Evaluating Risk Management Strategies for Third Party Payment Networks

A Case Study from the Payment Service Provider's Perspective

J.W. van Driel





Evaluating Risk Management Strategies for Third Party Payment Networks

A Case Study from the Payment Service Provider's Perspective

by

J.W. van Driel

to obtain the degree of Master of Science at the Delft University of Technology, to be defended publicly on Monday December 7, 2015 at 13:00.

Student number:4010833Project duration:May 1, 2015 – December 1, 2015Thesis committee:Prof. dr. M.J.G. van Eeten,
Dr. ir. C.H. Gañán,TU Delft, chairman
TU Delft, daily supervisor
TU Delft, first supervisor
TU Delft, first supervisor
TU Delft, second supervisor
Ir. L.W.M. Lobbezoo,

This thesis is confidential and cannot be made public until December 1, 2018.

An electronic version of this thesis is available at http://repository.tudelft.nl/.



If you think nobody cares about you, try missing a couple of payments

— Steven Wright

Copyright (c) 2015 Willem van Driel and Adyen B.V.

No part of this thesis may be reproduced, stored or transmitted without the written permission of Willem van Driel and Adyen B.V. The views expressed in this thesis are those of Willem van Driel and do not reflect the official policy or position of Adyen B.V. or any other party mentioned in this thesis.

Summary

Like most consumers, you probably take the **payment process** when checking out at your favorite webshop, for granted. Although the payment process may look simple from the consumer's perspective, a lot of complexity is hidden behind the scenes. Over the past decades an entire financial sector has emerged to facilitate the payment process: the **payment industry**. The industry is dominated by a small number of players, like MasterCard and Visa, that operate the **schemes** for the facilitation of credit and debit card payments. Regardless of which **payment method** you select; PayPal, Google Wallet, Apple Pay or you directly provide your card details, chances are big that the schemes of MasterCard or Visa are used. A scheme can be regarded as a standardized way of communication between the consumer's bank, the **issuer**, and the merchant's bank, the **acquirer**. We refer to the part of the payment industry that facilitates credit and debit card payments as the **card network**.

In order to create more competition for the card network, the European Commission is working on Payment Service Directive II (PSDII) which introduces the **third party payment service provider** (TPP). With this introduction the **TPP network** is created, next to the existing card network, providing the payment industry the opportunity to develop payment methods that circumvent the dominant position of MasterCard and Visa. The parties involved in the TPP network are visualized in Figure 1. We already touched on the role of the acquirer and issuer, in the remainder of this summary we will explain the roles of the other parties.

Just like for the payment process in the card network, the process in the TPP network comprises of two main phases: the **authorization** and **settlement**. During the authorization phase, the merchant obtains the confirmation whether the consumer has sufficient fund availability for the purchase. During the settlement phase, the funds are moved from the consumer's bank account at the issuer to the merchant's bank account at the acquirer. The main difference between the payment processes in the card and TPP network resides in the fact that in the card network the schemes are used for the authorization, while in the TPP network the authorization is provided by the TPP. This difference in authorization has a complication which we have investigated in our research.



Figure 1: Extended four party model representing the TPP network

Complication

A **payment service provider** (PSP) provides aggregation services to merchants, enabling them to accept payments from consumers independent of the payment method they use. Payment methods broadly fall into two categories: **guaranteed** and **non-guaranteed**. Non-guaranteed payment methods offer the consumer the option to reverse a payment after it has taken place, while guaranteed payment methods in the card network are typically non-guaranteed, while payment methods in the TPP network should be guaranteed. Practice has shown however that merchants that make use of payment methods in the TPP network are confronted with reversed payments — i.e. the TPP has provided an authorization, however no settlement has been received. We refer to these transactions as **non-payments**. There is a lack of knowledge about the reasons of the occurrence of the non-payments. We do know that their occurrence is triggered either intentionally or unintentionally, which we linked to fraud respectively technical malfunctioning in our conceptual model. In our research we have aimed to understand how big the financial risk of the occurrence of non-payments is and what could be strategies a PSP can use to manage this risk, as PSPs have the responsibility towards their customers, the merchants, to provide guaranteed payment methods that are truly guaranteed.

Approach

The goal of our research is two-fold. The primary goal is **G1**. To design strategies that can be used by the PSP to manage the financial risk of non-payments in the TPP network efficiently. The secondary goal is **G2**. To evaluate whether the introduction of the TPP network helps to reach the second goal of PSDII: the creation of more competition within the European market for payment methods. In order to achieve these goals, we formulated the following research question **RQ**. How can a PSP manage the financial risk of non-payments in TPP networks efficiently, such that it becomes more efficient than the card network?

For our research we have been provided the opportunity by PSP Adyen to make use of transaction data of German Sofort, one of the first TPPs active in the European Union and currently the Union's biggest TPP. After an initial data exploration, we have identified two preventive strategies and one reactive strategy to manage the risk of the occurrence of non-payments. The preventive strategies are based on the application of data mining techniques, random forest and break point analysis, enabling us to block transactions with a high probability that they will result in a non-payment. For the reactive strategy we have designed a dunning process, enabling us to recover the due funds from the consumer after the non-payment has taken place. We have applied the three strategies in practice and evaluated their financial efficiency (contributing to G1.). For the evaluation we have used the return on security investment model (ROSI) combined with the loss distribution approach (LDA) to quantify the risk exposure resulting from the occurrence of non-payments. In the last phase of our research we have compared the merchant's transaction costs in the card and TPP networks (contributing to G2.).

Results

We have concluded that the reactive strategy comprising the **dunning process** has a higher ROSI value than the preventive strategies. As such, the dunning process should be preferred over the preventive strategies. When considering the expected value of the risk exposure of the non-payments, the ROSI of the dunning process is four, which means that for every euro invested in the strategy, four euro revenue is generated. By comparing the merchant's transaction costs in the card network and the TPP network, with the application of the dunning process, we have observed that transaction costs in the TPP network can be lower than in the card network. As such, we believe that the TPP network is a cost competitor of the card network when considering the merchant's transaction costs.

Are these results **generalizable**? To answer this question we distinguished the (1) generalizability of the Sofort case and (2) generalizability of the results of the application of the risk management strategies. For the generalizability of the Sofort case, we derived that Sofort is (1) the biggest TPP within the European Union and (2) Adyen's Sofort portfolio comprises **of** of the total Sofort transaction volume. As such, we believe that our case study has the potential to produce results which can be generalizable to other TPPs. For the generalizability of the results of application of the risk management strategies, we focused on the generalizability of the most promising strategy, the dunning process. The merchant that cooperated in the dunning process sells digital gift cards which are prime subjects of fraudsters, as they are easy to cash out. We believe that if dunning works for this type of *high risk* products, it will also work for products or services which are less easy to cash out.

We see four main **contributions** in our work. The contributions are mainly driven by the uniqueness of the case study, as to the best of our knowledge we are the first ones to have investigated TPP transaction data in an academic setting. The first contribution encompasses describing the payment process in the TPP network and contrasting it to the process in the card network. The second contribution entails the analysis of the Sofort transaction data and the identification of patterns that distinguish payments from non-payment. The third contribution entails the design and application of risk management strategies for the TPP network. Finally, we see a contribution in the cost comparison of the merchant's transaction costs in the card and TPP networks.

Next Steps

We have identified eight **opportunities for future research**. The first opportunity encompasses the investigation of different TPP models. Although we believe that the overlay model — in which the consumer passes his credentials through the TPP to the issuer and that is currently used by TPPs in the European Union — has quite some potential, country specific implementations of the PSDII directive might push alternative models to the market. We see the second opportunity in the identification of different causes of non-payments besides the ones we have considered: consumer fraud and technical malfunctioning caused by the issuer. Third, we see potential for investigating data from other TPPs than Sofort, to create more, and more representative, insights into the TPP network. As the fourth opportunity we see potential for investigating TPP data from outside the DACH-region which we have considered (Germany, Austria and Switzerland), for the same reasons.

Besides, we see opportunities for future research into the investigation of different subsets of the Sofort dataset, especially subsets that are characterized by a higher imbalancedness in class type prevalence — i.e. whether a transaction is a payment or non-payment. Such research can be beneficial to increase the efficiency of the strategy application, not only in our research, but in data mining in general. Sixth, we see potential for optimization of our strategy implementation, as this has not been the main focus of our research. Seventh, we would recommend to incorporate indirect costs in the strategy evaluation, as this will provide more insight into the full cost of the application of the strategies. At last, we see potential for optimization of our evaluation implementation, mainly focused on optimizing the distribution fitting for calculating the risk exposure in the loss distribution approach.

Preface

When I started my graduation project at Adyen in April this year, I entered a world unknown to me until that moment: the payment industry. Looking back, I can say it has been a great time in which I was given the opportunity to explore the industry from various perspectives. The main industry lesson I learned can be captured by comedian Steven Wright's saying: 'If you think nobody cares about you, try missing a couple of payments.' A successful payment process means a lot to the customers served by the payment industry: the merchants.

The last seven months have been a time of transition. The transition started when the graduation committee changed early in the process. A little while later, I moved from Delft to Amsterdam. Mean-while, my research assessed the transition that is currently taking place in the payment industry. I believe we are at the start of a new payments era. Unbundling, as already witnessed in many other sectors, is now affecting our payment infrastructure.

I would like to start by thanking Jan van den Berg, Dina Hadžiosmanović and Hadi Asghari for their initial guidance of the project. Moreover, I would like to thank Carlos Gañán and Michel van Eeten, for offering me the opportunity to graduate within the Economics of Cybersecurity group. I really enjoyed working together! And of course — many thanks to Mark de Reuver and Wolter Pieters for their guidance during the project. Special thanks go to Maikel Lobbezoo, for offering me the opportunity to perform my research at Adyen. I admire your drive and your great dedication to build such a company! Besides, my project would not have been possible without the dedication of my colleagues at Adyen. Special thanks go to Christopher, Tony, Bert, Brian, Steffen, Guus, Ingo, Huub — and my fellow intern Ruben.

