so-called bathtub curve. The bathtub curve shows a *TEST/DEBUG* area where software failures are non-stochastic and can be attributed to a common source. In this area, it is possible to identify the underlying errors and correct them. In the *USEFUL LIFE* area new errors are introduced by software upgrades. As before, many software failures can be backtracked to their source and consequently fixed. In the *OBSOLENCE* area stochastic software failures remain, which cannot be corrected (occur randomly) and which occur at a constant rate λ.

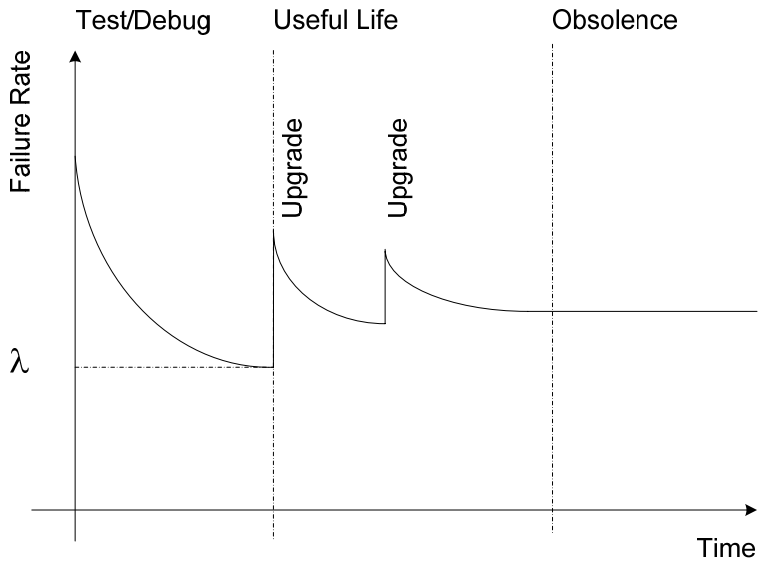

Figure A.3: Bathtub Curve.

Appendix B: Calculations in the Function Module

Appendix B reports the intermediary steps of Chapter 3.9 and continues the numbering of the equations.

B.1 Combining *e* and *u*

The following equations for *e*, *u*, and *z* show the starting situation:

$$e(t) = \begin{cases} \lambda \cdot e^{-\lambda \cdot t} & , \quad t \geq 0 \\ 0 & , \quad else \end{cases}, \tag{3.14}$$

$$u(t) = \begin{cases} \dfrac{1}{t_{\max} - t_{\min}} & , \quad t_{\min} \leq t \leq t_{\max} \\ 0 & , \quad else \end{cases}, \tag{3.22}$$

$$z(t) = \int_{-\infty}^{+\infty} e(t_t) \cdot u(t_t - t) dt_t \,. \tag{3.30}$$

Integrating *z(t)* over the range of interest where $T_t \geq T_r$, i.e. *0* ... ∞, (**3.31**) is obtained:

$$\Pr(T_t \geq T_r) = \Pr(Z \geq 0) = \int_0^{\infty} z(t) dt \,, \tag{3.31}$$

Conversely, in case the probability that the security mechanism prevails is of interest, then *z(t)* for $T_t < T_r$ is integrated over the range -∞ ... *0* and **(3.32)** is obtained:

$$\Pr(T_t < T_r) = \Pr(Z < 0) = \int_{-\infty}^{0} z(t) dt \,, \tag{3.32}$$

Inserting **(3.14)** and **(3.22)** into **(3.30), (3.33)** is obtained:

$$z(t)=\int_{-\infty}^{+\infty}\left[\lambda e^{-\lambda t_t}\cdot \mathrm{I}_{[t_t\geq 0]}(t_t)\right]\cdot\left[\frac{1}{t_{\max}-t_{\min}}\cdot \mathrm{I}_{[t_{\min}\leq t_t-t\leq t_{\max}]}(t_t-t)\right]\cdot dt_t\,, \tag{3.33}$$

$$\mathrm{I}_{[a\leq x\leq b]}(x)=\begin{cases}1 & a\leq x\leq b\\ 0 & else\end{cases}, \tag{3.34}$$

where $\mathrm{I}_{[a\leq x\leq b]}(x)$ is the indicator function with control variable x and value *1* in an interval $a\leq x\leq b$ and value 0 else

x is the control variable.

Factoring out $\frac{\lambda}{t_{\max}-t_{\min}}$ of **(3.33)** yields:

$$z(t)=\frac{\lambda}{t_{\max}-t_{\min}}\cdot\int_{-\infty}^{+\infty}\left[e^{-\lambda t_t}\cdot \mathrm{I}_{[t_t\geq 0]}(t_t)\right]\cdot\left[\mathrm{I}_{[t_{\min}\leq t_t-t\leq t_{\max}]}(t_t-t)\right]\cdot dt_t\,. \tag{3.35}$$

Isolating t_t in $\mathrm{I}_{[t_{\min}\leq t_t-t\leq t_{\max}]}(t_t-t)$ yields:

$$z(t)=\frac{\lambda}{t_{\max}-t_{\min}}\cdot\int_{-\infty}^{+\infty}\left[e^{-\lambda t_t}\cdot \mathrm{I}_{[t_t\geq 0]}(t_t)\right]\cdot\left[\mathrm{I}_{[t_{\min}+t\leq t_t\leq t_{\max}+t]}(t_t)\right]\cdot dt_t\,. \tag{3.36}$$

Applying case differentiation for $t_{\max}+t<0$, $-t_{\max}\leq t\leq -t_{\min}$, $t_{\min}+t>0$ yields:

$$z(t)=\begin{cases}0 & t_{\max}+t<0\\ \frac{\lambda}{t_{\max}-t_{\min}}\cdot\int_{-\infty}^{+\infty}e^{-\lambda t_t}\cdot \mathrm{I}_{[0\leq t_t\leq t_{\max}+t]}(t_t)\cdot dt_t & -t_{\max}\leq t\leq -t_{\min}\\ \frac{\lambda}{t_{\max}-t_{\min}}\cdot\int_{-\infty}^{+\infty}e^{-\lambda t_t}\cdot \mathrm{I}_{[t_{\min}+t\leq t_t\leq t_{\max}+t]}(t_t)\cdot dt_t & t_{\min}+t>0\end{cases}. \tag{3.37}$$

Inserting the integration boundaries yields:

$$z(t)=\begin{cases} 0 & t_{\max}+t<0 \\ \dfrac{\lambda}{t_{\max}-t_{\min}}\cdot\int_0^{t_{\max}+t} e^{-\lambda t_t}\cdot dt_t & -t_{\max}\le t\le -t_{\min}\cdot \\ \dfrac{\lambda}{t_{\max}-t_{\min}}\cdot\int_{t_{\min}+t}^{t_{\max}+t} e^{-\lambda t_t}\cdot dt_t & t_{\min}+t>0 \end{cases} \tag{3.38}$$

Integration yields:

$$z(t)=\begin{cases} 0 & t_{\max}+t<0 \\ \dfrac{\lambda}{t_{\max}-t_{\min}}\cdot\left[-\dfrac{1}{\lambda}\cdot e^{-\lambda t_t}\right]_0^{t_{\max}+t} & -t_{\max}\le t\le -t_{\min}\cdot \\ \dfrac{\lambda}{t_{\max}-t_{\min}}\cdot\left[-\dfrac{1}{\lambda}e^{-\lambda t_t}\right]_{t_{\min}+t}^{t_{\max}+t} & t_{\min}+t>0 \end{cases} \tag{3.39}$$

Evaluation of the boundaries yields:

$$z(t)=\begin{cases} 0 & t_{\max}+t<0 \\ \dfrac{\lambda}{t_{\max}-t_{\min}}\cdot\left[-\dfrac{1}{\lambda}\cdot e^{-\lambda(t_{\max}+t)}+\dfrac{1}{\lambda}\cdot e^{0}\right] & -t_{\max}\le t\le -t_{\min}\cdot \\ \dfrac{\lambda}{t_{\max}-t_{\min}}\cdot\left[-\dfrac{1}{\lambda}e^{-\lambda(t_{\max}+t)}+\dfrac{1}{\lambda}e^{-\lambda(t_{\min}+t)}\right] & t_{\min}+t>0 \end{cases} \tag{3.40}$$

Factoring out $\frac{1}{\lambda}$ and changing the algebraic sign $\cdot(-1)$ yields:

$$z(t)=\begin{cases} 0 & t_{\max}+t<0 \\ \dfrac{1}{t_{\min}-t_{\max}}\cdot\left[e^{-\lambda(t_{\max}+t)}-1\right] & -t_{\max}\le t\le -t_{\min}\cdot \\ \dfrac{1}{t_{\min}-t_{\max}}\cdot\left[e^{-\lambda(t_{\max}+t)}-e^{-\lambda(t_{\min}+t)}\right] & t_{\min}+t>0 \end{cases} \tag{3.41}$$

B.2 Calculating the Success Probability

The probability is obtained by integrating **(3.41)** from $-\infty$ to the point x of interest, i.e.:

$$F_Z(x) = \Pr(Z \le x) = \int_{-\infty}^{x} z(t) \cdot dt \,. \tag{3.42}$$

For $t_{\max} + x < 0$ of the above case differentiation results in:

$$\Pr(Z \le x) = 0 \,. \tag{3.43}$$

Taking case differentiation into account for $-t_{\max} \le x \le -t_{\min}$:

$$\Pr(Z \le x) = \frac{1}{t_{\min} - t_{\max}} \cdot \int_{-t_{\max}}^{x} \left[e^{-\lambda(t_{\max}+t)} - 1 \right] \cdot dt \,, \tag{3.44}$$

$$= \frac{1}{t_{\min} - t_{\max}} \cdot \left[-\frac{1}{\lambda} e^{-\lambda(t_{\max}+t)} - t \right]_{-t_{\max}}^{x} , \tag{3.45}$$

$$= \frac{1}{t_{\min} - t_{\max}} \cdot \left[\left[-\frac{1}{\lambda} e^{-\lambda(t_{\max}+x)} - x \right] - \left[-\frac{1}{\lambda} e^{-\lambda(t_{\max}-t_{\max})} + t_{\max} \right] \right], \tag{3.46}$$

$$= \frac{1}{t_{\min} - t_{\max}} \cdot \left[\left[-\frac{1}{\lambda} e^{-\lambda(t_{\max}+x)} - x \right] - \left[-\frac{1}{\lambda} + t_{\max} \right] \right], \tag{3.47}$$

$$= \frac{1}{t_{\min} - t_{\max}} \cdot \left[-\frac{1}{\lambda} e^{-\lambda(t_{\max}+x)} - x + \frac{1}{\lambda} - t_{\max} \right], \tag{3.48}$$

$$= \frac{1}{t_{\min} - t_{\max}} \cdot \left[\frac{1}{\lambda} \cdot \left(1 - e^{-\lambda(t_{\max}+x)} \right) - x - t_{\max} \right], \tag{3.49}$$

$$= \frac{\frac{1}{\lambda} \cdot \left(1 - e^{-\lambda(t_{\max}+x)} \right) - x - t_{\max}}{t_{\min} - t_{\max}} . \tag{3.50}$$

Taking case differentiation into account for $t_{\min} + x > 0$:

$$\Pr(Z \le x) = \frac{1}{t_{\min} - t_{\max}} \cdot \left[\int_{-t_{\max}}^{-t_{\min}} \left[e^{-\lambda(t_{\max}+t)} - 1 \right] \cdot dt + \int_{-t_{\min}}^{x} \left[e^{-\lambda(t_{\max}+t)} - e^{-\lambda(t_{\min}+t)} \right] \cdot dt \right], \tag{3.51}$$

$$= \frac{1}{t_{\min} - t_{\max}} \cdot \left[\left[-\frac{1}{\lambda} e^{-\lambda(t_{\max}+t)} - t \right]_{-t_{\max}}^{-t_{\min}} + \left[-\frac{1}{\lambda} e^{-\lambda(t_{\max}+t)} + \frac{1}{\lambda} e^{-\lambda(t_{\min}+t)} \right]_{-t_{\min}}^{x} \right], \tag{3.52}$$

$$= \frac{1}{t_{\min} - t_{\max}} \cdot \left[\left[-\frac{1}{\lambda} e^{-\lambda(t_{\max}-t_{\min})} + t_{\min} \right] - \left[-\frac{1}{\lambda} e^{-\lambda(t_{\max}-t_{\max})} + t_{\max} \right] + \left[\left[-\frac{1}{\lambda} e^{-\lambda(t_{\max}+x)} + \frac{1}{\lambda} e^{-\lambda(t_{\min}+x)} \right] - \left[-\frac{1}{\lambda} e^{-\lambda(t_{\max}-t_{\min})} + \frac{1}{\lambda} e^{-\lambda(t_{\min}-t_{\min})} \right] \right] \right], \tag{3.53}$$

$$= \frac{1}{t_{\min} - t_{\max}} \cdot \left[\left[-\frac{1}{\lambda} e^{-\lambda(t_{\max}-t_{\min})} + t_{\min} \right] + \frac{1}{\lambda} - t_{\max} + \left[\left[-\frac{1}{\lambda} e^{-\lambda(t_{\max}+x)} + \frac{1}{\lambda} e^{-\lambda(t_{\min}+x)} \right] + \frac{1}{\lambda} e^{-\lambda(t_{\max}-t_{\min})} - \frac{1}{\lambda} \right] \right], \tag{3.54}$$

$$= \frac{\left[\left[\frac{1}{\lambda} \left(1 - e^{-\lambda(t_{\max}-t_{\min})} \right) + t_{\min} - t_{\max} \right] + \left[\frac{1}{\lambda} \cdot \left(-e^{-\lambda(t_{\max}+x)} + e^{-\lambda(t_{\min}+x)} + e^{-\lambda(t_{\max}-t_{\min})} - 1 \right) \right] \right]}{t_{\min} - t_{\max}}, \tag{3.55}$$

$$\frac{t_{\min} - t_{\max} + \frac{1}{\lambda} \cdot \left(1 - e^{-\lambda(t_{\max}-t_{\min})} - e^{-\lambda(t_{\max}+x)} + e^{-\lambda(t_{\min}+x)} + e^{-\lambda(t_{\max}-t_{\min})} - 1 \right)}{t_{\min} - t_{\max}}, \tag{3.56}$$

$$\frac{t_{\min} - t_{\max} + \frac{1}{\lambda} \cdot \left(-e^{-\lambda(t_{\max}+x)} + e^{-\lambda(t_{\min}+x)} \right)}{t_{\min} - t_{\max}}. \tag{3.57}$$

Summarizing,

$$\Pr(Z \le x) = \begin{cases} 0 & t_{\max} + x < 0 \\ = \dfrac{\frac{1}{\lambda} \cdot \left(1 - e^{-\lambda(t_{\max}+x)}\right) - x - t_{\max}}{t_{\min} - t_{\max}} & -t_{\max} \le x \le -t_{\min} \\ \dfrac{t_{\min} - t_{\max} + \frac{1}{\lambda} \cdot \left(-e^{-\lambda(t_{\max}+x)} + e^{-\lambda(t_{\min}+x)}\right)}{t_{\min} - t_{\max}} & t_{\min} + x > 0 \end{cases} \quad \textbf{(3.58)}$$

Appendix C: Introduction to Rough Sets Theory (RST)

In 1982, Zdzislaw Pawlak [111, 164, 165] gave birth to the Theory of Rough Sets (RST) which yielded a well-organized research community (www.roughsets.org). RST has been applied to, e.g., medicine, economics, decision analysis and social sciences (for a comprehensive list see Komorowski et al [166]).