I would like to thank Ross Anderson for the inspiring conversations we have had over the past few months. Also, many thanks to Georg Schardt for inviting me over to the Sofort office and providing valuable insights. Special thanks also to the other people I had the privilege to interview for my research.

I would not have been able to finish the project without the support of my family and friends. First of all I would like to thank my parents, Harry and Anita, and my brother, Berend Jan, for their support. Many thanks to my friends and housemates for their support and the necessary distraction when needed. Special thanks go to Roy, for inspiring me to take on the challenge of performing my graduation project at Adyen.

As I've came to realize over the past seven months, the payment industry is a complex world. Although I hope that all the used concepts and terminology are explained well enough in the main text, I've added a glossary in Appendix G. Besides, important notions are displayed in blue boxes to highlight their emphasis. I hope that our work as presented in this thesis can inspire both academia and decision makers in the payment industry. Although my graduation project is coming to an end, I am sure it is the beginning of much more. Enjoy reading!

Willem Amsterdam, November 2015

Contents

1	Intro	oduction 1
	1.1	Research Problem and Goal
	1.2	Research Opportunity
	1.3	Research Gap and Questions
	1.4	Research Methodology
	1.5	Document Structure
2	Bac	karound
-	2 1	Theory on the Payment Industry
	2.1	2 1 1 The Payment Process
		2.1.2 Historical Development
		2.1.2 Instance Developments
		2.1.6 Recent Folicy Developments
	22	Theory on the Economics of Cybersecurity
	2.2	2.2.1 Security Aspects of Payments
		2.2.1 Security Aspects of Fayments
		2.2.2 Occurry Methods
	23	Theory on Knowledge Discovery 21
	2.5	Concentual Framework 23
	2.7	
3	Cas	e Study 25
	3.1	Case Study Introduction
		3.1.1 Characterization of the German Banking Landscape
		3.1.2 The Sofort Payment Process
	3.2	3.1.2 The Sofort Payment Process 26 Case Study Generalizability 29
	3.2 3.3	3.1.2 The Sofort Payment Process 26 Case Study Generalizability 29 Initial Data Exploration 30
	3.2 3.3	3.1.2 The Sofort Payment Process 26 Case Study Generalizability 29 Initial Data Exploration 30 3.3.1 Transaction 31
	3.2 3.3	3.1.2 The Sofort Payment Process 26 Case Study Generalizability 29 Initial Data Exploration 30 3.3.1 Transaction 31 3.3.2 Merchant 33
	3.2 3.3	3.1.2 The Sofort Payment Process26Case Study Generalizability29Initial Data Exploration303.3.1 Transaction313.3.2 Merchant333.3.3 Consumer33
	3.2 3.3	3.1.2 The Sofort Payment Process26Case Study Generalizability29Initial Data Exploration303.3.1 Transaction313.3.2 Merchant333.3.3 Consumer333.3.4 Issuer35
4	3.2 3.3 Stra	3.1.2 The Sofort Payment Process 26 Case Study Generalizability 29 Initial Data Exploration 30 3.3.1 Transaction 31 3.3.2 Merchant 33 3.3.3 Consumer 33 3.3.4 Issuer 35 teav Proposition 39
4	3.2 3.3 Stra 4.1	3.1.2 The Sofort Payment Process26Case Study Generalizability29Initial Data Exploration303.3.1 Transaction313.3.2 Merchant333.3.3 Consumer333.3.4 Issuer35tegy Proposition39Preventive Strategies39
4	3.2 3.3 Stra 4.1	3.1.2 The Sofort Payment Process26Case Study Generalizability29Initial Data Exploration303.3.1 Transaction313.3.2 Merchant333.3.3 Consumer333.3.4 Issuer35Itegy Proposition39Preventive Strategies394.1.1 Methods39
4	3.2 3.3 Stra 4.1	3.1.2 The Sofort Payment Process 26 Case Study Generalizability 29 Initial Data Exploration 30 3.3.1 Transaction 31 3.3.2 Merchant 33 3.3.3 Consumer 33 3.3.4 Issuer 35 tegy Proposition 39 4.1.1 Methods 39 4.1.2 Selection 41
4	3.2 3.3 Stra 4.1	3.1.2 The Sofort Payment Process 26 Case Study Generalizability 29 Initial Data Exploration 30 3.3.1 Transaction 31 3.3.2 Merchant 33 3.3.3 Consumer 33 3.3.4 Issuer 35 ttegy Proposition 39 4.1.1 Methods 39 4.1.2 Selection 41 Reactive Strategies 44
4	3.2 3.3 Stra 4.1 4.2	3.1.2 The Sofort Payment Process 26 Case Study Generalizability 29 Initial Data Exploration 30 3.3.1 Transaction 31 3.3.2 Merchant 33 3.3.3 Consumer 33 3.3.4 Issuer 35 tegy Proposition 39 4.1.1 Methods 39 4.1.2 Selection 41 Reactive Strategies 44 4.2.1 Dunning 44
4	3.2 3.3 Stra 4.1 4.2	3.1.2 The Sofort Payment Process 26 Case Study Generalizability 29 Initial Data Exploration 30 3.3.1 Transaction 31 3.3.2 Merchant 33 3.3.3 Consumer 33 3.3.4 Issuer 35 ttegy Proposition 39 Preventive Strategies 39 4.1.1 Methods 39 4.1.2 Selection 41 Reactive Strategies 44 4.2.1 Dunning 44
4	 3.2 3.3 Stra 4.2 Stra 	3.1.2 The Sofort Payment Process 26 Case Study Generalizability 29 Initial Data Exploration 30 3.3.1 Transaction 31 3.3.2 Merchant 33 3.3.3 Consumer 33 3.3.4 Issuer 35 ttegy Proposition 39 Preventive Strategies. 39 4.1.1 Methods. 39 4.1.2 Selection 41 Reactive Strategies. 44 4.2.1 Dunning. 44 ttegy Application 47
4	 3.2 3.3 Strat 4.1 4.2 Strat 5.1 	3.1.2 The Sofort Payment Process 26 Case Study Generalizability 29 Initial Data Exploration 30 3.3.1 Transaction 31 3.3.2 Merchant 33 3.3.3 Consumer 33 3.3.4 Issuer 35 ttegy Proposition 39 Preventive Strategies 39 4.1.1 Methods 39 4.1.2 Selection 41 Reactive Strategies 44 4.2.1 Dunning 44 ttegy Application 47 Preventive Strategies 47
4	 3.2 3.3 Strat 4.2 Strat 5.1 	3.1.2 The Sofort Payment Process 26 Case Study Generalizability 29 Initial Data Exploration 30 3.3.1 Transaction 31 3.3.2 Merchant 33 3.3.3 Consumer 33 3.3.4 Issuer 35 ttegy Proposition 39 Preventive Strategies 39 4.1.1 Methods 39 4.1.2 Selection 41 Reactive Strategies 44 4.2.1 Dunning 44 ttegy Application 47 Preventive Strategies 47 5.1.1 Random Forest 48

	5.2	Reactive Strategies. 5 5.2.1 Dupping	3 ⊿
		0.2.1 Dummig	-
6	Eva	luation Proposition 5	7
	6.1	Methods for Quantifying the Risk Exposure	7
	6.2		9
7	Eva	luation Application 6	3
	7.1	Determining the Aggregated Loss Distribution	3
	7.2	Evaluation of the Strategies	5
8	TPP	Networks versus Card Networks 6	9
	8.1	Costs in the Card Network	9
		8.1.1 Differentiating and Estimating Cost Components	9
	8.2	Costs in the TPP Network	1
		8.2.1 Differentiating and Estimating Cost Components	1
	8.3	Comparison	2
9	Con	clusions 7	3
	9.1	Answering the Subquestions	3
	9.2	Answering the Main Question	4
10	Dier	nuesion 7	7
10	10 1	Reflection and Limitations 7	7
		10.1.1 Conceptual Phase	7
		10.1.2 Converging Phase	8
		10.1.3 Design Phase	8
		10.1.4 Evaluation Phase	9
	10.2	Contributions	9
		10.2.1 Conceptual Phase	0
		10.2.2 Converging Phase	0
		10.2.3 Design Phase	0
		10.2.4 Evaluation Phase	0
	10.3	Recommendations for Future Research	1
		10.3.1 Conceptual Phase	1
		10.3.2 Converging Phase	1
		10.3.3 Design Phase	2
		10.3.4 Evaluation Phase	2
Α	Arti	cle 8	3
в	Soft	ware Implementation 8	9
·	B.1	Data Collection	9
	B.2	Data Analysis	9
c	الماما	al Data Exploration	
C		a Data Exploration 9	1
	0.1		ו כו
	0.2	Decision Hee	4

D	Code Implementations	95				
	D.1 Decision Tree.	. 95				
	D.2 Random Forest	. 96				
	D.3 Break Point Analysis	. 97				
	D.4 Loss Distribution Approach	. 98				
Е	Dunning Process	101				
	E.1 Notification of Retry	.101				
F	Interviews	103				
G	Glossary	105				
Bil	3ibliography 1					

Introduction

Do you want to buy a plane ticket? Just provide your credit card details online. Or do you want to check out at your favorite webshop? Just login with your PayPal account and press the *Pay Now* button. An entire financial sector has emerged to facilitate these **payment processes**: the **payment industry**. Although the payment process, also referred to as a **transaction**, may look simple to consumers like you and me, a lot of complexity is hidden behind the scenes.

A transaction typically involves four parties: the **consumer**, the **issuer** (the consumer's bank), the **merchant** and the **acquirer** (the merchant's bank) [1] — as displayed in Figure 1.1. When processing the transaction, the acquirer collects the money on behalf of the merchant at the issuer. The issuer and acquirer are linked by the **schemes**, for example provided by MasterCard and Visa. In addition, the merchant can use the **payment service provider** (PSP) to establish a connection with the acquirer [2].