RST operates on 2-dimensional data tables where the rows represent *OBJECTS* that are related among each other and the columns contain *ATTRIBUTES* (set *A*) describing the objects. Two types of attributes are distinguished: *DECISION ATTRIBUTES* (set *B*), which depend on *CONDITIONAL ATTRIBUTES*. In RST, attribute values are allowed to be incomplete, inconsistent or vague.

In essence, RST discerns objects. Its main concept is an indiscernability relation *IND(.)* applied to a set of attributes *B* describing two objects. *IND(B)* verifies whether the two objects are equivalent (indiscernable) by their attribute values and produces *EQUIVALENCE CLASSES,* which are called *ELEMENTARY SETS* of *B*. **Figure C.1** exemplifies the above.

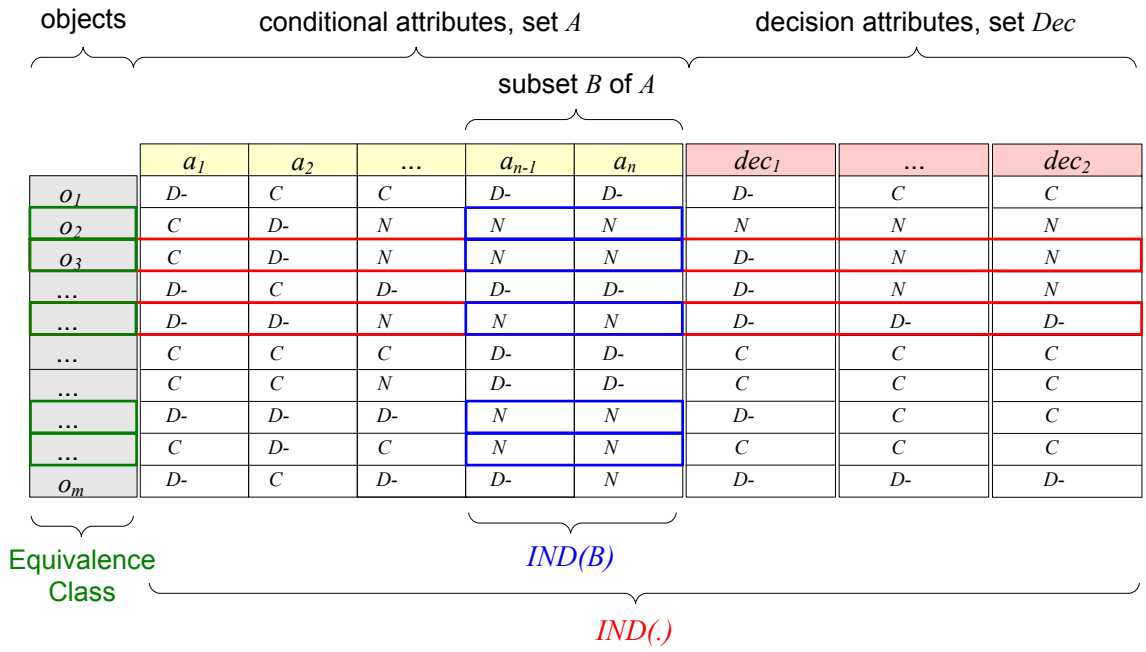

Figure C.1: 2-Dimensional Data Table in Rough Sets.

In the above data table, the objects $o_1, o_2, ..., o_m$ are described by conditional attributes $a_1, a_2, ..., a_n$. Two objects, e.g., o_2 and o_3, are indiscernable by a subset B, e.g., a_{n-1} and a_n, if their attributes display the same values.

According to Pawlak [167] the indiscernability relation is at the basis of:

- describing a set of objects by their attributes
- recognizing (full or partial) dependencies between and among attributes
- reducing the number of attributes (e.g., to a minimum number)
- analyzing the significance of attributes
- inducing *IF ... THEN* rules from the data table.

In contrast to Classical Set Theory where objects are classifiable as strictly belonging to a set X or not, RST approximates the set X by another set $X`$. Imagine a target set $X$ of objects defined by decision attributes. Further, imagine a set $X`$ defined by conditional attributes. In principle, the idea is to determine how well $X`$ fits the target set $X$ by the number of overlapping objects in the two sets. Conditional attributes, which create an $X`$ such that the objects of $X`$:

- *STRICTLY* belong to X are said to form a lower approximation of X, i.e. all the objects in $X`$ also belong to $X$; conversely, objects of $X$ do not necessarily belong to $X`$
- *POSSIBLY* belong to X are said to form an upper approximation, i.e. some objects of $X`$ belong to X but others do not.

If equivalence classes (created by arbitrarily chosen conditional attributes) are thought of as the smallest unit for describing a *UNIVERSE*, then all of its objects are contained in the union of all of them. Consequently, a target set X (see area evidenced in the upper left corner of **Figure C.2**) encompasses a subset of all objects of a universe. In the upper right of **Figure C.2**, a *LOWER APPROXIMATION* of X is displayed. It is formed by equivalence classes containing all objects, which strictly belong to X. Conversely, the *UPPER APPROXIMATION* of X (lower left hand side) indeed displays all objects that strictly belong to X but in addition it displays equivalence classes containing objects of which some belong to X but others do not.

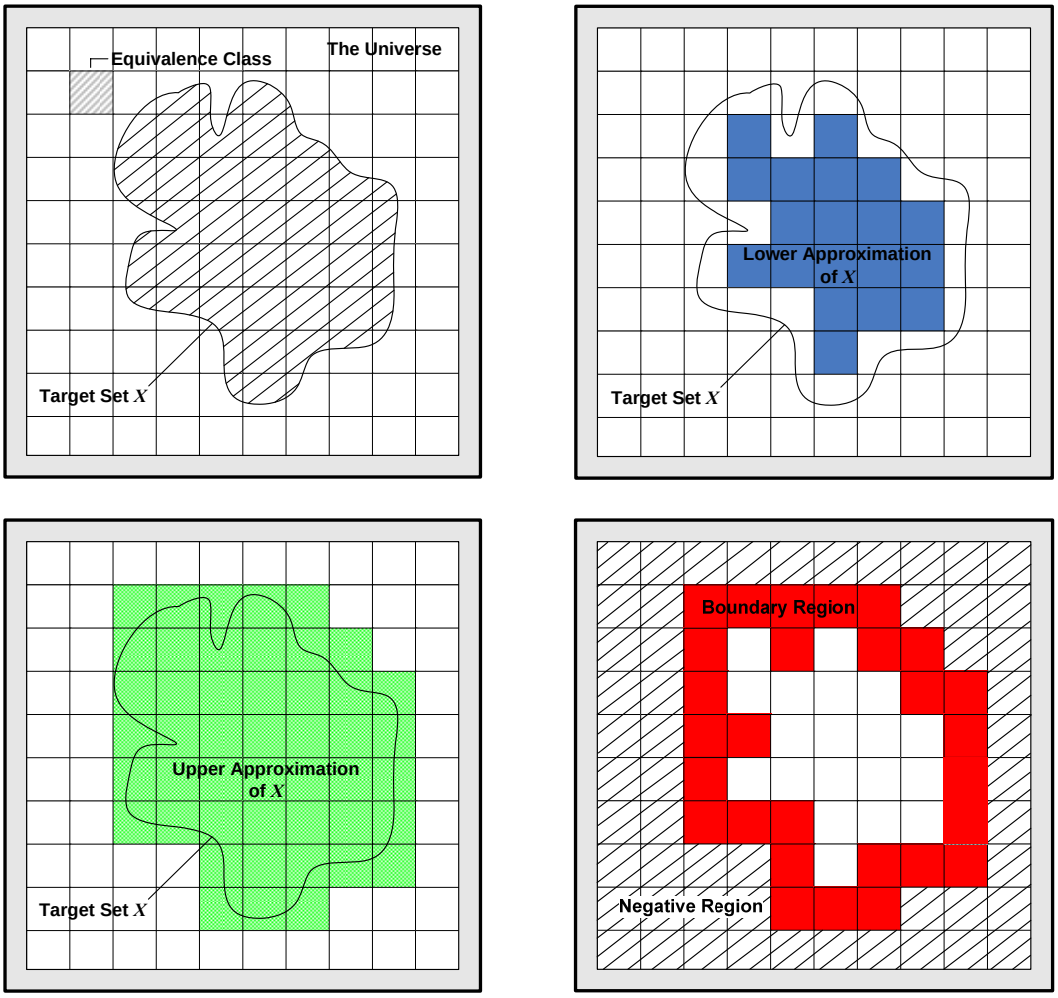

Figure C.2: Graphic Interpretation of Rough Sets.

The difference between the upper and the lower approximation is called a *BOUNDARY REGION* (see lower right hand side of **Figure C.2**). This is the very notion of RST: a non-empty boundary region indicates that the sets are rough and an empty boundary region indicates that the sets are crisp (as in Classical Set Theory). Pawlak argues that this notion of Rough Sets is useful when attributes are vague [168]. Keefe defines vagueness as *admitting borderline cases*[86]*, lacking (or at least apparently lacking) sharp boundaries*[87] *and being susceptible to the Sorites Paradox*[88] [171]. Finally, the region outside of the boundary region is called the *NEGATIVE REGION* while inside lies the *POSITIVE REGION*.

[86] Borderline height: people may be not clearly tall and not clearly not tall.

[87] Sharp/vague boundaries: vague boundaries indicate the lack of well-defined extensions, see Dubois [169].

[88] The Sorites Paradox: (1) Base step: A one day year old human being is a child. (2) Induction step: If an *n* day old human being is a child, then that human being is also a

The process of modeling data in RST is two-fold: in the *PRE-PROCESSING STAGE*, the table is searched in order:

- to select attributes for synthesizing approximations of a set X (feature selection)
- to extract new features by combining existing attributes. Feature extraction usually synthesizes a target X more accurately than feature selection alone.

The aim of the pre-processing stage is to provide an extensional description of the data table. Consequently, tasks related to the treatment of attributes are at the centre of attention, i.e.:

- the discretization of attributes with real or symbolic values
- the elimination of "noise", or
- the treatment of missing data.

The aim of *RULE EXTRACTION* is to provide an intensional description of the data table by Boolean operations while the extensional form of a data table is represented in disjunctive normal form (DNF) of propositional logic. The goal is to find a set of logical implications (if-then rules) that characterize the data table. An example of a rule is *if* (conditional) $attribute_1 = x$ and (conditional) $attribute_2 = y$ *then* (decision) $dec = z$.

In the rule extraction stage, Stefanowski distinguishes two types of induction: a *CLASSIFICATION-ORIENTED INDUCTION* and a *DISCOVERY-ORIENTED INDUCTION* [172]. A classification-oriented induction is used with the aim to classify future examples of data from a set of learning examples and is usually evaluated by one criteria: the classification or predictive accuracy of the rules.

The aim of discovery-oriented induction is to find information patterns and regularities along with exceptions and anomalies. Such patterns support the clarification of dependencies between attributes (i.e. column-wise dependencies among conditional attributes and row-wise dependencies among decision attributes). Discovery-oriented rule induction is evaluated by individual rules; thus, ending up with multiple evaluation criteria.

In this work, the focus was put on the pre-processing stage. However, in **Appendix D** information related to the rule extraction stage is provided, in particular, for discovery-oriented induction. An easy first reader on RST is found in Communication of the ACM [173]. Further reading is provided by Pawlak [168], Komorowski et al [166] or Grzymala-Busse [174]. RST software is available from ProSoft [117].

---

child when it is $n + 1$ days old. (3) Conclusion: Therefore, a 36,500 days old human being is a child. Clearly, this conclusion must be false if it is considered that 36,500 days are equal to 100 years! (example according to Sorensen [170])

Appendix D: Applying RST Rule Extraction to Security Information

Chapter 4.8 evidences the security processes *policies*, *technical* and *resources*. The threat *brute force attack* is not considered as it is assumed that the security processes do not affect it. In contrast, the security mechanism *password quality is* considered and shows three different values: *C, D-* and *N*. The following data table is used for the remainder of this appendix (**Figure D.1**):

set U	reduced set A			reduced set Dec
	A_1	A_3	A_4	Dec
	Adoption of Policies	Technical Assessments	Resource Allocation	Password Quality
Frankfurt	*D-*	*C*	*D-*	*D-*
Madrid	*C*	*N*	*N*	*N*
Paris	*C*	*N*	*N*	*D-*
Milan	*D-*	*D-*	*D-*	*D-*
Guernsey	*D-*	*N*	*N*	*D-*
Luxembourg	*C*	*C*	*D-*	*C*
New York	*C*	*N*	*D-*	*C*
Nassau	*D-*	*D-*	*N*	*D-*
Singapore	*C*	*C*	*N*	*C*
Sydney	*D-*	*D-*	*D-*	*D-*

Figure D.1: Reduced Data Table.

From this new data table, aim is to extract a set of *IF-THEN RULES*. These if-then rules convey additional information as they collect security processes, which "point" at specific values of a security mechanism. The if-then rules are obtained by:

- transforming the reduced data table into Disjunctive Normal Form (DNF) of propositional logic (Chapter D.1 in this appendix) and
- by applying a "suitable" algorithm for rule generation (Chapter D.2 in this appendix).

D.1 Security Processes and Rule Induction

Rule extraction in RST is based on *INDUCTIVE REASONING*[89]. Its application yields an intensional characterization of the security information in the data table. For a set of security processes $A = \{A_1, A_2 \ldots A_N\}$, and a set of security mechanisms $Dec = \{Dec_1, Dec_2 \ldots Dec_M\}$ these rules have the form:

$$(A_1 = a) \wedge (A_2 = b) \wedge \ldots (A_N = c) \rightarrow (Dec_1 = d_1), \ldots, (Dec_M = d_M), \quad \textbf{(D.1)}$$

which can also be expressed as:

$$A_1^{a_i} \bullet A_2^{a_i} \bullet \ldots \bullet A_N^{a_N} \rightarrow Dec_1^{d_1}, \ldots, Dec_M^{d_M}, \quad \textbf{(D.2)}$$

where $A_i^{a_i}$ is the *i*-th security process with value a_i, $1 \le i \le N$

$D_k^{d_k}$ is the *k*-th security mechanism with value d_k, $1 \le k \le M$,

Many rule induction algorithms are available depending on the aim, which is driven by three basic purposes. According to Stefanowski [175] these are:

- obtaining the smallest number of rules to describe the entire data table (minimum set approach)
- obtaining the greatest number of rules, which can possibly be generated from the available data table (exhaustive set approach), or
- obtaining a number of rules that satisfy requirements defined by the user/analyst (satisfactory set approach).

Rule induction in RST belongs to the *SUPERVISED LEARNING TECHNIQUES*[90] and is also one fundamental tool for data mining. In RST, the security mechanism usually classifies the branches (via the security processes). The rules may be inconsistent with respect to unseen branches and inconsistencies may be considered acceptable if they can be resolved

---

[89] Inductive reasoning infers a conclusion from the premises of an argument with some probability but does not ensure it. In contrast, deductive reasoning leads to a true conclusion in all cases, given a true premise that is.

[90] Two techniques of machine learning are distinguished among others:

- Supervised learning creates a function from input and output data (training data), which has previously been classified by an expert or gained through measurement
- Unsupervised learning is distinguished from supervised learning by the fact that there is no *a priori* output data. Unsupervised learning treats input objects as random variables for which a density model is built.

by some method (e.g., by voting). The number of branches matching the rules is called the support.

In our ongoing governance example the DNF is applied to the reduced data table and the following is obtained:

$$(A_1^{D\text{-}} A_3^{C} A_4^{D\text{-}} Dec^{D\text{-}}) \vee (A_1^{C} A_3^{N} A_4^{N} Dec^{N}) \vee$$
$$(A_1^{C} A_3^{N} A_4^{N} Dec^{D\text{-}}) \vee (A_1^{D\text{-}} A_3^{D\text{-}} A_4^{D\text{-}} Dec^{D\text{-}}) \vee$$
$$(A_1^{D\text{-}} A_3^{N} A_4^{N} Dec^{D\text{-}}) \vee (A_1^{C} A_3^{C} A_4^{D\text{-}} Dec^{C}) \vee$$
$$(A_1^{C} A_3^{N} A_4^{D\text{-}} Dec^{C}) \vee (A_1^{D\text{-}} A_3^{D\text{-}} A_4^{N\text{-}} Dec^{D\text{-}}) \vee$$
$$(A_1^{C} A_3^{C} A_4^{N} Dec^{C}) \vee (A_1^{D\text{-}} A_3^{D\text{-}} A_4^{D\text{-}} Dec^{D\text{-}})$$

Table C.1: Disjunctive Normal Form (DNF).

D.2 Security Mechanisms and Rule Induction

A branch *o* is covered by a rule *r* if every security process of the branch is satisfied by the corresponding value for *o*. The concept *C* defined by the right hand side of rule *r* is indicated by *r*. A rule set *R* is:

- complete, if for every branch *o* in *C* there is a rule *r* in *R* such that *r* covers *o*
- consistent, if for every branch *o* in *C*, o is a member of the concept *C* indicated by *r*.

The most frequent task of rule induction is to induce a rule set *R* that is consistent and complete. There are many types of rules, e.g., strong rules (where the rules cover many cases) and associative rules (the left and right side of a rule both contain security processes). For a systematic overview on induction algorithms see Stefanowski and Bazan [176, 177]. For a discussion refer to Grzymala-Busse [178].

The left hand side of the if-then rules (the security processes) corresponds to a target, which the branches are mandated to achieve. Achieving this target yields a *PROBABILITY* that the right hand side (the security mechanisms) is satisfied. However, as the number of branches is usually small (in the ongoing example a total of ten branches), the term probability must be put into perspective. Therefore, the primary interest is to analyze the dependency, dispensability and significance among security processes and the author prefers using the term *INFLUENCE* rather than probability.

For the sake of completeness, the rule induction algorithm by Ziarko and Shan [179] is presented. Ziarko and Shan form one distinct decision matrix for each value d_k of the security mechanism Dec_k **(Table D.2)**. This matrix orders all branches with $Dec_k = d_k$ in the most left column and all objects with $Dec_k \neq d_k$ in the topmost row. The intersecting cells list the security processes of a branch (row), which differs from another branch (column). They list all the differences between branches having *password quality = D-* and branches having *password quality ≠ D-*.

For example, for the *password quality = D-*, the decision matrix takes the form:

	Madrid	Luxembourg	New York	Singapore
Frankfurt	$A_1^{D-}, A_3^{C}, A_4^{D-}$	A_1^{D-}	A_1^{D-}, A_3^{C}	A_1^{D-}, A_4^{D-}
Paris	—	A_3^{N}, A_4^{N}	A_4^{N}	A_3^{N}
Milan	$A_1^{D-}, A_3^{D-}, A_4^{D-}$	A_1^{D-}, A_3^{D-}	A_1^{D-}, A_3^{D-}	$A_1^{D-}, A_3^{D-}, A_4^{D-}$
Guernsey	A_1^{D-}	$A_1^{D-}, A_3^{D-}, A_4^{D-}$	A_1^{D-}, A_4^{N}	A_1^{D-}, A_3^{N}
Nassau	A_1^{D-}, A_3^{D-}	$A_1^{D-}, A_3^{D-}, A_4^{N}$	$A_1^{D-}, A_3^{D-}, A_4^{N}$	A_1^{D-}, A_3^{D-}
Sydney	$A_1^{D-}, A_3^{D-}, A_4^{D-}$	A_1^{D-}, A_3^{D-}	A_1^{D-}, A_3^{D-}	$A_1^{D-}, A_3^{D-}, A_4^{D-}$

Table D.2: Decision Matrix for *password quality = D-* (Ziarko and Shan).

Next, for the decision matrix, the items[91] within each cell are aggregated disjunctively and the individual cells are aggregated conjunctively, i.e.:

Frankfurt	$(A_1^{D-} \vee A_3^{C} \vee A_4^{D-}) \wedge A_1^{D-} \wedge (A_1^{D-} \vee A_3^{C}) \wedge (A_1^{D-} \vee A_4^{D-})$
Paris	$(A_3^{N} \vee A_4^{N}) \wedge A_4^{N} \wedge A_3^{N})$
Milan	$(A_1^{D-} \vee A_3^{D-} \vee A_4^{D-}) \wedge (A_1^{D-} \vee A_3^{D-}) \wedge (A_1^{D-} \vee A_3^{D-}) \wedge (A_1^{D-} \vee A_3^{D-} \vee A_4^{D-})$
Guernsey	$A_1^{D-} \wedge (A_1^{D-} \vee A_3^{N} \vee A_4^{N}) \wedge (A_1^{D-} \vee A_4^{N}) \wedge (A_1^{D-} \vee A_3^{N})$
Nassau	$(A_1^{D-} \vee A_3^{D-}) \wedge (A_1^{D-} \vee A_3^{D-} \vee A_4^{N}) \wedge (A_1^{D-} \vee A_3^{D-} \vee A_4^{N}) \wedge (A_1^{D-} \vee A_3^{D-})$
Sydney	$(A_1^{D-} \vee A_3^{D-} \vee A_4^{D-}) \wedge (A_1^{D-} \vee A_3^{D-}) \wedge (A_1^{D-} \vee A_3^{D-}) \wedge (A_1^{D-} \vee A_3^{D-} \vee A_4^{D-})$

Table D.3: Disjunctive and Conjunctive Aggregation.

Simplifying the above table yields the following rules for *password quality = D-*:

1. *if A_1^{D-} or A_3^{D-} then D-*

[91] The underlying assumption is that each item can be described by the others, by their presence or absence.

2. *if* (A_3^N *and* A_4^N) *then D-*.

The approach by Ziarko and Shan requires a decision matrix for each value of *password quality*. Analogously, for *password quality = N* the following is obtained:

1. *if* (A_1^C *and* A_3^N *and* A_4^N) *then N*.

Analogously, for *password quality = C* the following is obtained:

1. *if* (A_1^C *and* A_3^C) *then C*
2. *if* (A_1^C *and* A_4^{D-}) *then C*
3. *if* (A_3^C *and* A_4^N) *then C*
4. *if* (A_3^N *and* A_4^{D-}) *then C*.

The above rules show no inconsistencies. Refer to Ziarko and Shan [179] for resolving inconsistencies.

Appendix E: Howard's Decision Model

E.1 Decision Situation

Suppose an *INDIVIDUAL* is at a point in time, which separates the past from the future. A decision needs to be made, i.e. to choose among *ALTERNATIVES*. Each alternative offers *PROSPECTS* whose occurrence the individual is uncertain about. This uncertainty is expressed by degrees of belief, i.e. by subjective probabilities of an event occurring (based on current knowledge and experience).

The decision maker then chooses one alternative. Eventually, one prospect of the chosen alternative is realized by an event. Prospects, should they become real, are of *VALUE* to the decision maker. This value is usually expressed by a commodity such as money or time. However, there is no need for prospect value to be commensurate (expressed in the same dimension).

Looking from the present into the future, each alternative can be interpreted as a *LOTTERY* (or *DEAL*) the decision maker holds. This lottery (ticket) offers a chance to the decision maker to either make money or loose money. The value the lottery has for the decision maker is expressed in terms of its *CERTAIN EQUIVALENT*. The certain equivalent signifies the minimum (monetary) value the lottery ticket has for a decision maker. If the decision maker was offered any value higher than the certain equivalent, the lottery ticked would be exchanged for that value, for any lower value the ticket holder would keep it and for the same value the ticket holder would be indifferent between choosing the certain equivalent and the lottery ticket.

Typically, the certain equivalent is valued less than the expected value[92] of the lottery ticket, because the decision maker is *RISK AVERSE*, i.e. the decision maker is willing to forego a *RISK PREMIUM* in order to avoid risk. Just how high the certain equivalent of a lottery is, depends on the *RISK PREFERENCES* of the decision maker.

[92] The expected value denotes the individual prospects in relation to their probabilities of occurrence.

In essence, deciding among alternatives becomes a problem of assessing their certain equivalents and then selecting the alternative, which relates to the highest certain equivalent. Finally, accepting an alternative means to irreversibly allocate limited resources. **Figure E.1** displays the above by a so-called decision tree.

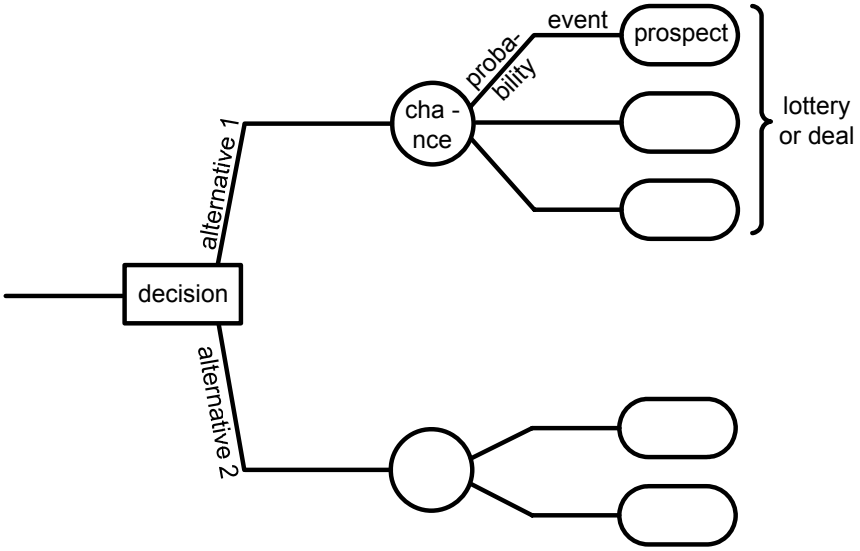

Figure E.1: Decision Tree.

Figure E.1 shows a decision tree with two alternatives. The (red) rhombus symbolizes the pending decision for which there are two alternatives. The (blue) boxes on the right hand side of the picture show the prospects, which are subject to chance (circle). If the probabilities related to the prospects are single values, then a decision tree is suited to represent them. Conversely, if the probabilities are (continuous) random variables, then a probability distribution function or a cumulative probability distribution are used.