Figure 1.1: Extended four party model representing the card network

Although they are not part of the traditional four party model, PSPs play an important role by providing interoperability to the industry. They provide aggregation services to merchants, enabling them to accept payments from consumers independent of the payment method they use. A **payment method** can be regarded as the way the consumer chooses to compensate the merchant for the delivered

products or services. Payment methods are offered by a variety of organizations like MasterCard, Visa, PayPal, iDEAL, etc. In this research, we will only consider payment methods for **e-commerce**.

Domination of the Card Network

The market for payment methods for e-commerce is characterized by the domination of a small number of players that operate schemes in the **card network**, like MasterCard and Visa [3]. As a result, e-commerce payment method providers are forced to use these schemes in their payment process. As argued by Anderson, this has resulted in a high concentration of market power for the scheme operators [4]. The presence of such a concentration however, is not an uncommon phenomenon in two-sided markets, such as the market for payment methods [5].

Introduction of the Third Party Payment Service Provider

In order to diversify the market, the European Commission is working on Payment Service Directive II (PSDII) [6]¹. PSDII introduces a new type of payment method provider: the **third party payment service provider** (TPP). With the introduction of the TPP a new type of network is created: the **TPP network**. TPPs are enabled to offer innovative payment methods to compete with established payment methods that are based on the card schemes.

A TPP serves as a hub between the consumer, the issuer and the PSP — as displayed in Figure 1.2. A TPP offers services based on direct access to a consumer's online banking environment. These services are based on (1) account information and/or (2) payment initiation [8]. This research focuses on the latter: the **payment initiation service** (PIS). A PIS is a payment method that enables a consumer to initiate a **real time** payment from his online banking environment, without accessing this environment himself. To do this, the consumer has to provide the TPP the required authentication to access his online banking environment. Rather than using the card schemes, as provided by MasterCard and Visa, the TPP uses the SEPA credit transfer schemes to facilitate the transaction. A further elaboration on the payment process in the TPP network, SEPA and its schemes is presented in Chapter 2.



Figure 1.2: Extended four party model representing the TPP network

¹At the start of this research PSDII was a proposal, it was adopted on the 8th of October 2015 [7]

The proposed introduction of the TPP has led to a number of early entrants into the European market like German Sofortüberweisung (Sofort) and Swedish Trustly. There are also examples of TPPs that are active outside of Europe like Australian POLi. More TPPs are expected to enter the European market shortly after the formal adoption of PSDII [9].

1.1. Research Problem and Goal

Payment methods for e-commerce broadly fall into two categories: **guaranteed** and **non-guaranteed**. Non-guaranteed payment methods offer the consumer the option to reverse a payment after it has taken place [10]. These **chargebacks** create uncertainty for the merchant whether he will actually receive the funds of the consumer after the transaction has taken place [1]. A chargeback enables the consumer to retrieve his money, for example when he is dissatisfied with the purchased product or service. However, chargebacks can also be initiated for non-genuine reasons, which enables consumers to commit fraud. It is the merchant who typically bears the financial consequences of these fraudulent chargebacks. This financial consequence arises because typically the merchant has already delivered his products or services, and gets the notification of the chargebacks afterwards. Guaranteed payment methods do not offer the possibility to issue a chargeback and provide more certainty for the merchant [11]. A further elaboration on chargebacks is presented in Chapter 2.

The new TPP networks are less mature than the established card networks that have dominated the payment industry for decades. Although they bring innovation to the market, experts from the industry [12], government [13] and academia [14] have indicated their concerns regarding the new networks around **security**, **privacy** and **chargeback handling**. If these issues are not handled properly, they can undermine the successful acceptance of TPPs by the market [5]. In this research, we will limit our scope to chargeback handling for the TPP network.

Payment methods in the card network are non-guaranteed, because of the consumer's right to issue a chargeback. Payment methods in the TPP network should be guaranteed, because the consumer does not have the right to reverse a SEPA credit transfer — the TPP scheme — after it has taken place [15]. PSPs have a responsibility towards their customers, the merchants, to provide guaranteed payment methods that are truly guaranteed — i.e. no reversals should occur. Practice has shown that merchants that make use of payment initiation services are confronted with reversals. Therefore we need to understand how big this financial risk is and what could be strategies a PSP can use to manage this risk.

With the incorporation of PSDII in the European Union member states' national legislation, banks will be legally obliged to provide the TPPs the access they need to initiate a payment on a consumer's online banking account [15] — for example by providing access to an API. But until that moment, and even after, obtaining the access to successfully initiate a payment poses its challenges. Existing TPPs like Sofort use screen scraping method to obtain the access [14].

From Chargebacks to Non-payments

From now on, we refer to reversals in the TPP network as **non-payments**, since the chargeback is a concept that originates from the card network [1]. Chargebacks in the card network are often associated with fraud [16]. Since we are not certain whether the chargebacks in the TPP network are fraud related, we decide to use the more neutral term non-payment for the TPP network. For the card network, we will keep referring to chargebacks.

Non-payment Liability

In the current situation, there is unclarity about the liability for the financial loss that results from non-payments in the TPP network. In the European Commission's PSDII impact assessment, the current liability situation is characterized as a **legal vacuum**. In practice however, it is the merchant who typically bears the financial loss [17].

As already mentioned, PSPs have the responsibility to provide guaranteed payment methods that are truly guaranteed. Therefore, this research will be conducted from the PSP's perspective. One of the benefits of taking the PSP's perspective is that its customers, the merchants, have the capability to influence the payment process which enhances the opportunity to conduct experiments during the research process. The primary goal of this research is **G1**. *To design strategies that can be used by the PSP to manage the financial risk of non-payments in the TPP network efficiently.* The secondary goal of this research is **G2**. *To evaluate whether the introduction of the TPP network helps to reach the second goal of PSDII*²: the creation of more competition within the European market for payment methods. The primary goal makes the research **relevant** from a **business** perspective, the secondary goal adds to its **societal** relevance.

1.2. Research Opportunity

Opportunities for doing quantitative research in the payment industry are rare due to legal and competition constraints [18, 19]. The majority of quantitative research in the payment industry is conducted using synthetic data. There is an ongoing discussion about the validity of research using synthetic data, for example found in the work of Barse et al. [20]. Our research, however, is based on real transaction data. This opportunity is provided by Adyen, one of the PSPs that facilitates Sofort transactions. As mentioned, Sofort is one of the first TPPs active in the European Union. Besides being one of the first, they are currently the European Union's biggest TPP. Estimates show that they processed about € in 2014. For this research, data spanning one year of Sofort transactions, which were processed on the Adyen platform, are available. The dataset contains around **Sofort transactions**. More about the **generalizability** of the used data in this research can be found in Section 3.2.

1.3. Research Gap and Questions

A lot of research has been dedicated to the card network, and into the design of strategies to detect and manage chargebacks in the network [21]. Research into the TPP network however, is practically non-existent. To the best of our knowledge, this research is the second academic effort to explore the TPP network. The first effort was conducted by Anderson, who characterized Sofort as an innovative payment method competing with the '... slow-moving cartels with high barriers to entry' that are present in the market for payment methods [4]. Anderson identified the presence of a security threat in Sofort's payment process as a result of its position as a man-in-the-middle. In Anderson's research however, no elaboration on this security threat was provided. Our research will be an effort to fill this gap. Because the novelty of the application domain of this research, it is of exploratory nature. As defined by Baxter et al. [22] an exploratory research is '... used to explore those situations in which the intervention being evaluated has no clear, single set of outcomes'. In line with our research goals, the gap as identified by Anderson and the exploratory nature of the research the research question is formulated as:

²The other goals of PSDII are presented in Section 2.1.3

RQ. How can a PSP manage the financial risk of non-payments in TPP networks efficiently, such that it becomes more efficient than the card network?

A note on Efficiency

As mentioned in the primary research goal we will (1) design strategies [...] (2) to manage [...] (3) efficiently. Efficiency refers to how well something is suited to achieve a goal, while effectiveness refers to its usefulness in achieving the goal [23]. For example, a coal plant can be very useful to generate electricity, but a windmill might be more efficient in terms of CO2-emissions. In this thesis, we will only include strategies that are capable to reduce the probability or impact of occurrence of non-payments in the TPP network. We presume all strategies will be useful as a result. The question remains, how efficient are the strategies in achieving this usefulness. That is why we focus on efficiency, from a financial point of view, rather than on effectiveness. We use the following definition of financial efficiency '... how well the dollars invested in each alternative produce revenues to the agency' [24] — in which the PSP represents the agency.

In order to answer the main research question, subquestions are formulated. The first four subquestions will be answered to achieve the primary research goal, the last subquestion will be answered to achieve the secondary research goal. Given the fact that we can both use **preventive** and **reactive** strategies to manage the risk of non-payments in the TPP network, the first question is which strategies can be used. Preventive strategies are aimed at reducing the probability of occurrence of the non-payments, reactive strategies are aimed at reducing their impact. The first subquestion is formulated as:

SQ1. What kind of preventive and reactive strategies can be used to reduce the probability or impact of occurrence of non-payments in TPP networks?

In order to test the strategies that have been selected, we will apply the strategies in practice. Because of the limited time span of our research, it is important to keep the setup of the application minimal. Therefore, the second subquestion is formulated as:

SQ2. How can the preventive and reactive strategies be applied in a minimal setup?

Now we have applied a selection of the strategies, we come to the point of evaluation. Before being able to evaluate the strategies' performance, we need to determine how the evaluation will take place. Therefore, we will investigate which evaluation methods can be used:

SQ3. What kind of evaluation methods can be used to evaluate the financial efficiency of the preventive and reactive strategies?

Given the selection of evaluation methods, we will apply them to evaluate the efficiency of the strategies that were applied:

SQ4. How can the evaluation methods be applied to evaluate the financial efficiency of the preventive and reactive strategies?

At last, we will compare the TPP network with the card network, to evaluate whether the introduction of the TPP network is beneficial to achieve the earlier mentioned goal of PSDII:

SQ5. Assuming that the most efficient strategy is deployed by the PSP, does this make the TPP network more financial efficient than the card network?

1.4. Research Methodology

We will design strategies that make use of the real life transaction data that is provided by Adyen. However, what can you do with data? According to Frawley et al. [25] data can be used to extract knowledge. This **knowledge discovery from data** (KDD) entails the '... nontrivial extraction of implicit, previously unknown, and potentially useful information from data'. The extraction concerns the identification of patterns within the data. A pattern that is interesting, according to a defined interest measure, and certain enough is called knowledge. Frawley at al. [25] argued that solely focusing on the application of **data mining** techniques without looking at the process of their application leads in many cases to sub-optimal results. To overcome this they introduced the concept of a **knowledge discovery process model** [25] that describes the different steps required to transform data into knowledge.

We will make use of a knowledge discovery process model as our research methodology. It is based on the Cross Industry Standard Process for Data Mining model (CRISP-DM) — visualized in Figure 1.3. According to a comparative study by Marbán et al. [26] CRISP-DM is the '... de facto standard' for KDD. Following the model, the first step is to obtain an *understanding of the business*. This will be obtained by presenting the case introduction in Chapter 3. In the same chapter, an initial data exploration will be presented, as part of the second step of CRISP-DM, the *data understanding*. The exploration is conducted to get a better understanding of the possible risk management strategies that can be used.

Subsequently, we move into the third and fourth steps of CRISP-DM, *data preparation* and *modeling*. In Chapter 4 we present the models that will be used as the preventive strategies, and we define a reactive strategy. This chapter aims to answer **SQ1**. concerning the strategy proposition. In the next chapter we present the application of the strategies in practice. This chapter aims to answer **SQ2**. concerning the strategy application. Next, we move into the fifth step of CRISP-DM, the *evaluation*. As part of the evaluation, we start by proposing different strategy evaluation methods in Chapter 6. This chapter aims to answer **SQ3**. concerning the evaluation method selection. Next, we evaluate the applied strategies using the strategy evaluation method of choice. This chapter aims to answer **SQ4**. concerning the evaluation method application. Subsequently, we answer **SQ5**. concerning the comparison of the TPP and card network in Chapter 8.



Figure 1.3: Research methodology with the research subquestions

1.5. Document Structure

Although we already touched upon the chapters that will answer the subquestions in the previous section, we will present the full document structure in this section. The document is partitioned in five different phases, inspired on the work of Van der Valk [27]. A schematic overview of the research approach and its representation in this document is presented in Figure 1.4.

This document continues with defining the context of the research in Chapter 2 by presenting its theoretical background. Theories from three domains are used: the payment industry, the economics of cybersecurity and knowledge discovery. The chapter concludes with defining the conceptual framework. Chapter 3 introduces the case study that will be used and relates it to similar cases in order to show the generalizability of the research. The chapter concludes with an initial data exploration. The next chapter presents a proposition of strategies that can be used in order to manage the risk of non-payments. The application of these strategies will be presented in the subsequent chapter.

In Chapter 6 different methods will be proposed for the evaluation of the strategies. The actual evaluation will be presented in the subsequent chapter. In Chapter 8 the conclusions of the strategy application and evaluation will be presented and the research questions will be answered. In the next chapter, a reflection on the research process will be presented, along with the research limitations. Also, contributions of our research will be highlighted and recommendations for future research are presented.



Figure 1.4: The elements of the research approach and the phases in which they are presented in this document

2

Background

With this chapter we move into the **conceptual phase** of our research. The chapter starts by providing theory on the payment industry, to get a better understanding of how the card and TPP networks are organized. Subsequently, theory on the economics of cybersecurity is presented that will help us in designing and evaluating risk management strategies. Third, theory on knowledge discovery is provided, for a more applied view on the strategy design. Elements from the three domains are combined in a conceptual framework, that is presented in the last section of this chapter. The conceptual framework will be used to clarify our understanding of the research problem.

2.1. Theory on the Payment Industry

In the previous chapter we introduced the four party model and distinguished the card and TPP network. In this section, we will first explain the payment process for both networks. Then we will describe the historical development of the payment industry and relate this to the structure of its underlying infrastructure. Next, an introduction to recent policy developments concerning TPPs in the European Union is presented along with its implications for the infrastructure. Building upon this, we continue by elaborating on risk management strategies that are typically used in the industry for the card network, as we might be able to learn from this for the design of the risk management strategies for the TPP network.

2.1.1. The Payment Process

The payment process can be divided in two subprocesses: **authorization** and **settlement** [1]. In the authorization process, the merchant obtains the confirmation whether the consumer has sufficient fund availability for the purchase. A positive authorization yields an **approval**, a negative authorization results in a **refusal**. In the settlement process, which takes place after an approved authorization, the funds are moved from the consumer's bank account at the issuer to the merchant's bank account at the acquirer. First we will explain both the authorization and settlement processes for the card network, then contrast it with the process for the TPP network. For the process descriptions, we assume that the merchant uses a PSP to connect with the acquirer rather than connecting with the acquirer directly — as this is the common practice in today's payment industry [1].

Payment Process in the Card Network

The payment process in the card network starts at the consumer who presents his card details to the merchant, for example by submitting these in the check out procedure in a merchant's webshop. The full process is visualized in Figure 2.1. The authorization process is captured by the first ten steps:

- 1. The consumer submits his card details to the merchant
- 2. The merchant connects with the PSP to initiate the payment
- 3. The PSP submits an authorization request to the acquirer
- 4. The acquirer sends the authorization request to the card scheme, corresponding to the scheme of the consumer's card
- 5. The card scheme submits the authorization request to the card issuer
- 6. If the authorization request is successful, the issuer creates a reservation for the due funds on the card, and returns an approval to the card scheme. If the request is not successful, the issuer return a refusal to the card scheme.
- 7. The authorization response is transmitted to the acquirer
- 8. The acquirer notifies the PSP of the authorization response
- 9. The PSP notifies the merchant of the authorization response
- 10. The merchant notifies the consumer of the authorization response

Typically, the merchant delivers his products or services after step ten and only in case the authorization was approved [1]. During the settlement process, the funds travel from the consumer to the merchant as visualized by the last five steps.



Figure 2.1: Payment process in the card network

Payment Process in the TPP network

The payment process in the TPP network also starts at the consumer. The consumer however, does not present his credit card details to the merchant, but merely indicates that he wants to pay with the TPP. The full process is visualized is visualized in Figure 2.2. The autorization process is captured by the first ten steps:

- 1. The consumer indicates to the merchant he wants to pay using the TPP of his choice
- 2. The merchant connects with the PSP to initiate the payment
- 3. The PSP sends the payment initiation to the TPP
- 4. The TPP asks the consumer for the credentials of his online banking environment
- 5. The consumer submits these credentials to the TPP
- 6. The TPP connects with the issuer's online banking environment and initiates the payment
- 7. If the payment initiation is successful, the issuer sends an approved authorization to the TPP. If the initiation is not successful, the issuer sends a rejected authorization.
- 8. The TPP sends the authorization response to the PSP
- 9. The PSP notifies the merchant of the authorization response
- 10. The merchant notifies the consumer of the authorization response

Just like for the card network, the merchant typically delivers his products or services after step ten and only in case the authorization was approved. During the settlement process, the funds travel from the consumer to the merchant as visualized by the last five steps, just like in the card network. However, for the settlement process in the TPP network, the SEPA credit transfer schemes are used, rather than the card schemes as provided by, for example, MasterCard and Visa. More about the SEPA credit transfer schemes later in Section 2.1.3.



Figure 2.2: Payment process in the TPP network

One of the main differences between the payment processes resides in the fact that in the card network the schemes are involved in the authorization, while in the TPP network, they are not. For the non-payments in the TPP network, an authorization was generated by the issuer, however, no settlement took place. To get a better understanding of the settlement process, also referred to as **clearing**, we will look into the historical development of the payment industry.

2.1.2. Historical Development

The formation of the payment industry as we know it today started in the 17th century London. As described by Quinn [28] goldsmiths introduced various forms of demandable debt, like notes, orders and bills, as a medium of exchange. These goldsmith-bankers formed the early foundation of the payment industry. As trade increased, consumers and merchants demanded for the transfer of bank notes between goldsmith-banks. A banking system evolved as individual goldsmith-banks started to accept notes of its competitors. This interbank traffic led to the development of a **banking layer**.

As the interbank traffic increased, so did the the administrative task of the banks. The banks not only had to administer their own issued notes, they also had to keep track of their competitors' notes they accepted. Simultaneously, the banks were confronted with an increased credit risk [28]. Credit risk concerns 'the risk of loss of principal or loss of a financial reward stemming from a borrower's failure to repay a loan or otherwise meet a contractual obligation' [29]. Before the interbank traffic emerged, banks were only confronted with credit risk as a result of the relation with their own customers. With the growing interbank traffic, banks became increasingly exposed to the credit risk of their competitors' customers. In order to manage the interbank activities, a **clearing layer** emerged [28].

How are the processes in the clearing layer organized? To understand this we can use a simple example. On a given day consumer A at bank A would pay a merchant at bank B 100 pounds. On that same day, consumer B at bank B would pay a merchant at bank A 80 pounds. Without clearing, these two transactions would have been executed individually. With clearing, the transactions would be consolidated: at the end of the day, only 20 pounds are transfered from bank A to bank B.

This example, however, neglects some complications such as the presence of an imbalanced credit risk of the two transactions. Let us assume that the notes issued to consumer A are subject to a higher risk than the notes issued to consumer B. Now, bank B is confronted with a higher credit risk as a result of the clearing process. To mitigate this risk, bank B can apply various strategies. One strategy is charging a risk premium in order to cover the imbalanced risk exposure. Another strategy is delaying the clearing process in order to reduce the risk. A way to do this is the T + x method in which x represents the number of days between booking a transaction and clearing [30].