In Howard's Model, probabilities of events are assigned by:

$$\{B|\&\} = p\ , \qquad \textbf{(E.1)}$$

where $\{.\}$ the braces indicate a probability assignment

B is an event yielding a prospect

$|\&$ indicates the current knowledge and experience of the decision

maker

p is the probability (subjective degree of belief).

Decision trees may be more complex than are displayed in **Figure E.1**. For example, they may be composed of other decision trees (so-called compound trees).

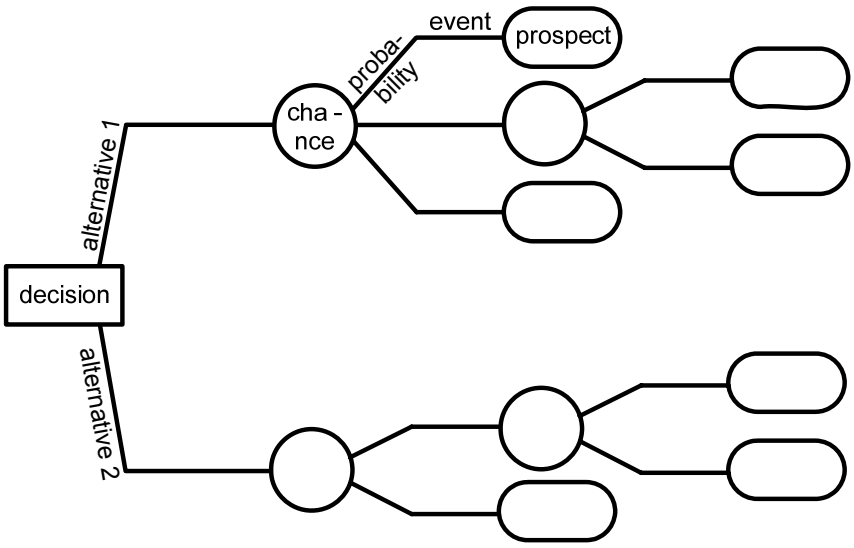

Figure E.2: Compound Decision Tree.

In (compound) decision trees, probabilities are calculated by the Sum and the Product Rules, which were brought forward by Richard Cox in 1946 [180]. The Sum Rule states that specifying one's belief about a prospect occurring implies having also specified one's belief about the prospect not to occur, i.e.:

$$\{B|\&\} + \{B`|\&\} = 1 \,, \qquad \textbf{(E.2)}$$

where $\{.\}$ the braces indicate a probability assignment

B is the event denoting the occurrence of a prospect

$B`$ is the opposite event, which denotes the prospect not occurring

$|\&$ indicates the current knowledge and experience .

For independent events, the Product Rule states that if we first specify our belief that event B occurs, and then specify our belief that event A occurs (given that event B has occurred), then our belief has implicitly been specified that the events A and B occur:

$$\{A, B|\&\} = \{A|\&\} \cdot \{B|\&\} , \qquad \textbf{(E.3)}$$

where $\{.\}$ the braces indicate a conditional probability

A, B indicate the events of prospect *B* occurring followed by prospect *A*

$|\&$ indicates the current knowledge and experience .

In a decision tree, the expected value of an alternative with discrete prospect values is obtained by:

$$E(X) = \sum_{x \in \Omega} x m(x) , \qquad \textbf{(E.4)}$$

where $E(X)$ is the expected value

X is the discrete random variable

$m(x)$ is the distribution function

$\boldsymbol{\Omega}$ is the sample space.

Once a coherent decision tree can be drawn from a finite number of (relevant) alternatives and their prospects, the decision problem passes the so called *CLARITY TEST*. In this case, events in the decision tree are:

- mutually excluding, i.e. if an event occurs then another can not
- collectively exhausting, i.e. summing up the probabilities among all prospects in the same alternative yields *1* ($\sum p_i = 1$).

Next, Howard introduces Five Rules of Actional Thought for processing the prospects.

E.2 Five Axioms of Risk Preference

1. *The ORDERABILITY AXIOM* arranges the prospects of the decision tree in a list of descending order from best to worst. This implies that a decision maker must be able to order the prospects according to individual preferences, i.e. $A > B$, $A \sim B$ or $A < B$ and that these preferences are transitive, i.e. if $A > B$ and $B > C$ then $A > C$ or if $A \sim B$ then $B \sim A$). Note that the prospects may be valued in multiple dimensions or differing dimensions for which the orderability axiom applies as well.

2. Let $A > B > C$. Further, let p_p a preference probability and p an event probability. The *CONTINUITY AXIOM* allows for a given prospect B to be compared with the lottery comprising prospects A and C in the following way. The axiom states that a preference probability p_p exists with which A should occur (conversely, a preference probability $(1 - p_p)$ with which C should occur), so that a decision maker is indifferent between receiving B with certainty and A with probability p_p (C with $(1 - p_p)$ respectively) (see upper graph of **Figure E.3**). B is called the certain equivalent and the equivalence rule is about personal preferences. Alternatively, the event probabilities p and $(1 - p)$ of A and C are given and the decision maker chooses the value of the certain equivalent B (lower graph of **Figure E.3**).

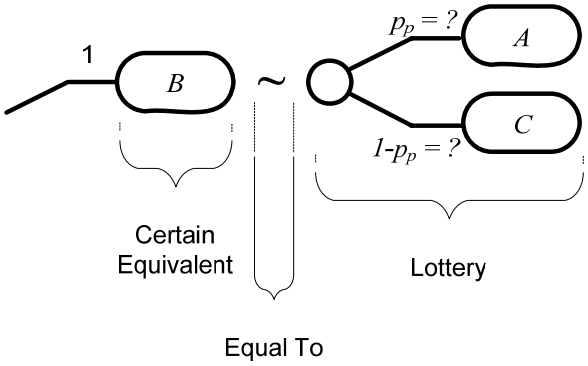

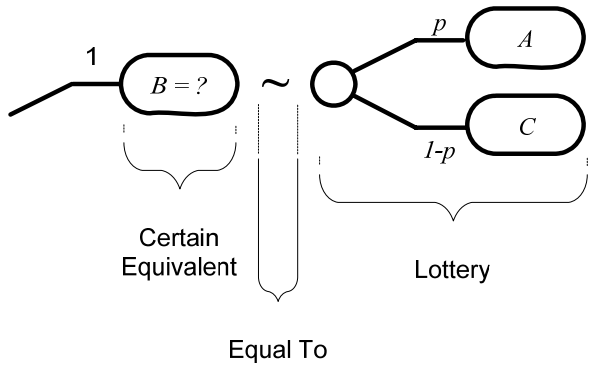

Figure E.3: Continuity Axiom (Certain Equivalent).

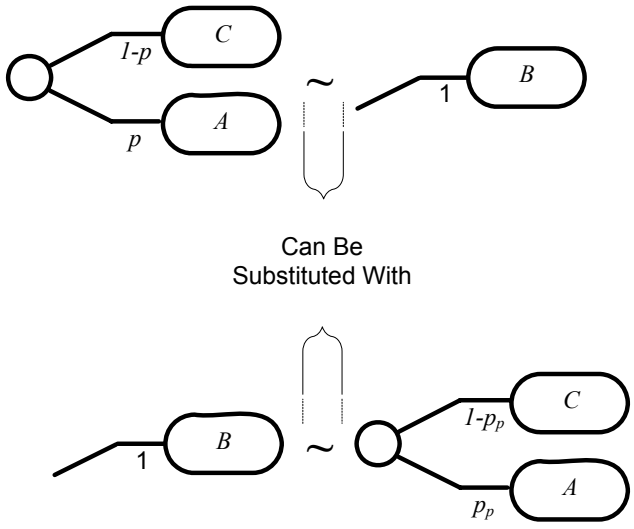

Figure E.4: Substitutability Axiom.

3. Let $A > B > C$. The *SUBSTITUTABILITY AXIOM* states that, in a decision tree, the lottery comprising A and C can be substituted by its certain equivalent B and vice versa. The lottery may actually occur with event probabilities p and $1 - p$ (upper graph in **Figure E.4**) or may be tagged with preference probabilities p_p an $1 - p_p$ expressing the preference of the decision maker on a given lottery to occur (lower graph in **Figure E.4**). By respecting the substitution rule the reader implicitly agrees to treat preference probabilities and event probabilities alike.

4. When facing a compound lottery (i.e. a lottery composed of other lotteries), the *DECOMPOSABILITY AXIOM* specifically only considers the final prize (or loss) for decision

making (sometimes also referred to as the Probability Rule). Consequently, a decision maker computes the (intermediate and final) probabilities by using Cox's Sum and Probability Rules.

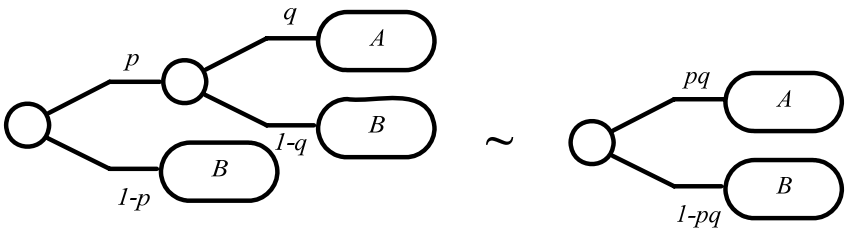

Figure E.5: Decomposability Axiom.

In the simplest case, applying the first four axioms to any decision situation produces a construct where alternatives are valued in terms of two lotteries each of which contains the same prospects but different probabilities for them to occur. This is shown in **Figure E.6**.

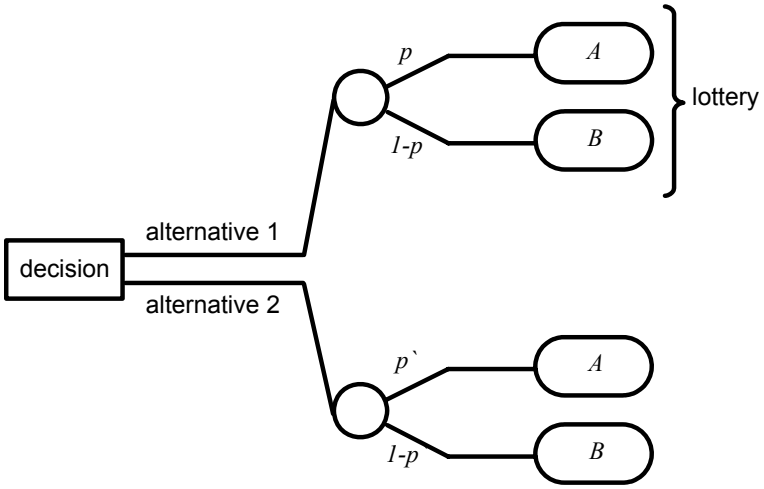

Figure E.6: Alternatives Represented by Two Lotteries.

5. Given the decision situation of **Figure E.6** and letting $A > B$, the *MONOTONICITY AXIOM* requires the decision maker to choose the alternative with the higher probability of achieving A. For example, given $A > B$ in the upper lottery of **Figure E.6** with *(p, A; (1 – p), B)* and the lower lottery with *(p`, A; (1 – p`), B)* and $p > p`$ then the decision maker chooses the upper lottery over the lower lottery.

E.3 Utility Function and Utility Curve

Utility Function and Properties: The *UTILITY FUNCTION u(.)* relates a commodity (e.g., money) to a utility number. The purpose of utility numbers is to compare the desirability of lotteries. They have two important properties:

1. The utility number u of a lottery, which yields a prospect A with probability p and a prospect B with probability $(1 - p)$, is:

$$u(A,B) = p \cdot u(A) + (1-p) \cdot u(B) \ , \qquad \textbf{(E.5)}$$

where u is the utility

p is the probability.

Further, the certain equivalent $\tilde{x}$ of the lottery has the same utility as the lottery, i.e.:

$$u(\tilde{x}) = p \cdot u(A) + (1-p) \cdot u(B) \ , \qquad \textbf{(E.6)}$$

where $\tilde{x}$ is the certain equivalent.

2. Between two lotteries, the decision maker prefers the lottery with the higher utility. A decision maker is assumed to always prefer more money. Consequently, the *U-CURVE* never decreases as the money value increases. Utility curves can be transferred to decision problems of different nature than they were previously constructed for provided they remain within the original monetary boundaries set out by the decision maker.

In addition to the utility function, the certain equivalent can also be determined graphically. For this, we apply the commodity to the x-axis and arbitrarily let $u(0) = 0$. Further, we let $u(100) = 1$ on the y-axis. In **Figure E.7**, the line consisting of 1 and 2 reflects the ratio of the probabilities for $u(A)$ and $u(B)$ respectively.