Figure 2.3: Different layers in the payment infrastructure

As clearing activities developed, the clearing layer became subject to government regulation in order to control the increasing risks associated with the interbank traffic [30, p. 5]. The banking and clearing layers as they emerged in the 17th century are still part of the foundation of today's payment infrastructure, as presented in Figure 2.3. Recent policy developments in the European Union are adding another layer on top of the existing infrastructure: the **third party layer**. An elaboration on the working of this layer is presented in the following subsection.

2.1.3. Recent Policy Developments

The payment industry has over time grown into a complex network of interconnected organizations and systems [1]. In order to create a more unified and efficient payment infrastructure, the European Commission has been working on two legislative packages: Payment Service Directive I and II. This subsection introduces both packages and relates them to the development of the third party layer in the payment infrastructure.

Payment Service Directive I

In the beginning of 2014 the Single European Payments Area (**SEPA**) was introduced. One of the goals of the SEPA was to enable banks and payment services providers to expand their services with new, innovative, products and compete within the European retail payments market. The SEPA was codified in the PSDI [6]. PSDI introduced standardization rules for the European retail payment market. A universal system for (1) bank account numbers was introduced: the International Bank Account Number (IBAN), and for (2) bank identification: the Business Identifier Code (BIC) [15].

With the introduction of the SEPA two new types of payment schemes were introduced [31]. The first scheme is referred to as the **SEPA credit transfer**. Using the standardized IBAN and BIC every business or consumer can transfer funds to a counter party within the European Union. The second scheme is referred to as the SEPA direct debit. Using the standardized IBAN and BIC and the consumer's consent, the merchant is enabled to debit an amount from a consumer's bank account. The SEPA credit transfer can be regarded as the scheme that is used by the TPP network.

Payment Service Directive II

Currently, the European Commission is working on the successor of PSDI: PSDII. This new directive is expected to be adopted by the European Commission by the end of 2015 [6]. The goals of PSDII are to create (1) better integration, (2) more competition and (3) more innovation within the European payment industry [12]. The TPP is introduced as one of the means to achieve these goals. A TPP offers services based on direct access to a consumer's bank account. These services can be based on (1) account information (XS2A) and/or (2) payment initiation (PIS) [8].

A payment initiation service as provided by a TPP can be regarded as a SEPA credit transfer initiation service. What differentiates a TPP from the other parties in the industry, as displayed in Figure 1.2, is the fact that no actual money flows through a TPP. A TPP serves as a proxy of the consumer. Therefore, a TPP typically relies on the infrastructure of a PSP to provide its services. Although PSDII is not yet embedded within the European Union member states' legislative frameworks, some TPPs have emerged in its anticipation [15].

PSDI and PSDII combined: Introduction of the Third Party Layer

As mentioned before, with the introduction of PSDI and PSDII another layer is added on top of the banking layer: the third party layer. Both the consumer and merchant are enabled to connect to this layer directly. This creates the opportunity for these parties to make use of banking services without having to connect to the bank — and thus the banking layer — directly. Both PSPs and TPPs reside

in the third party layer. Figure 2.4 displays how PSPs and TPPs interact via the three layers. The TPP connects with the issuing bank in order to initiate a transaction. After the transaction is authorized by the issuing bank, the transaction is sent to the clearing agency. When the clearing agency has cleared the funds, they are transferred to, or consolidated with, the acquiring bank. The acquiring bank allocates the funds to the PSP's bank account. After all these steps have taken place, the transaction is finished and considered to be final. This finality is an important aspect, as at this point no reversal of the payment is possible anymore [1].



Figure 2.4: Position of TPPs and PSPs in the payment infrastructure

We continue this chapter by providing an introduction into risk management strategies that are used by the payment industry to deal with chargebacks in the card network. This will help us to get a better understanding of the risk challenges in the payment industry, and might help us in the process of designing risk management strategies for non-payments in the TPP network.

2.1.4. Risk Management Strategies

Chargebacks in the Card Network

A chargeback is the return of funds to a consumer initiated by the consumer [32]. A chargeback enables the consumer to retrieve his money, for example when he is dissatisfied with the purchased product or service, or when he discovers that someone else has made a transaction on his behalf. To initiate a chargeback, the consumer has to start a dispute process. This process entails the communication between the consumer, merchant and issuer about the question if the charge was legitimate. Chargebacks can be initiated for non-legitimate reasons and these are typically associated with fraud [1].

Over the last decade significant advancements are made to the **risk process** in the payment industry in order to reduce chargeback rates [33]. A combination of various techniques like data mining and more secure authentication have reduced the overall chargeback rate. Due to the growing e-commerce market however, the absolute monetary impact is increasing. A survey by Visa-owned CyberSource [34] revealed these trends for the card networks, the outcome is displayed in Figure 2.5 (adapted visualization of original).

Chargebacks, no matter if they are fraud or not fraud related, are typically regarded by the industry as a financial loss [1] and are treated as such in the risk process. The industry has adopted various strategies for chargeback handling, which are fairly industry standard according to Anderson [35] and



Relative revenue loss due to chargebacks

Figure 2.5: Historical overview of chargeback impact for the U.S. and Canadian market

Bolton [33]. Most of these risk strategies rely on KDD, as introduced in Chapter 1 and further elaborated on in Section 2.3. Parties within the payment industry store transaction information in databases on a large scale. Stored data includes attributes such as merchant details, account details, payment method, type of purchase, client name, amount of the transaction, date of the transaction, etc. [21]. This data is used by the industry to develop and apply a '... wide variety of statistical, machine learning and data mining methods' [33]. The patterns that are identified by these method can be used to block transactions that have a high risk of resulting in a chargeback. Although there is a lot of literature that describes the methods, most literature lacks detailed method descriptions [20]. One of the reasons for this is that detailed method descriptions can help fraudsters to evade detection [19].

Which methods are used for pattern identification? In a comparative study, Bolton and Hand argue that **supervised classification** and **anomaly detection** are two of the foremost used method types to create models to estimate the risk that a transaction will result in a chargeback. Supervised classification methods make use of a set of observations, transactions in our case, with known classifications and try to identify patterns that can be used to predict to which class a transaction belongs to. Anomaly detection methods are unsupervised — i.e. they do rely on previously known classifications. A further elaboration on methods for KDD can be found in Section 2.3.

Interaction of Risk Management Strategies

The different parties in the payment industry operate their own risk management systems in order to manage the risk of chargebacks. When a transaction is blocked because there is a high likelihood that it will result in a chargeback, this is referred to as a refusal. In Figure 2.6 the refusal ratio over time for an undisclosed merchant with a high volume is displayed, using anonymized data from Adyen [17]. The refusal ratio is calculated by dividing the number of transactions that are refused by the total number of transactions. The merchant started using Adyen's risk system on the first of January 2015. It can be observed that after a short increase in refusals by the PSP, the refusal ratio of both the PSP and issuer decrease. Because the PSP refuses high risk transactions before sending them to the issuer, the issuer's risk system adapts by lowering its risk thresholds. This allows more transactions to go through the system, as argued by Van der Valk [27].

Absolute revenue loss due to chargebacks



Figure 2.6: Refusal ratio of the issuer and PSP for a large volume merchant

Overall, the payment industry makes uses of risk management strategies that are aimed at the prevention of chargebacks [18]. But what are the costs of prevention, assuming that preventive strategies also block transactions that would not have resulted in a chargeback? We will continue this chapter with theory on the economics of cybersecurity, in an effort to put this trade-off in an economical perspective.

2.2. Theory on the Economics of Cybersecurity

This section aims to introduce relevant concepts in the field of economics of cybersecurity. First, security aspects of payments are presented, to get a better understanding of the threats present in the payment process. Second, an overview of information security investment metrics is presented. Third, an introduction to decision theory is provided. Theory from the last two sections will be combined to help us decide which risk management strateg(ies/y) can be used best in the TPP network.

2.2.1. Security Aspects of Payments

The concept of security is often confused with the concept of safety. As argued by Van den Berg et al. [36], security threats are caused by intentional triggered events, while safety threats are caused by unintentional triggered events. When the security threats are caused by fraudulent behavior, this can be regarded as cybercrime. As argued by Moore [37] cybercrime concerns '... any crime that involves a computer and a network, where a computer may or may not have played an instrumental part in the commission of the crime'. Anderson et al. [38] distinguish three types of cybercrime:

- Traditional crime that is now cyber because they are conducted online
- · Transitional crime which modus operandi has changed as a result of the move online
- · New crime that owes its existence to the Internet

Cybercrime in the domain of payments can be regarded as a form of transitional crime, as payment crime already existed before the move online. As mentioned before, chargebacks in the card network that are issued for non-genuine reasons are typically associated with fraud. Kahn and Roberds [11] identify three different types of payment fraud on card networks:

- **New account fraud** where the fraudster obtains someone's identity to apply for a new account. Chargebacks occur if the victim reverses payments that are made using his identity.
- Existing account fraud where the fraudster initiates a transaction using someone's account illicitly. If the legitimate account holder reverses the payments, this results in chargebacks.
- Friendly fraud, in which the fraudster orders goods and/or services and later denies having placed the order (intentionally).

All these three forms of fraud have a financial impact on the industry. However, how can this impact be measured? Detica [39] identified four different types of costs associated with cybercrime:

- Costs in anticipation of cybercrime, such as investments in anti-virus software, insurance and compliance
- Costs as a consequence of cybercrime, such as direct losses and indirect costs like weakened competitiveness as a result of intellectual property compromise
- Costs in response to cybercrime, such as compensation payments to victims and fines paid to regulatory bodies
- Indirect costs, such as reputation damage, loss of confidence in cyber transactions by individuals and businesses, reduced public-sector revenues and the growth of the underground economy

As already mentioned, various strategies are applied by the industry to manage chargeback risk in the card network. Just like for the card network, we will design and evaluate strategies for the non-payment risk in the TPP network. However, how much and where should be invested in relation to the presence of security threat(s)?

2.2.2. Security Metrics

Over the years, various models have been developed in an effort to help decision makers optimize their cybersecurity investments. One of best known models is the **Gordon-Loeb model** [40]. The model reflects the optimal investment to be made by the defender to protect against different classes of security incidents. The classes are represented in the form of security breach probability functions. Gordon and Loeb assume that the defender has a fixed security budget that has to be allocated in an optimal way over the different classes. The model shows that for one class, the optimal amount to invest in security increases with the level of vulnerability of the to be protected asset. For another class, the optimal amount to invest in security initially increases with an increase in the level of vulnerability, but eventually decreases. The Gordon-Loeb model shows that a defender should only invest a fraction of the expected value of the losses resulting from cybersecurity incidents. More specifically, the model shows that in general spending more than 37% of the expected value of the losses is uneconomical.

Another model is the **Iterated Weakest Link model** as proposed by Böhme and Moore [10], which can be regarded as an extension of the Gordon-Loeb model. The model reflects the interaction between the defender and attacker in relation to the optimal investment. Böhme and Moore assume that the defender is faced with uncertainty about where the attacker will strike — and that the attacker will repeatedly target the weakest link in the defender's defense. The model shows that even when effective defenses are available, the defender can still rationally decide not to invest in such defenses. One of the reasons for this is that the defender can learn from the attacker — by allowing a certain degree of insecurity. The defender can use this knowledge to design preventive strategies that are aimed at the attackers preferred weakest link.

Shortcoming of Cybersecurity Investment Models

The models, like those proposed by Gordon & Loeb and Böhme & Moore can help decision makers to optimize their investments. The models are limited, however, as they only consider the size of the investments and not where to invest. A model that can be used to answer this question is the Return On Security Investment (**ROSI**) model, as presented by Sonnenreich, Abanese and Stout [41]. As the authors argue: 'Before spending money on a product or service, decision-makers want to know that the investment is financially justified. Security is no different — it has to make business sense. What decision-makers need are security metrics that show how security expenditures impact the bottom line. There's no point in implementing a solution if its true cost is greater than the risk exposure.'

The ROSI model as presented in Equation 2.1 is an extension of the Return On Investment (ROI) model. In the ROSI model, the expected returns are replaced by the savings that are generated by the diminished risk exposure, as a result of the risk mitigation. So rather than returns which are considered in the ROI model, the ROSI model considers the risk cost as an opportunity cost, which can be minimized as a result of an effective risk management strategy.

$$ROSI = \frac{(Risk\ exposure\ *\ Percentage\ risk\ mitigated) - Solution\ cost}{Solution\ cost}$$
(2.1)

Sonnenreich et al. further propose that the risk exposure can be calculated by '... multiplying the projected cost of a security incident (Single Loss Exposure, or SLE) with its estimated annual rate of occurrence (ARO). The resulting figure is called the Annual Loss Exposure (ALE).' This formula is displayed in Equation 2.2.

$$Risk\ exposure = ALE = SLE * ARO$$
(2.2)

The ROSI model can be used to evaluate the efficiency of risk management strategies. However, simply comparing ROSI values of different strategies neglects the presence of uncertainty in the risk exposure. Therefore, we continue our exploration of the literature by investigating decision theory in an approach to identify how to deal with this uncertainty.

2.2.3. Decision Theory

Decision theory concerns the theory behind the decision making process [42]. But what is a decision making process? According to the Cambridge Online Dictionary¹ a decision concerns '... the choice that you make about something after thinking about several possibilities'. In a broader context, *you* can represent various entities: an individual, a firm, a public organization, etc. From now on we will refer to these entities as agents. Agents can be faced with various types of decisions. According to Wikipedia, decision can broadly be distinguished in the following categories [43]:

- **Choice under uncertainty** involves making a decision considering the expected value or utility that is accompanied with the outcome of the choice.
- **Intertemporal choice** involves making a decision where different actions lead to the realization of outcomes at different points in time.

¹http://dictionary.cambridge.org/

- **Interaction of decision makers** involves making a decision where different agents interact, this type of decisions is studied in the domain of game theory.
- **Other-regarding preferences** involves making a decision in which an agent can give up direct benefits in order to achieve a fair outcome.
- **Complex decisions** involves other types of decisions which are characterized by the complexity of the environment in which they are made.

Risk involves situations in which we do not know what the outcome will be, but we do know the distribution of the outcomes. Uncertainty involves situations in which we do not know that the outcome will be, and we also do not know the distribution of the outcomes [44]. The concepts of risk and uncertainty are closely related however. An important reason for this is the fact that you can not predict the future based on the past. Because of this, we further investigate theories behind choice under uncertainty. Because uncertainty involves risk, we will start by looking at some of the early definitions of the risk concept.

A Note on why we cannot predict the Future

In respect to this statement we follow Heisenberg in his definition of the uncertainty principle. Heisenberg argues that '... any of a variety of mathematical inequalities asserting a fundamental limit to the precision with which certain pairs of physical properties of a particle, known as complementary variables, such as position x and momentum p, can be known simultaneously' [45]. Because of the bounded scope of this research however, we will leave a detailed analysis of Heisenberg's work out of this document.

Risk according to Pascal

One of the founders of the risk concept is Blaise Pascal. In his 1670 published *Pensées* [46] he introduced the concept of expected value. He defined that expected value is '... that when faced with a number of actions, each of which could give rise to more than one possible outcome with different probabilities, the rational procedure is to identify all possible outcomes, determine their values (positive or negative) and the probabilities that will result from each course of action, and multiply the two to give an expected value' [43]. Let *E* be the expected value, x_i be the value of the outcome of *i* and $p(x_i)$ be the probability of outcome *i*, then Pascal's definition can be expressed as presented in Equation 2.3.

$$E[x] = \sum_{i=1}^{n} p(x_i) * x_i$$
(2.3)

To illustrate his concept of expected value, Pascal provided an example of a lottery, referred to as **Pascal's wager**, concerning the consequences of (dis)belief in God. In the wager, he postulates that a person's belief in God, given the fact that God exists, leads to an infinite gain. However, disbelief, given the fact that God exists, leads to a finite loss. On the contrary, belief or disbelief in God's existence, given the fact that God does not exist, leads to a finite loss respectively a finite gain. The full wager is presented in Table 2.1.

	God exists	God does not exist
Belief in God	Infinite gain, +∞	Finite loss, $-f_1$
Disbelief in God	Finite loss, $-f_2$	Finite gain, f_3

Table 2.1: Pascal's wager accompanied with values of the outcomes

Which action (belief or disbelief) will yield the highest expected value? We assign the chance p to the probability that God exists, subsequently (1 - p) represents the chance that God does not exist. The values of the outcomes as presented in Table 2.1 are used. Based on Pascal's definition of expected value as presented in 2.3, the following expected values for belief and disbelief in God can be derived:

$$E[Belief in God] = p * \infty + (1-p) * -f_1 = \infty$$
(2.4)

$$E[Disbelief in God] = p * -f_2 + (1-p) * f_3 < \infty$$

$$(2.5)$$

Pascal argues that a rational person should believe in God, because the expected value of belief (Equation 2.4) is bigger than of disbelief (Equation 2.5). Pascal's assumption that the value of an outcome is similar for everyone became subject to debate, as will be presented in the following section.

Risk according to Bernoulli

In 1738 Bernoulli published the paper *Exposition of a New Theory on the Measurement of Risk* [47] in which he argued that the expected value theory of Pascal is normatively wrong. To illustrate his claim, Bernoulli provided an example of a game of chance, referred to as the **St. Petersburg paradox**, in which at each stage a coin is tossed. The pot starts with 2 dollars and is doubled every time head appears. The first time the tail appears, the player wins the pot. The question is what the player should pay to play the game. Let *i* be the number of coin tosses, then $1/2^i$ represents $p(x_i)$ and 2^1 represents x_i . If we use Pascal's theory, then the expected value can be calculated as in 2.6

$$E = \sum_{i=1}^{\infty} \frac{1}{2^i} 2^i = 1 + 1 + \dots + 1 + 1 = \infty$$
(2.6)

Following Pascal's expected value theory, the amount a player should be willing to pay to play the game should be infinite. Bernoulli made an attempt to solve the paradox by replacing Pascal's definition of value of an outcome by utility of an outcome. He noted that '... the determination of the value of an item must not be based on the price, but rather on the utility it yields. There is no doubt that a gain of one thousand ducats is more significant to the pauper than to a rich man though both gain the same amount.'

Then the question arises, how do you define utility? Bernoulli came up with the idea that a person's utility diminishes as his wealth increases. Let the person's wealth be w, then his utility u(w) can be expressed as $\ln(w)$. Furthermore, he argues that a utility function exists whose expected net change is an indicator for the person's behavior. Let c be the cost of playing the earlier mentioned game of chance. The expected value now converges to a finite value, as expressed in 2.7.

$$E[u] = \sum_{i=1}^{\infty} \frac{(\ln(w+2^{i+1}-c) - \ln(w))}{2^i} < \infty$$
(2.7)
As mentioned before, we explored decision theory in an approach to identify how to deal with uncertainty in the ROSI's risk exposure. When we follow Pascal's theory, the risk exposure can be captured by its expected value. However, when we follow Bernoulli the risk exposure can vary per agent based on his own utility function. We will elaborate on the subjectivity of the risk exposure in Chapter 6. For now the question remains, how can we design the strategies? As already mentioned, the payment industry primarily uses preventive strategies that make use of KDD for the card network. Therefore, we will continue this chapter by presenting different types of KDD, to explore how they might be used to design preventive strategies in the TPP network.