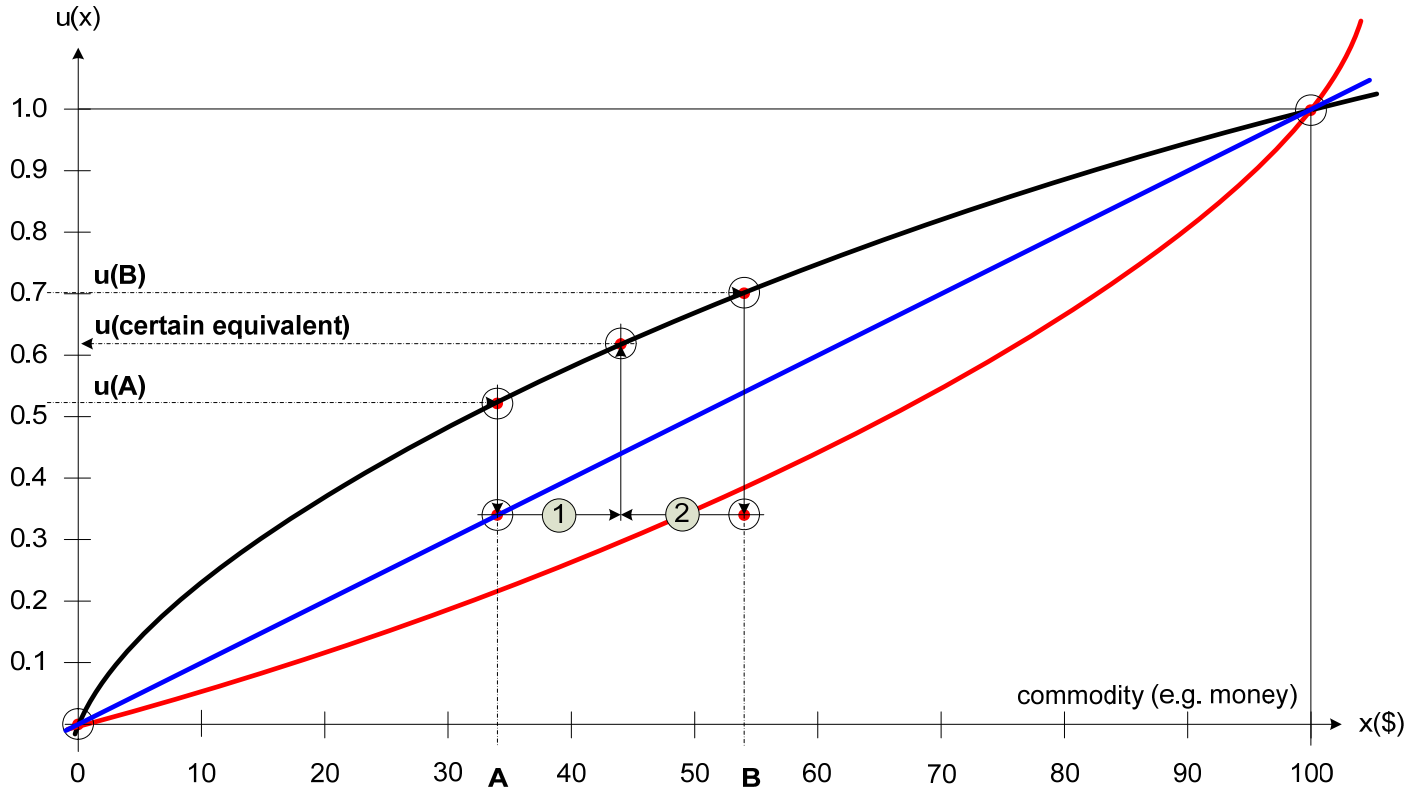

Figure E.7: Utility Curves.

Each individual has a unique *UTILITY CURVE*. Three types of utility curves are distinguished:

- a concave curve downwards is typical of a *RISK AVERTER* (**black curve**)
- a straight line indicates an individual, which is *RISK NEUTRAL* (**blue curve**)
- a concave curve upwards is typical of a *RISK PREFERRING* individual (**red curve**).

Assessment Methods for Utility Curves: There are two methods for assessing utility curves of individuals (i.e. their risk preferences). The first one is the *METHOD OF EQUIPROBABLE LOTTERIES* and the second is the method of probability assignment. To begin with, let $u(0) = 0$ and $u(100) = 1$ be an arbitrary normalization of the utility curve. With the method of equiprobable lotteries $p = 0.5$ is chosen and two out of three of A, B, and $\tilde{x}$ are stated. This approach distinguishes three cases:

$A = 100$ and $B = 0$, which yields $\tilde{x}$ by interpolation (left hand of **Figure E.8**)

1. $\tilde{x} = 0$ and $B = 100$, which yields A by extrapolating downward (middle of **Figure E.8**)
2. $\tilde{x} = 100$ and $A = 0$, which yields B by extrapolating upward (right hand of **Figure E.8**).

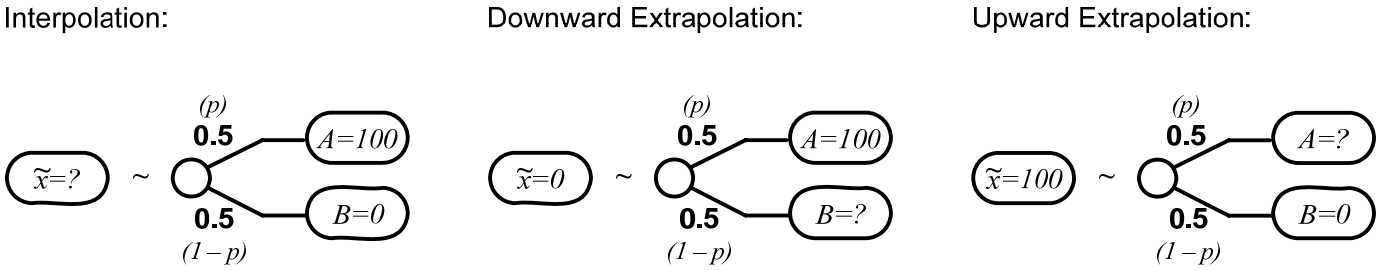

Figure E.8: Utility Assignment via the Method of Equiprobable Lotteries.

The *METHOD OF PROBABILITY ASSIGNMENT* requires the individual to state the preference probability of winning that makes the individual indifferent between a lottery and its certain equivalent. With this method A, B, and $\tilde{x}$ are fixed and yield p_p as shown by the following three cases:

1. $A = 100$, $B = 0$, $\tilde{x} = 50$, which yields p_p for the dotted line by interpolation (left hand of **Figure E.9**)
2. $A = 100$, $B = -100$, $\tilde{x} = 0$, which yields p_p for the dotted line by extrapolating downwards (middle of **Figure E.9**)
3. $A = 200$, $B = 0$, $\tilde{x} = 100$, which yields p_p for the dotted line by extrapolating upwards (right hand of **Figure E.9**).

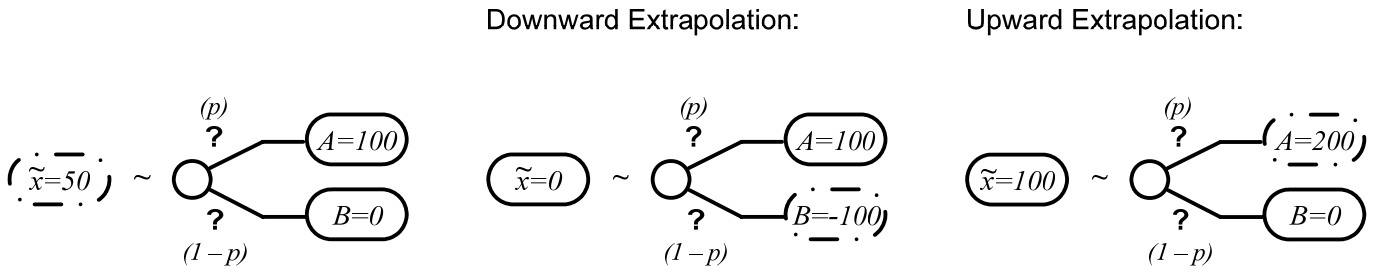

Figure E.9: Utility Assignment via the Method of Probability Assignment.

Appendix F: Classic Phishing Scenario

The following refers to the classic phishing scenario Chapter 7.2.4.

Classic Phishing	The classic phishing scenario starts with a fraudulent e-mail, which is sent randomly to potential victims. In turn, the victim surrenders online banking user credentials related to his/her bank account. Then, the attacker impersonates the user and cashes out money by using money mules.	
Threat Agent	Motivation	The attacker is motivated by illegitimate personal gain.
	Resources	The attacker uses technical equipment, which is commonly available on the market. The phisher commands a botnet.
	Location	The attacker is external to the targeted company.
	Knowledge	The attacker has specialized IT knowledge, especially in the area of networks. S/He understands how money forwarding channels like Western Union and E-gold work.
Phishing Chart (IS Context)	Victim / Attacker Interaction	This scenario is frequent in the US and is often operated from countries where the governments are unlikely to prosecute the phishers, e.g., Russia. The attacker sends an e-mail to the victim upon which the user visits the website of the attacker and surrenders his/her Internet banking login credentials. The attacker collects the credentials and executes the attack.

	Number of Targeted Victims	The number of targeted victims is usually large.	
Phishing Chart (Business Context)	Number of Money Mules	The number of money mules is low, e.g., one to a few.	
	Mule Attacker Interaction	The attacker interacts with the mule by phone calls and e-mails. The attacker pressures the mule to cash out the money and forward it via a channel other than Internet banking.	
	Consequence	Financial	Data is not published upon request of the remitter[93]. Losses comprise: refunds paid to customers, investigating fraudulent cases, enhancing internal control processes, training the internal work force or deploying security awareness among customers, projects to counteract phishing attacks, coordination efforts, etc.
		Reputation	None with the current *modus operandi* (refund the victims in turn for signing a non-disclosure agreement)
		Legal	Negligible loss

Asset	E-mail address **Event Number** ①
Threat Action	The attacker acquires the e-mail address of the potential victim.
Security Mechanism	Not investigated further.
Impact	loss of confidentiality of user e-mail address
Frequency	unknown (not investigated further)

[93] Original figure from interview with Menotti / Walder; refer to Salvati [181].

Probability	unknown (not investigated further)

Asset	User e-mail inbox **Event Number** (2)
Threat Action	The phisher spams the inbox of users at their Internet Service Provider (ISP), e.g., via a botnet to hide the tracks of the attacker and being in a position to deliver mass spamming to victims.
Security Mechanism	Spam filter at the ISP
Impact	loss of integrity of the user e-mail inbox
Frequency	*294* e-mails for the month of October 2007[94]
Probability	*0.0068 %* (*2* spam e-mails passed the filter / *294* overall spam e-mails)[95]

Asset	Internet banking login credentials **Event Number** (3)
Threat Action	The potential victim opens the e-mail and is subject to social engineering to deceive the victim into unveiling personal Internet banking login credentials to a malicious Internet site under the control of the attacker. The victim surrenders user credentials to the attacker.
Security Mechanism	Whether the threat action is successful depends on the victim recognizing the ongoing phishing attack. According to Gutermann (reported in Salvati [182]), once the attack is recognized, a small part of the users will report it to the internal help line within *15* to *27* minutes, which, in turn, sets forth a process for verifying and shutting down the website

[94] Measurement refers to the author's own mailbox; the result is not representative and must be put into perspective. Roughly, it is safe to assume that a filter blocks a spam email with a success probability of around 98%.

[95] Dito.

	of the attacker.
Impact	loss of confidentiality of Internet banking login credentials
Frequency	• unique phishing reports ascertained by APWG[96]: *31,650* (October 2007); *28,074* (November 2007; *25,688* (December 2007). • For our remitter: *0.48* successful phishing attacks per month or *29* single occasions where a customer lost money in a period of *60* months against the remitter of this study[97]
Probability	$p(C_{1159}, C_{34}) = 0.984$ (probability of threat succeeding in overcoming security mechanism) $p(C_{1159}, C_{240}) = 0.950$ $p(C_{20667}, C_{34}) = 0.948$ $p(C_{20667}, C_{240}) = 0.861$

[96] Anti Phishing Working Group: The Anti-Phishing Working Group (APWG) is the global pan-industrial and law enforcement association focused on eliminating the fraud and identity theft that result from phishing, pharming and e-mail spoofing of all types (taken from www.apwg.com). For further information refer to APWG [183].

[97] Original figure from interview with Menotti / Walder in November 2007; refer to Salvati [181].

Asset	Electronic account **Event Number** ⑧
Threat Action	Attacker transfers money
Security Mechanism	Company employs fraud detection system
Consequence	Data is not published upon request of the remitter[98]. A distinction is made between the money which was stolen (original losses) and the money that could not be recovered (actual losses).
	Reputation ok
	Legal ok
Frequency	• phishing reports received by APWG per month[99] • *0.48* phishing attacks per month or *29* successful phishing attacks in *60* months against the remitter of this study[100]
Probability	unknown

[98] Original figure from interview with Menotti / Walder in November 2007, refer to Salvati [181].

[99] Refer to www.apwg.com.

[100] Original figure from interview with Menotti / Walder; refer to Salvati [181].

Asset	Electronic account **Event Number** (9)
Threat Action	Attacker transfers money
Security Mechanism	Company reclaims money
Consequence	Data is not published upon request of the remitter[101]. Losses comprise: refunds paid to customers, investigating fraudulent cases, enhancing internal control processes, training the internal work force or deploying security awareness among customers, projects to counteract phishing attacks, coordination efforts, etc.
	Reputation ok
	Legal ok
Frequency	• phishing reports received by APWG per month[102] • data concerning the remitter of the case study is not published upon request
Probability	unknown

Table F.1: Classic Phishing Scenario.

[101] Original figure from interview with Menotti / Walder; refer to Salvati [181].

[102] Refer to www.apwg.com.

Appendix G: Phishing with Malicious Software

The following refers to the phishing with malware scenario of Chapter 7.2.4.

Phishing with Malicious Software	Spear phishers personalize their attacks as much as possible by using available information about the victim. They send e-mails that appear to come from the employer, a colleague, the head of human resources or a computer systems administrator. Because it comes from a known and trusted source, the request to execute some malware on the victim's computer is more plausible. The malware records the Internet banking login credentials and communicates them back to the attacker. Then, the attacker impersonates the user and cashes out money by using money mules.	
Threat Agent	Motivation	The attacker is motivated by illegitimate personal gain.
	Resources	The attacker uses technical equipment, which is commonly available on the market. The phisher commands a botnet.
	Location	The attacker is external to the targeted company.
	Knowledge	The attacker has specialized IT knowledge, especially in the area of networks. He understands how money forwarding channels like Western Union and E-gold work.
Phishing Chart (IS Context)	Victim / Attacker Interaction	This scenario is frequent in the US and is often operated from countries where the governments are unlikely to prosecute the phishers, e.g., Russia. The attacker sends an e-mail to the victim upon which the user visits the website of the attacker and surrenders Internet banking login credentials. The attacker collects the credentials and executes the attack.