2.3. Theory on Knowledge Discovery

As already mentioned in Chapter 1 KDD entails the extraction of patterns from data in order to capture knowledge. As argued by Phua et al. '.... fraud detection, being part of the overall fraud control, automates and helps reduce the manual parts of a screening/checking process. This area has become one of the most established industry/government data mining applications.' [18]. There are various types of data mining methods, Han et al. [48] argue that five categories can be distinguished:

- 1. **Characterization and discrimination** of classes and concepts encompasses the summarization of general characteristics of specific subsets of data. This type of analysis is exploratory of nature.
- 2. **Mining frequent patterns, associations and correlations** is about finding patterns that occur frequently in the data. A pattern consists of one (single-dimensional) or multiple (multidimensional) attributes that lead to a certain outcome.
- 3. Classification and regression for predictive analytics has its similarities with mining frequent patterns, however the focus of classification and regression is on model finding. The dataset, containing class-labeled data, is split into a training and test set. The training set is used to train a model (or so-called classifier), the test set is used to evaluate the performance of the model.
- 4. **Cluster analysis** requires no class-labeled data, unlike classification and regression. Clustering can be used to generate class labels for a group of data.
- 5. Outlier analysis can be used to find objects that do not comply with the general behavior of the data. In some applications these rare objects are more interesting than the general objects. In this context, outlier analysis is referred to as anomaly detection.

Usage of Patterns for Preventive Strategies in the TPP Network

Patterns that are obtained by using a data mining method, can be used to block transactions that are likely to result in a non-payment, before they take place. The question remains, how do you know if a data mining method will be able to produce such patterns? This is where the concept of **interestingness** comes in.

Han et al. [48, pp. 21-22] argue that one of the most important aspects of a successful data mining application is its capability to extract **interesting** knowledge. According to Han et al. two **objective measures** of interestingness can be defined. The first measure is the **support** of a pattern, the second measure is the **confidence** of a pattern. Patterns with a high support and confidence can be used to create rules that capture the knowledge. Before explaining the support and confidence measures in more detail, we have to understand the concept of an itemset.

An itemset is a group of items that (frequently) appear together. The co-occurrence of items is referred to as an association. When considering a transactional database, as present in our research,

each item can be seen as an attribute of a transaction. Examples of such attributes are the transaction *amount*, *product category* or *day of the week*. The itemsets are the different combinations of attribute values. So for example, all transactions with a value of 10 euro, concerning iTunes gift cards that were bought on a Friday share a similar itemset.

The support represents the percentage of transactions from a transaction database that share a similar itemset. In order to illustrate this concept, we take a look at the earlier mentioned itemset. Assume the transaction database contains a total of 100 transactions. Out of these 100 transactions, 30 have a value of 10 euro, concern iTunes giftcards and were bought on a Friday — i.e. they share the same itemset. The support of the itemset equals 30/100=0.30. So the support is a measure of the generalizability of the itemset to the whole dataset. The formal definition of the support measure is displayed in Equation 2.8. Both *X* and *Y* represent an itemset.

$$support(X \Rightarrow Y) = P(X \cup Y)$$
(2.8)

The confidence represents the certainty of an itemset's occurrence along with another itemset's occurrence. To illustrate this concept, we take a look at the earlier mentioned itemset. However, now we do not consider the dataset as a whole, we only consider the subset of the 30 transactions that have a value of 10 euro, concern iTunes gift cards and were bought on a Friday. Of those 30 transactions, 20 might have lead to a successful payment, however, 10 led to a non-payment. In this case, the confidence equals 20/30=0.67. The definition of the confidence measure is displayed in Equation 2.9.

$$confidence(X \Rightarrow Y) = P(X|Y)$$
 (2.9)

The support and confidence of itemsets are objective measures of the capability of a data mining method to extract interesting knowledge. In an ideal situation, we would have a data mining method that generates rules with a high support and high confidence. However, both measures can only be determined after a data mining method has been implemented and applied onto a dataset. Therefore, we cannot use these measures to determine which data mining methods to use ex ante. Next to the two objective measures, Han et al. present **subjective measures** of interestingness which are based on the users' beliefs in the data [48]. We will use these measures as criteria in our data mining method selection process. As the definitions by Han et al. leave room for interpretation, we present our definition of the criteria below.

First, the data mining method must offer the opportunity to identify **actionable** patterns. Patterns are actionable if they are able to isolate groups of non-payments from the dataset. Moreover, there should be an opportunity to implement the rules in practice. If the method is not capable to generate rules based on such patterns, we consider it not to be useful.

Second, it is important that the data mining method generates **understandable** patterns. Patterns are understandable if they are not too complex to be understood by an expert. We find it important to be able to understand why a transaction was accepted or rejected.

Third, the results of the data mining method have to lead to **useful** patterns. Patterns are useful if they enable us to distinguish between payments and non-payments. Preferably, they would relate to the security and safety threats that are present in the TPP network, to be introduced in Section 2.4.

Fourth, the data mining method has to generate **valid** patterns. In this context, we define the validity of the method as its capability to detect rules that remain valid over time. When applying data mining methods it is common practice to split the dataset in two subsets. The first part of the set is used to **train** the model. The second part is used to **test** the model's validity. In order to generalize the obtained rules, it is important to have a good performance of the model on both the training and the test set.

2.4. Conceptual Framework

Now we have presented theories from various domains, we will combine them in a conceptual framework to present our belief of the research problem. According to Shields et al. [49] a conceptual framework represents '... the way ideas are organized to achieve a research project's purpose'. A bowtie diagram as presented by Van den Berg et al. [36] will be used to design the framework, as it enables us to separate (1) threats, (2) preventive strategies, (3) incidents, (4) reactive strategies and (5) the incidents' impacts. For each part of the bowtie, as presented in Figure 2.7, an overview of the theories from this chapter that will be used in our research is presented below.



Figure 2.7: Conceptual framework using the bowtie

- Threats in the TPP network, which cause the non-payments, can be caused either intentionally or unintentionally. We follow Van den Berg et al. [36] and link these types of threats to security respectively safety. For the security threats, we adopt the three types of payment fraud as identified by Kahn and Roberds [11] — as we assume that the causes for fraud in the TPP network will be similar to those in the card network. For the safety threats, we have not been able to obtain any literature. We will refer to safety threats as technical malfunctioning.
- 2. Preventive strategies will be based on the application of data mining methods, as they enable us to derive knowledge from data. In our research we will make a selection of the five types of methods as presented by Han et al. [48], as they cover the wide variety of available methods. Because we will not be able to apply methods from each type, we will have to make a selection of the methods. Unfortunately, we will not be able to use the Han's objective interestingness criteria, as they can only be used after a method has been applied. We will use the subjective interesting criteria instead. Interesting knowledge, represented by patterns, will be used to create rules to block transactions.
- 3. **Incidents** are regarded as the non-payments that occur in the TPP network. Because of the lack of research into the TPP network, we are not able to link existing theories to their occurrence.
- 4. **Reactive strategies** will be designed in cooperation with merchants who make use of Sofort as a payment method, as theory about reactive risk management in the TPP network is non-existent.
- 5. Impact assessment of the strategies will be done using security investment models. We decide to use the ROSI model as introduced by Sonnenreich et al. [41], because this is the only identified model that can help to answer the question *where* to invest rather than *how much* as answered by the Gordon-Loeb [40] and Iterated Weakest Link [10] models. We will make use of Pascal's expected value theory [46] and Bernoulli's expected utility theory [47] in order to deal with uncertainty in the strategy evaluation process.

3

Case Study

Now we have presented our theoretical background and the conceptual framework, we move into the **converging phase** of our research. This chapter starts by presenting an introduction to the Sofort case, related to the *business understanding* step of CRISP-DM. The next section elaborates on the generalizability of the case. We continue the chapter by presenting the results of an initial data exploration, in order to obtain a better *understanding of the data*, the second step of CRISP-DM.

3.1. Case Study Introduction

Like any other TPP, Sofort is built upon the payment infrastructure of issuing banks. Therefore, we will first investigate the characteristics of the banking landscape where Sofort operates. Although Sofort originates from Germany, it offers its payment initiation services in multiple countries in the European Union. Currently, Sofort is mostly active in the so-called DACH-region. This region consists out of Germany, Austria and Switzerland. The data that will be used for this research, to be introduced in Section 3.3, contains Sofort transactions that were initiated in these countries. From the data, it can be derived that Germany accounts for 90.9% of all Sofort transactions, Austria for 8.5% and Switzerland for 0.6%. So most transactions in the DACH-region are initiated from Germany. Therefore, we start by providing a characterization of the Germany banking landscape, as this has some unique characteristics which are different from other EU member states' landscapes.

3.1.1. Characterization of the German Banking Landscape

According to a comparative study performed by the International Monetary Fund, the German banking sector is characterized by a high number of credit institutions, both in absolute numbers and relative numbers when compared to other EU member states [50]. This comparison is visualized in Figure 3.1. One of the reasons for the high number of institutions in Germany is the decentralized banking model. The German banking system comprises of three pillars, according to a study of Krahnen and Schmidt [51]. 1) *private owned commercial banks*. These banks include the country's largest banking groups and account for a total of 36% of the total assets of the German banking system, 2) *public sector banks*. These banks include the country's Sparkassen, which are founded by the government in order to '... support economic development in the respective region, and also to subsidize local public goods' [51]. Information about how large their share of the total assets of the German banking system is could unfortunately not be obtained. The third, and last, pillar consists of the *cooperative banks*.



These banks consists of networks of smaller banks that operate together under unified brand names. They account for about two third of the total credit institution figure. However, they only account for 11% of the total assets of the German banking system.

Figure 3.1: German banking sector compared to other EU member states [50]

One of the implications of the three pillars in the German banking system is a diverse national banking landscape. The various banks use different standards and operate different types of processes and systems [52]. The domain of payment processing is no exception in this respect. Traditionally, the German market for payment processing has been dominated by payment methods that offer the consumer to possibility to buy on invoice [53]. However, this is changing as an increasing percentage of German consumers buy their products online. A survey of the German Trade and Invest Agency revealed that the use of online payment service providers (e.g. PayPal, Giropay and Sofort) rose from 26% in 2011 to 40% in 2014. In the same period purchases on invoice dropped from 40% to 29% [53]. Because within this research we will focus on Sofort, the Sofort payment process will be presented in the following subsection.