	Number of Targeted Victims	The number of targeted victims is usually large.	
Phishing Chart (Business Context)	Number of Money Mules	The number of money mules is low, e.g., one to a few.	
	Mule Attacker Interaction	The attacker interacts with the mule by phone calls and e-mails. The attacker pressures the mule to cash out the money and forward it via a channel other than Internet banking.	
	Consequence	Financial	Data is not published upon request of the remitter[103]. Losses comprise: refunds paid to customers, investigating fraudulent cases, enhancing internal control processes, training the internal work force or deploying security awareness among customers, projects to counteract phishing attacks, coordination efforts, etc.
		Reputation	None with the current *modus operandi* (refund the victims in turn for signing a non-disclosure agreement)
		Legal	Negligible loss

Asset	E-mail address **Event Number** ①
Threat Action	The attacker acquires the e-mail address of the potential victim.
Security Mechanism	Not investigated further.
Impact	Loss of confidentiality of user e-mail address.
Frequency	unknown (not investigated further)

[103] Original figure from interview with Menotti / Walder; refer to Salvati [181].

Probability	unknown (not investigated further)

Asset	User e-mail inbox	**Event Number**	2
Threat Action	The phisher spams the inbox of users at their Internet Service Provider (ISP), e.g., via a botnet to hide the tracks of the attacker and being in a position to deliver mass spamming to victims.		
Security Mechanism	Spam filter at the ISP		
Impact	loss of integrity of the user e-mail inbox		
Frequency	*294* e-mails for the month of October 2007[104]		
Probability	*0.0068 %* (*2* spam e-mails passed the filter / *294* overall spam e-mails)[105]		

Asset	Computer of user	**Event Number**	4
Threat Action	The victim visits the website of the attacker and downloads the malware (e.g., a web trojan) onto the victim's computer.		
Security Mechanism	Company response (C_{34} or C_{240}), anti-virus software		
Impact	Loss of integrity of the computer of the user		
Frequency	Malware is often transported by e-mails, nifty appliances, games, etc. Depending on the kind of malware, the contribution on behalf of the user may be as insignificant as starting the e-mail program. • Worldwide unique variants of keyloggers: *359* (October 2007); *338* (November 2007); *269* (December 2007)		

[104] Measurement refers to the author's own mailbox; the result is not representative and must be put into perspective. Roughly, it is safe to assume that a filter blocks a spam email with a success probability of around 98%.

[105] Dito.

	• The available phishing data for F_{1159} indicates that out of *1,159* users, *562* fell for the attack (*48%*).
Probability	Same as in ❸ of **Appendix F** - if attack is deployed by e-mail and signature of malware is not recognized by anti-virus software

Asset	Internet banking login information **Event Number** ❺
Threat Action	Malware communicates back Internet banking login information to attacker.
Security Mechanism	Company response (C_{34} or C_{240}), personal firewall or company firewall
Impact	Loss of confidentiality of Internet banking login credentials of user
Frequency	The available data for F_{1159} indicates that out of *1,159* users, *373* have downloaded and executed the malware (*32%*).
Probability	Same as in ❸ of **Appendix F** - if attack is deployed by e-mail, installation is successful and firewall is badly configured

Asset	Electronic account **Event Number** 8
Threat Action	Attacker transfers money
Security Mechanism	Company employs fraud detection system
Consequence	Data is not published upon request of the remitter[106]. A distinction is made between the money which was stolen (original loss) and the money that could not be recovered (actual loss).
	Reputation ok
	Legal ok
Frequency	• phishing reports received by APWG per month[107] • data concerning the remitter of the case study not published upon request
Probability	Unknown

[106] Original figure from interview with Menotti / Walder; refer to [181].

[107] Refer to www.apwg.com.

Asset	Electronic account **Event Number** (9)
Threat Action	Attacker transfers money
Security Mechanism	Company reclaims money
Consequence	Data is not published upon request of the remitter[108]. Losses comprise: refunds paid to customers, investigating fraudulent cases, enhancing internal control processes, training the internal work force or deploying security awareness among customers, projects to counteract phishing attacks, coordination efforts, etc.
	Reputation ok
	Legal ok
Frequency	• phishing reports received by APWG per month[109] • data concerning the remitter of the case study not published upon request
Probability	Unknown

Table G.1: Phishing with Malicious Software.

[108] Original figure from interview with Menotti / Walder; refer to Salvati [181].

[109] Refer to www.apwg.com.

Appendix H: Curve Fitting for Threats and Security Mechanisms

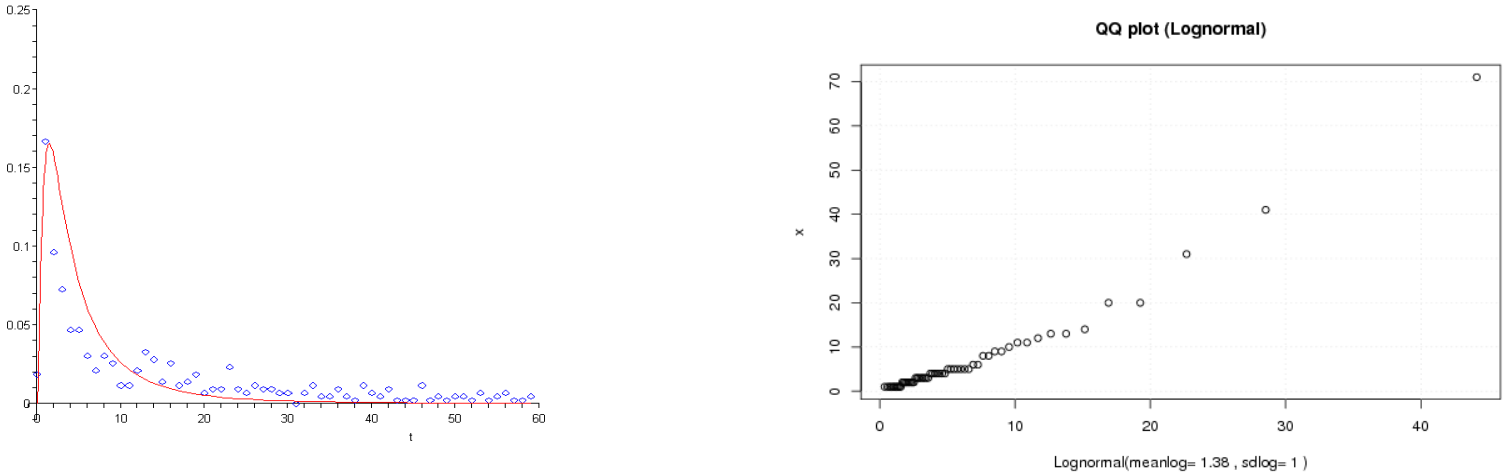

Figure H.1: Maximum Likelihood Curve Fitting and QQ-Plot for F_{1159} / C_{1159}

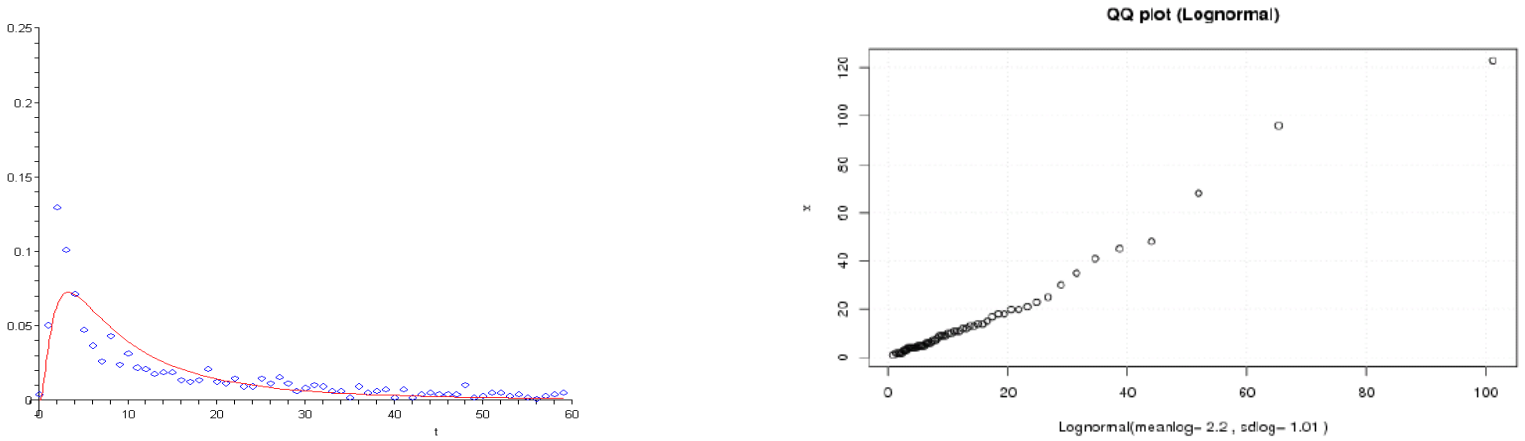

Figure H.2: Maximum Likelihood Curve Fitting and QQ-Plot for F_{20667} / C_{20667}

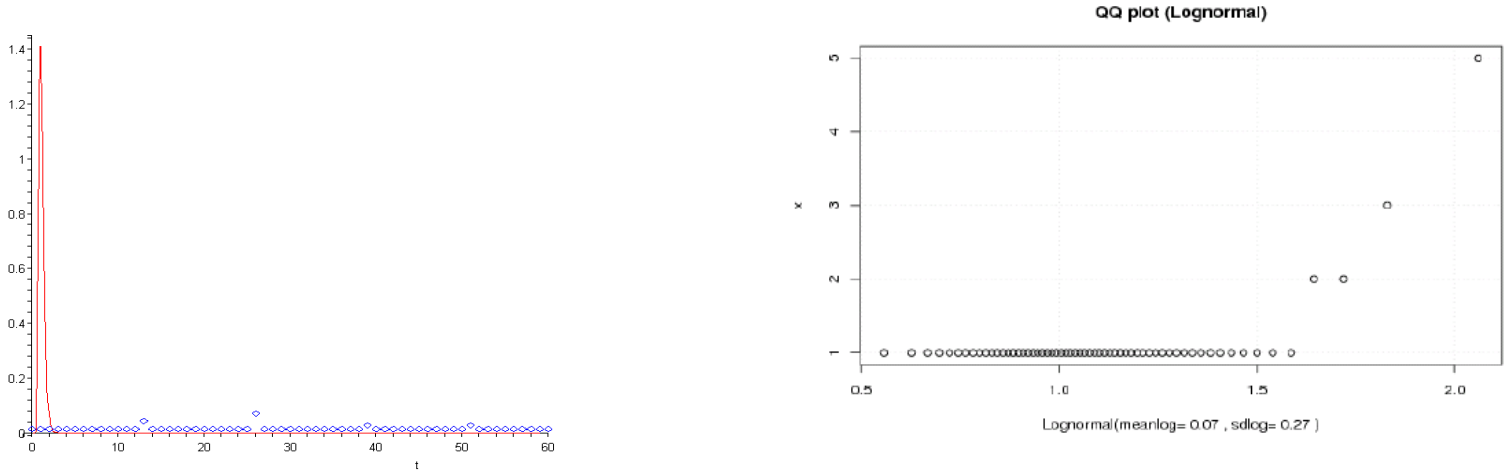

Figure H.3: "Manual" Curve Fitting and QQ-Plot for F_{34} / C_{34}

Appendix J: Probability Simulation

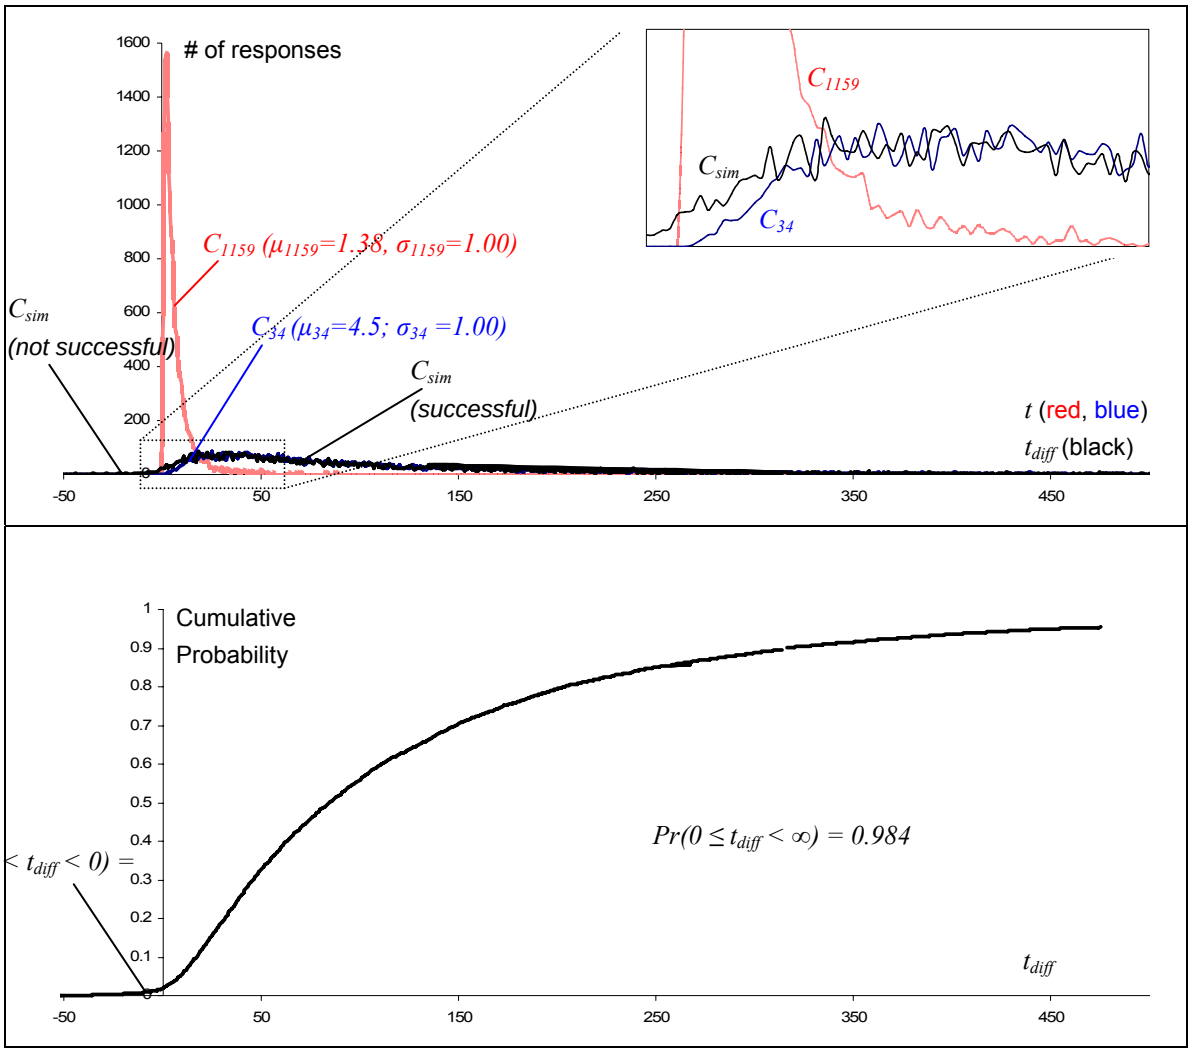

Figure J.1: Simulation and Probability showing C_{1159}, C_{34}, C_{sim}.

The upper part of **Figure J.1** shows C_{1159} and C_{34}. The convolution of the two curves yields C_{sim} which is then integrated to display the cumulative probability (lower part of **Figure J.1**). T_{diff} indicates the temporal difference between two arbitrary data points in C_{1159} and C_{34} respectively.

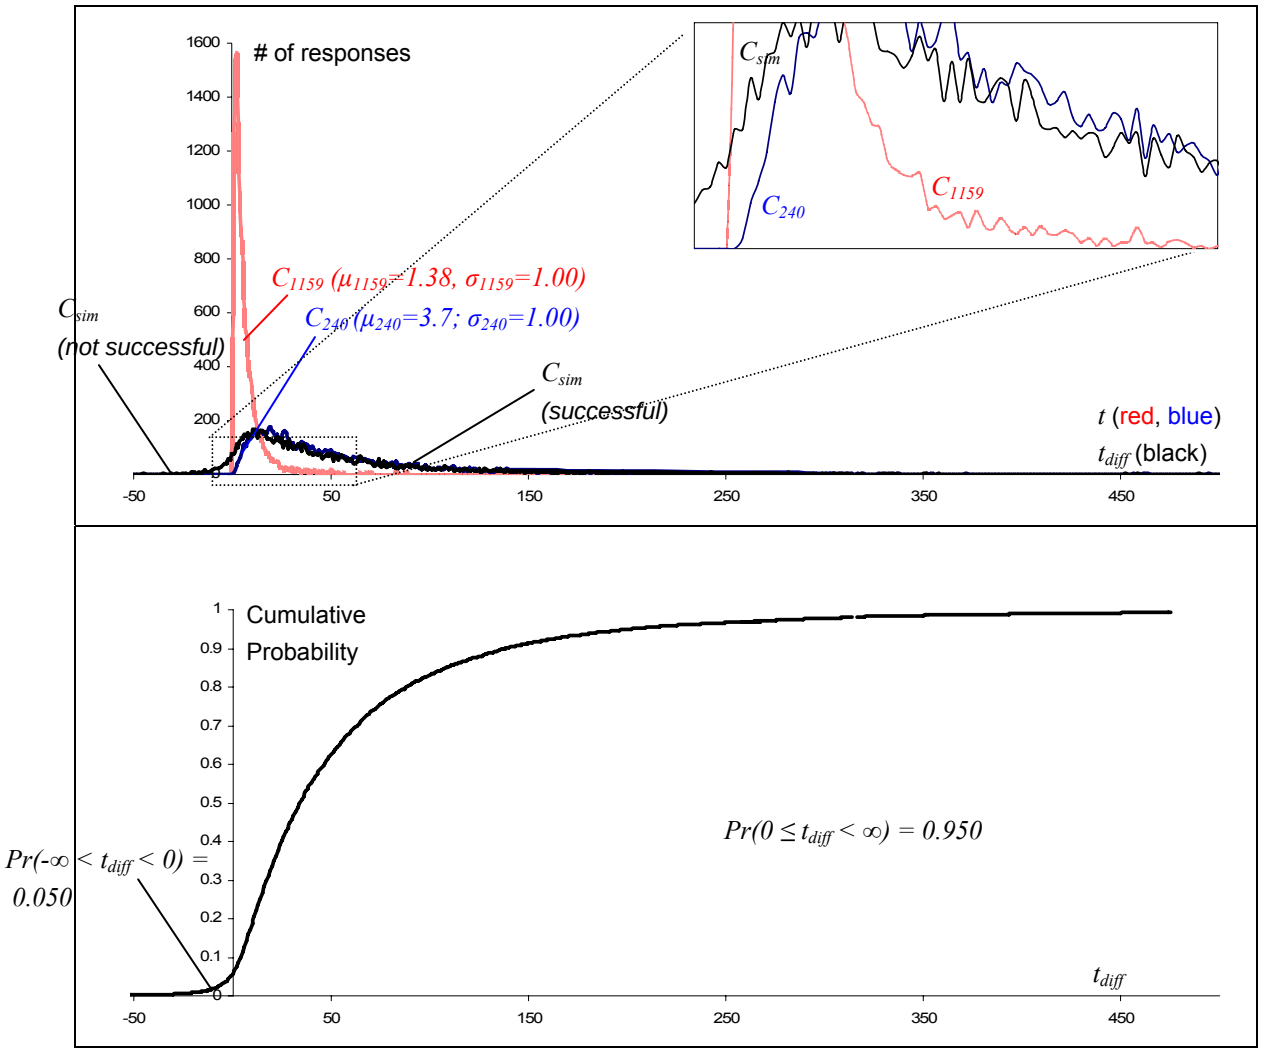

Figure J.2: Simulation and Probability for C_{1159}, C_{240}, C_{sim}.

The upper part of **Figure J.2** shows C_{1159} and C_{240}. The convolution of the two curves yields C_{sim} which is then integrated to display the cumulative probability (lower part of **Figure J.2**). T_{diff} indicates the temporal difference between two arbitrary data points in C_{1159} and C_{240} respectively.

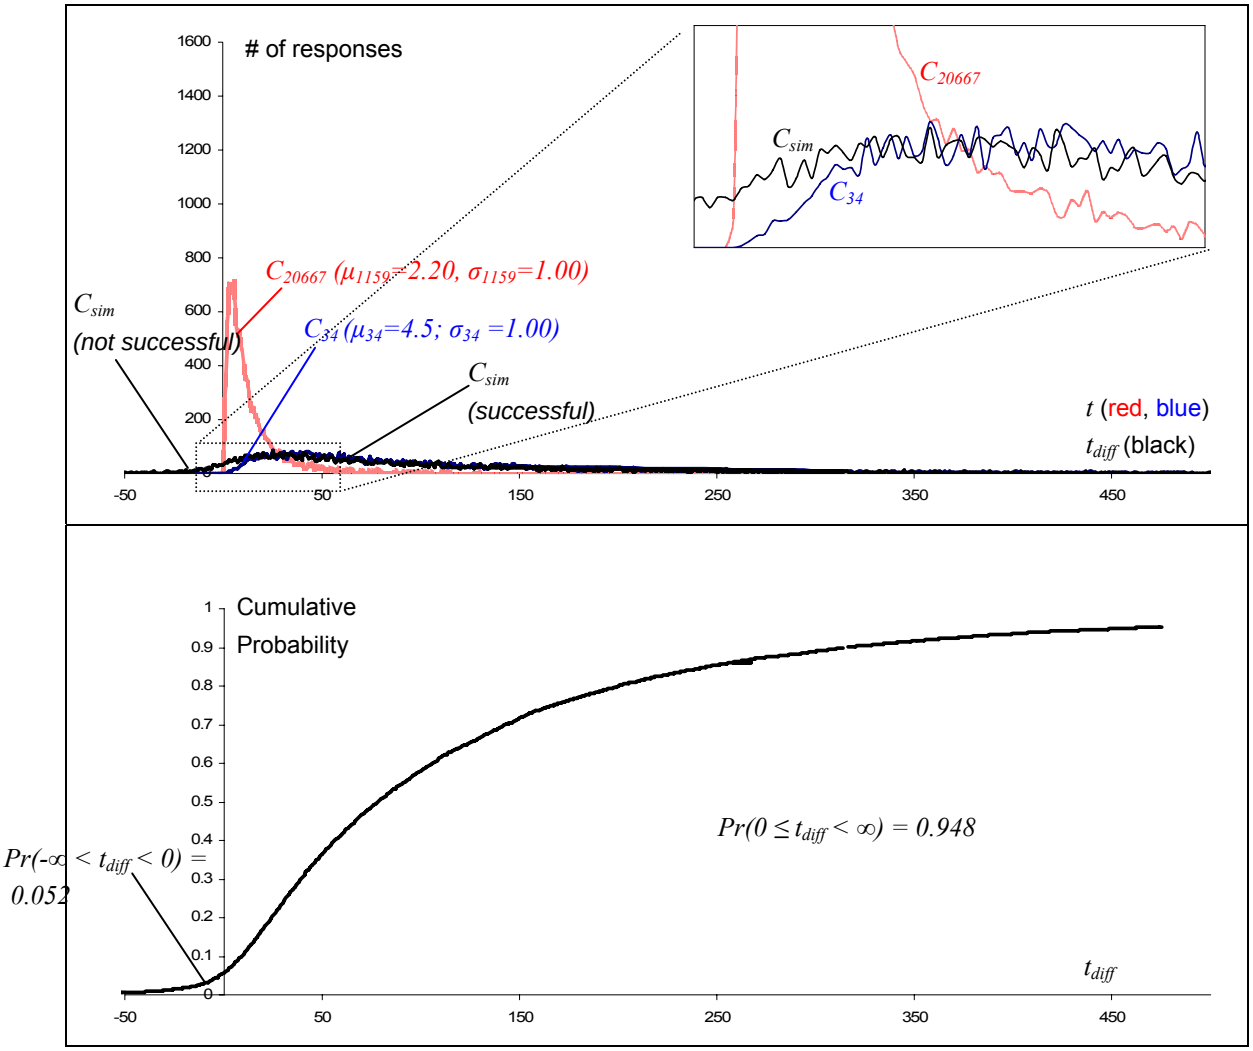

Figure J.3: Simulation and Probability for C_{20667}, C_{34}, C_{sim}.

The upper part of **Figure J.3** shows C_{20667} and C_{34}. The convolution of the two curves yields C_{sim} which is then integrated to display the cumulative probability (lower part of **Figure J.3**). T_{diff} indicates the temporal difference between two arbitrary data points in C_{20667} and C_{34} respectively.

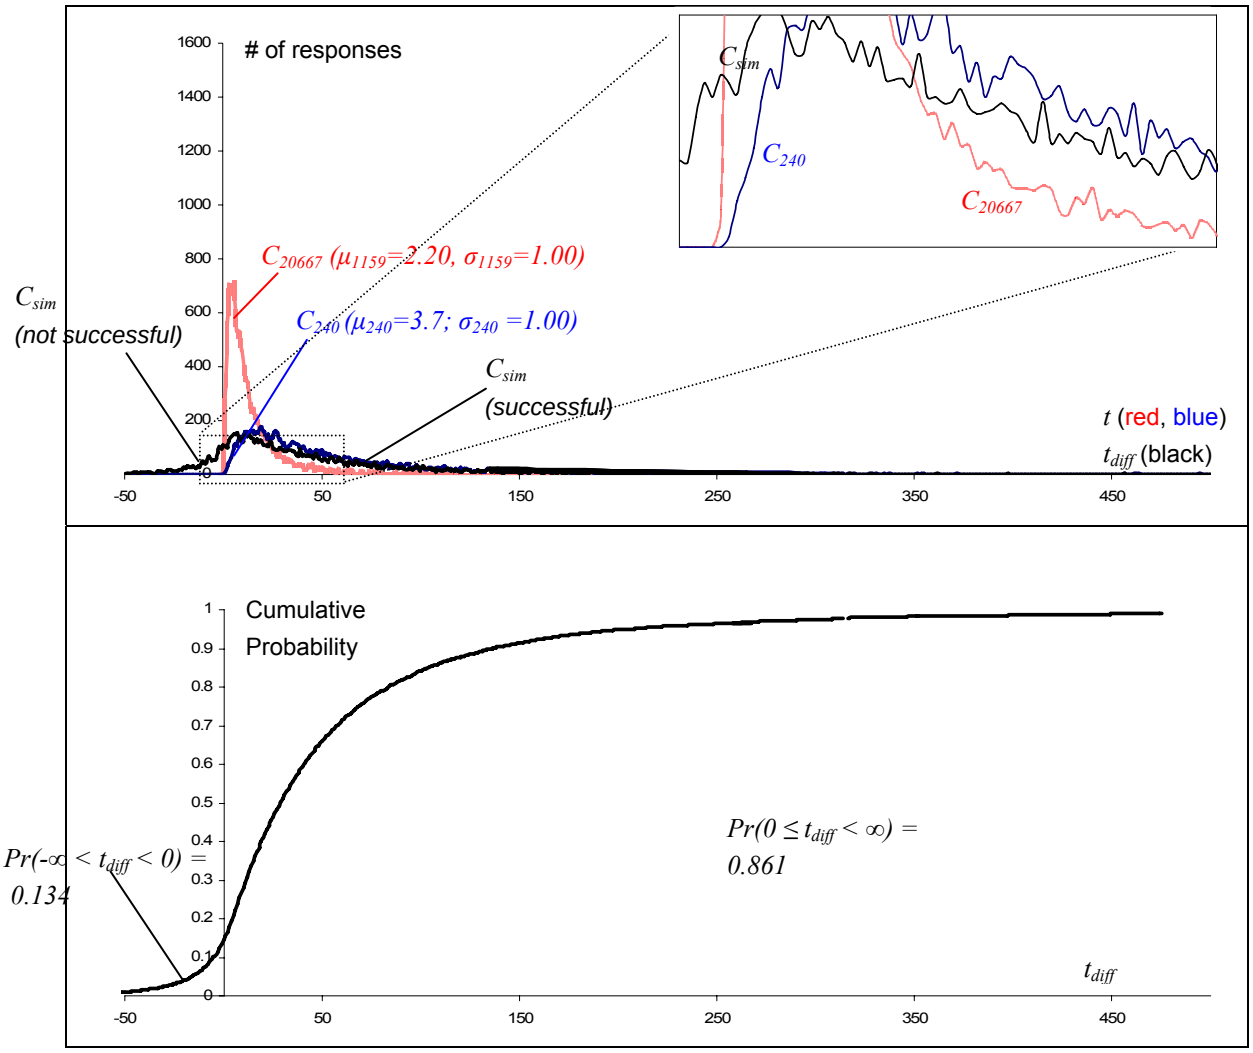

Figure J.4: Simulation and Probability for C_{20667}, C_{240}, C_{sim}.

The upper part of **Figure J.4** shows C_{20667} and C_{240}. The convolution of the two curves yields C_{sim} which is then integrated to display the cumulative probability (lower part of **Figure J.4**). T_{diff} indicates the temporal difference between two arbitrary data points in C_{20667} and C_{240} respectively.

Appendix K: Lognormal Distribution

$$\Lambda(\mu,\sigma\ ) = f_{u,\sigma}\ (t) = \begin{cases} \dfrac{1}{t\sigma\sqrt{2\pi}} \cdot e^{-\frac{(\ln(t)-\mu)^2}{2\sigma^2}} & (t > 0) \\ 0 & (t \le 0) \end{cases}, \tag{K.1}$$

where $\Lambda(.)$ is the lognormal probability density
μ is the mean
σ is the standard deviation
t is time.

mode at:

$$e^{\mu-\sigma^2}, \tag{K.2}$$

lower quartile at:

$$e^{\mu-0.67\sigma}, \tag{K.3}$$

median at:

$$e^{\mu}, \tag{K.4}$$

mean at:

$$e^{\mu+0.5\sigma^2}, \tag{K.5}$$

upper quartile at:

$$e^{\mu+0.67\sigma}. \tag{K.6}$$

Appendix L: Success Probabilities of Phishing Attacks (Internal Notification)

Appendix L shows the solution approach for the calculation of success probabilities of phishing attacks which are notified by users to company-internal authorities. This example differs from the Case Study in such that MELANI is not asked to request a hosting provider to shut down a fraudulent Internet site. Instead, the password notification is prevented within the company premises. The following steps are devised:

Step 1: Measure the frequency curve of users responding to the phishing attack. Presupposing that it follows a lognormal distribution, assess the parameters μ_{resp} and σ_{resp}.

Step 2: For the sake of simplicity, the continuous user response curve obtained under **Step 1** is assumed to represent the ongoing phishing attack as well as future attacks.

Step 3: Measure the frequency curve of users notifying internal authorities. Typically, only a small percentage of users will provide such a notification to internal authorities. These individual notifications related to the specific ongoing phishing attack.

Step 4: Presupposing that the notification curve follows a lognormal distribution, its parameters μ_{notify} and σ_{notify} are assessed by the means shown in Chapter 7.

Step 5: The continuous notification curve is related to the specific ongoing phishing attack. To obtain the company response curve to future phishing attacks the following steps are proposed (**Steps 8** though **9**).

Step 6: From the newly obtained continuous curve, generate a number of samples equal to the number of notifications received under the ongoing phishing attack.

Step 7: Repeat **Step 5** for a sufficient number of times, e.g., *1,000* times.

Step 8: From each sample set that was generated take the fastest notification. Fit the fastest sample of each generation into company response curve.

Step 9: Presupposing that the company response curve follows a lognormal distribution, its parameters μ_{comp} and σ_{comp} are assessed by the means shown in Chapter 7.

Step 10: Proceed in the calculation of probabilities as devised in Chapter 7.

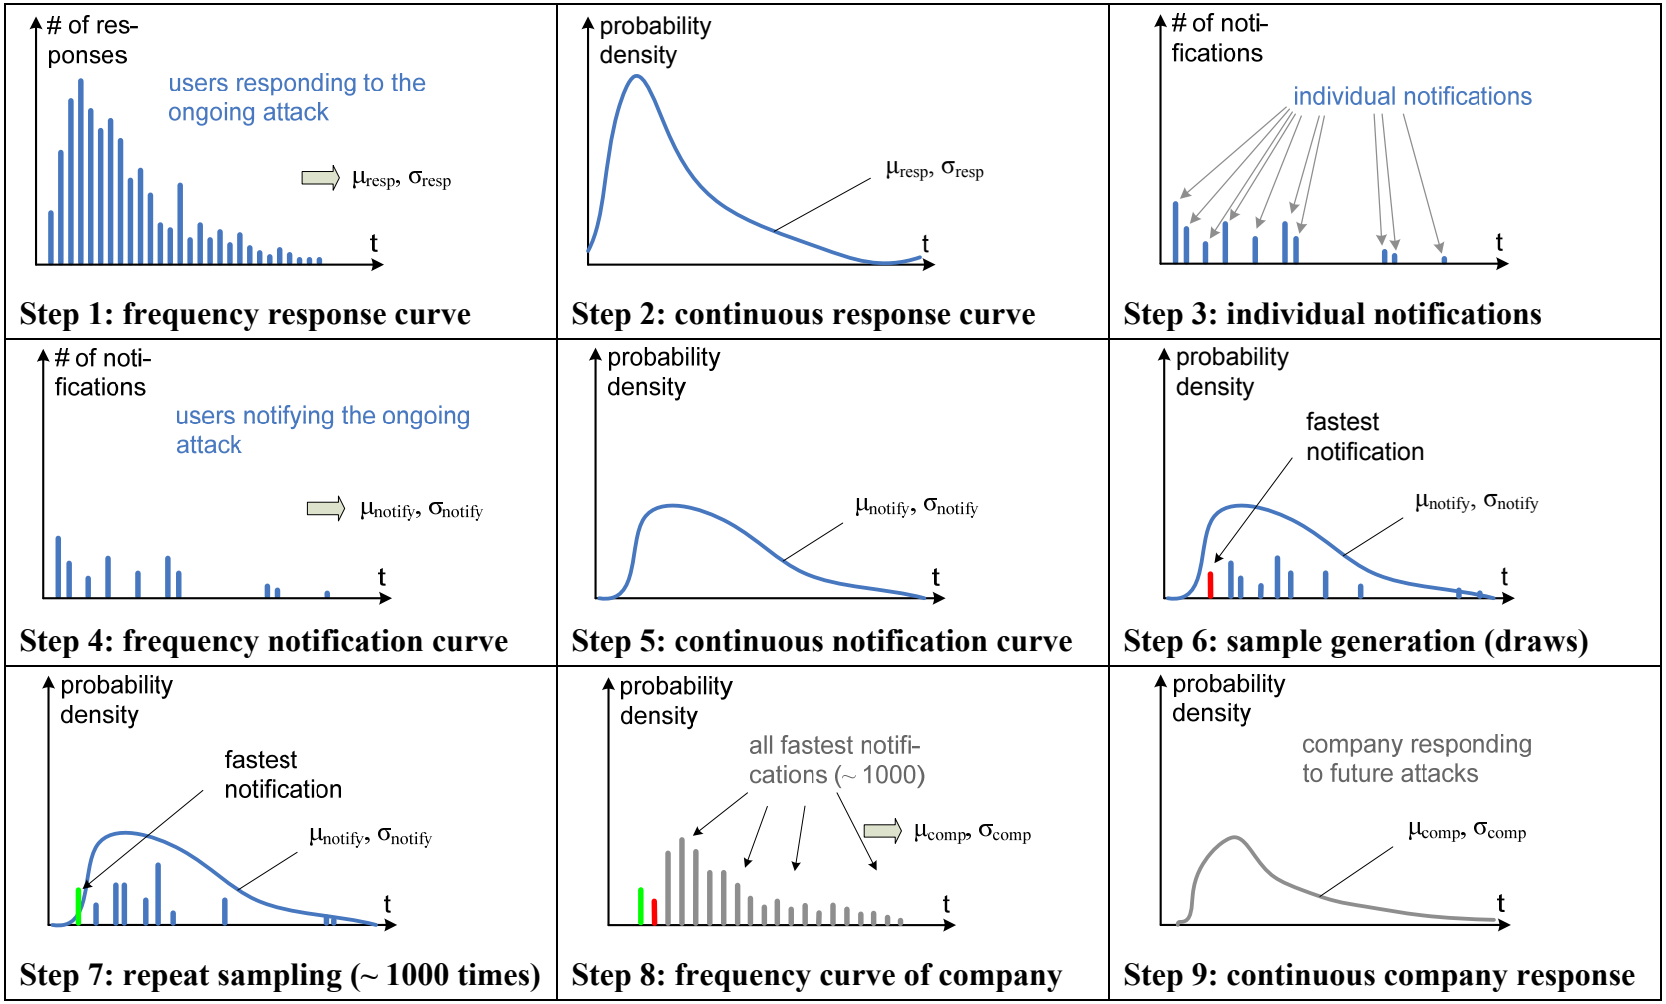

Figure L.1: 9 Steps to Resampling Internal Notifications.

Appendix M: Risk Preferences

The following shows parts of the questionnaire used in November 2007 and March 2008 to ascertain the risk preferences of decision makers.

Questions: Series 1

Question 1:

- We explore the shape of the utility curve between (0, 10).
- Question: "What is your certain equivalent if you stand to gain 10 with a probability of 50% and lose nothing with a probability of 50%? Answer here: "_____". Report the answer also on the line bordered by parentheses "< >" in questions 2 and 3.
- Please note that since you are likely to be risk averse, the certain equivalent is usually below the expected value of the lottery, i.e. it is below (0.5 * 10) + (0.5 * 0) = 5.
- Please note it is your preferences that are of interest which you assess by taking into consideration your business context.

Question 2:

- We explore the shape of the utility curve between (0, 10).
- Question: "What is your certain equivalent if you stand to gain <_____> with a probability of 50% and lose nothing with a probability of 50%? Answer here: "_____".

Question 3:

- We explore the shape of the utility curve between (0, 10).
- Question: "What is your certain equivalent if you stand to gain 10 with a probability of 50% and lose <_____> with a probability of 50%? Answer here: "_____".

Question 4:

- We explore the shape of the utility curve below the region 0.
- Assume the certain equivalent is 0. You face a lottery where both prospects occur with 50%. The winning prize is 10. How high is your losing prize? Answer here: "_____".

Question 5:

- We explore the shape of the utility curve above the region 10.
- Assume the certain equivalent is 10. You face a lottery where both prospects occur with 50%. The losing prize is 0. How high is your winning prize? Answer here: "_____".

Questions: Series 2

… [omitted; analogous to Series 1]…

Questions: Series 3

… [omitted; analogous to Series 1]…

Questions: Series 4

… [omitted; analogous to Series 1]…

Thank you for your time.

Domenico Salvati, October 25th, 2007

Appendix N: The Allais Paradox

The following exemplifies the Allais Paradox. Let:

- *A = u(*win ***100 Mio*** with a chance of ***89%***; win *100 Mio* with a chance of *11%)* and
- *B = u(500 Mio, 10%*; ***100 Mio***, ***89%***; *0 Mio, 1%)*

where *u(.)* signifies the utility function for the associated lottery.

Further let:

- *C = u(100 Mio, 11%*; ***0 Mio***, ***89%)*** and
- *D = u(500 Mio, 10%*; ***0 Mio***, ***89%***; *0 Mio, 1%).*

Choose between the lotteries *A* or *B* and then between *C* and *D*!

Most people will choose *A* over *B* and then choose *D* over *C*. However, this is inconsistent with the decomposability axiom! The inconsistency appears if the common consequence in the equations above is omitted (highlighted in **bold text**):

Let:

- *A' = u(100 Mio*, *11%)* and
- *B'= u(500 Mio, 10%*; *0 Mio, 1%)*

- *C' = u(100 Mio, 11%)* and
- *D' = u(500 Mio, 10%*; *0 Mio, 1%).*

By comparing the utility *A, B* and *C, D* as well as *A',B'* and *C',D'* it can easily be verified that the expected utility would suggest to choose differently than the actual observed behavior!

Curriculum Vitae

Domenico Salvati, 23.01.1968

Work Experience

since 11.04 **IT Risk Officer** Credit Suisse, Zurich

- IT Risk Reporting for Executive Management
- Execution of IT Risk assessments in the Americas, Europe and Middle East
- Support and consulting of banking branches in IT Security matters

06.01 – 10.04 **Head Operational Security Services** Credit Suisse, Zurich

- Project & network security, spot checks on production security (4 team members)
- Build-up and functional mgmt of the global IT Security officers network (team of 15)

11.99 – 06.01 **Chief Information Security Officer** MIGROSBANK, Wallisellen

05.98 – 04.99 **IT Audit and Consulting** Arthur Andersen, Oerlikon

Education (Excerpt)

11.04 – 12.08 **Doctoral Student** Swiss Federal Institute of Technology, Zurich

- The studies are sponsored by the Academic Research Programs of Credit Suisse
- Thesis is elaborated part time (40% doctoral student; 60% employee of Credit Suisse)

03.91 – 02.98 **Lic.oec.publ** University of Zurich

- Studies in Information Management with emphasis on IT Security
- Thesis "Org. Aspects of IT Security" awarded Best Paper in Stuttgart, Germany
- Jobs as working student

85 – 86 **High School Diploma** San Andreas, CA, USA

77 - 83 **Scuola di Lingua e Cultura Italiana** Cham, Switzerland

Languages

German	Mother tongue
Italian	Mother tongue
English, exchange student in 85 – 86	Grammar School Diploma, Zug, CH
French	Grammar School Diploma, Zug, CH