3.1.2. The Sofort Payment Process

First, the Sofort payment process will be presented from the payment industry's perspective. When the consumer requests the payment initiation, Sofort connects with the issuing bank. After the consumer has enabled Sofort to initiate the transactions (this process will be presented later), the issuer has to send the transaction to the clearing layer. In the clearing layer, dedicated clearing agencies operate, like Equens. The clearing agency clears the funds and transfers them to the acquiring bank. Adyen has dedicated bank accounts to facilitate the acquiring process. After the funds are received on one of these accounts, Adyen handles the further logistics that are needed to transfer the funds to the merchant. Figure 3.2 represents the position of Sofort and Adyen in the layered visualization of the payment infrastructure.

As visualized in Figure 3.1, Germany has around 2000 credit institutions. In order to have a broad coverage of the market, Sofort should be able to connect with a large number of issuing banks to facilitate payment initiation. Many of these banks operate their own systems. In many cases, Sofort obtains access to the issuer's online banking environment by using screen scraping methods.



Figure 3.2: Position of Sofort and Adyen

From the consumer's perspective, the Sofort payment process consists of six steps. Screenshots of the steps are displayed in Figure 3.3 [17]. The steps are:

- 1. The consumer selects Sofort as the payment method to use to check out at the merchant.
- 2. The consumer **selects his bank** and provides his IBAN. Based on this, Sofort derives which authentication credentials are needed to access the consumer's online banking environment.
- 3. The consumer **provides credentials** to authenticate at his online banking environment. With these credentials, Sofort initiates a screen scraping process to obtain the access.
- 4. If the access has successfully been obtained, Sofort initiates its own risk procedure and checks the liquidity of the consumer. Although the exact working of the risk checks are bank specific, two types of checks can be distinguished:
 - (a) Fund availability checks, e.g. what is the difference between the current balance and the value of outstanding payments?
 - (b) Velocity checks, e.g. how many payments has the consumer already made since the last confirmed successful Sofort payment?
- 5. When the consumer passes the checks, the consumer has to provide additional authentication like a transaction authentication number (TAN), if the bank requires so. TANs are frequently used by European banks as a means to perform two-factor authentication [4]. With the TAN, Sofort can confirm the payment initiation.
- The consumer gets a confirmation of the payment and gets redirected to the merchant. At this moment Sofort notifies the merchant that the funds have been captured and the product/service can be delivered.



Figure 3.3: Sofort payment process

Non-payments on Sofort

The working of the Sofort payment process creates opportunities for consumer to circumvent the standard TPP payment process. Currently — for about 0.3% of all Sofort payments the funds do not arrive at the merchant. There is a suspicion that these non-payments are either caused by *technical issues* or *fraud*. There can however be overlap between these two types which makes it hard to distinguish them. One of the issues of the Sofort payment process resides in a time delay between authorization and settlement. When Sofort receives an authorization of the issuer that the payment has successfully been initiated, this does not always result in a successful settlement by the issuer. This vulnerability can be exploited by consumers, either conscious or unconscious.

3.2. Case Study Generalizability

Sofort is the largest TPP that is currently active within the European Union in terms of transaction volume, followed by Swedish Trustly. To the best of our knowledge, there are no other TPPs active in the European Union with a significant market share — although multiple TPPs are expected to enter the European market shortly after the formal adoption of PSDII [9]. The largest TPP that operates outside the European Union is Australian POLi, of which the business models and processes are comparable a TPP in the European context. Table 3.1 provides an overview of the countries in which the three mentioned TPPs operate and presents an estimation of their 2014 transaction volumes. The estimations are based on internal figures of Adyen [17].

TPP	Active markets	Size (2014 estimate)
Sofort	Germany, Austria, Switzerland,	Number of transactions:
	Netherlands, Great Britain,	Transaction volume: €
	France, Spain, Italy, Poland,	
	Slovakia, Hungary, Czech	
	Republic	
Trustly	Sweden, Norway, Finland, Den-	Number of transactions:
	mark, Poland, Italy, Spain, Esto-	Transaction volume: €
	nia	
POLi	Australia, New Zealand	Number of transactions:
		Transaction volume: €

Table 3.1: Overview of TPPs

How representative is Sofort as a TPP in the European context? Since Sofort is the biggest TPP in the European Union, we believe that our research can produce results which are representative for European TPPs in general. As already mentioned, our dataset primarily contains transactions originating from Germany. We believe this does not necessarily decreases the representative nature of our research. As mentioned before, the German banking landscape is characterized by a high number of credit institutions who use different standards and operate different types of processes and systems. We believe that the European Union's banking landscape is also highly diverse, and has its similarities with the German banking landscape as such.

How representative is Adyen's Sofort portfolio for Sofort? Given the fact that our dataset spanning one year of Sofort transactions handled on the Adyen platform contains **and the second seco**

3.3. Initial Data Exploration

This section presents the insights that were obtained from the initial data exploration. The insights obtained from the exploration will solely be used as input for the design of our preventive and reactive strategies, as will be presented in the subsequent chapters. An introduction to the data, the way it was obtained and some additional analyses can be found in appendix B. Table 3.2 provides an overview of the data that is available in the Sofort dataset. The attributes are either *transaction*, *merchant*, *consumer* or *issuer*-related. This section presents the results from the data exploration following these categories.

Category	Attributes	Format
	ID	Integer
Transaction	Date and Time	Timestamp
	Amount	Integer
Merchant	Name	String
	Name	String
	Email	String
Canaumar	IBAN	String
Consumer	Country	String
	Device type	String
	BIC	String
loguer	Name	String
Issuel	Country	String
	Settlement delay	Timestamp
	Non-payment	Boolean

Table 3.2: Attributes as present in the Sofort dataset

For some of the analyses we will make use of the non-payment ratio. This ratio enables us to contrast the occurrence of non-payments to the total number of transactions. The method the non-payment ratio is calculated, is defined in Equation 3.1.

$$Non - payment \ ratio = \frac{Number \ of \ non - payments}{Number \ of \ payments + Number \ of \ non - payments}$$
(3.1)

3.3.1. Transaction

Figure 3.4 presents the non-payment ratio over the full year the data was obtained. Four spikes in the ratio can be observed. Investigating these spikes revealed that all of them were caused by a similar type of incident. Individual merchants changed the acquiring bank account number in the configuration of the Sofort payment method. As a result of this, the funds of the transactions were not transfered to Adyen's acquiring accounts. This resulted in the reporting of non-payments that were actually not non-payments. Because these four incidents were identified and solved quickly, and it only concerned a confined number of non-payments, their impact on the overall non-payment ratio is limited.



Figure 3.4: Non-payment ratio over one year (n_{transactions}=

A more detailed date and time analysis of the non-payment ratio is displayed in Figure 3.5. The left plot shows the average non-payment ratio per day of the week. It can be observed that the ratio is higher around and during the weekend than within the middle of the week. This might be caused by technical reasons, as most banks do not settle transactions in the weekend. The right plot shows the ratio per hour of the day. It can be observed that the non-payment ratio is higher during night time than during day time. At the peak, around 3am, the ratio is twice as high as during its lowest point around noon. This might indicate fraud, as fraudulent chargebacks in the card network show an increase during night time.

In Figure 3.6 a frequency distribution of the number of transactions per amount is displayed. The amounts are displayed in ranges spanning \in 10. The figure is based on a subset of transactions that have an amount less than \in 500. No clear difference can be observed in the frequency distribution between normal payments and non-payments. Figure 3.7 displays a similar visualization, however now a subset of transactions with an amount equal to or higher than \in 500 was used. It can be observed that there are some high value transactions in the normal payments, where these were not present for the non-payments. However, this can be caused by the fact that the total number of non-payments in this subset is relatively small.



Figure 3.5: Non-payment ratio per day of the week and hour of the day (n_{transactions}=



Figure 3.6: Payments and non-payments versus amount (< € 500) (n_{transactions}=



Figure 3.7: Payments and non-payments versus amount (>= € 500) (n_{transactions}=

3.3.2. Merchant

The data revealed that there is a concentration of non-payments on specific merchants. The top ten merchants that have generated most non-payments account for a total of 7,389 non-payments while they only accounted for 2,020,410 transactions. In other words, for a displayed in Table 3.3. Most of the merchants that only account for gambling related products or services.

Merchant	Transactions	Non-payments
Merchant A	779,794	2,755
Merchant B	349,412	1,479
Merchant C	98,067	572
Merchant D	199,963	541
Merchant E	155,453	457
Merchant F	81,599	450
Merchant G	93,089	441
Merchant H	67,977	265
Merchant I	119,014	244
Merchant J	76,042	185
Sum	2,020,410	7,389
Percentage of total		

Table 3.3: Transactions and non-payments per merchant

3.3.3. Consumer

To get a better understanding of the concentration of non-payments among individual consumers, we aggregated the non-payments on the IBANs. Assuming that each IBAN represents one consumer, we visualized the number and value of non-payments versus the number of consumers in Figure 3.8. It can be observed that both distributions are right skewed, i.e. the non-payments and their monetary values do not seem to concentrate on individual consumers. This might indicate that most of the non-payments are caused by technical reasons or friendly fraud, rather than new and existing account fraud. We base this hypothesis on the work of Kahn and Roberds [11] in which they argue that new and existing account fraud are more costly for the fraudster than friendly fraud. Therefore, these two types of fraud are typically accompanied with higher gains. Therefore, we assume that if new and existing account fraud had a dominant presence in the dataset, we would see more left skewed distributions in Figure 3.8. However, we cannot prove these arguments, so the line of reasoning remains hypothetical.

A note on Independence of the Observations

The observations in the Sofort dataset are not independent as one consumer can account for multiple transactions. In the case of fraud however, specifically **existing account fraud** [11], the independence of the observations is enhanced as one IBAN does not correspond to one individual anymore. As mentioned, we cannot prove whether transactions are fraud related. Combined with the fact that the data exploration in this chapter is solely used as input for the design of the strategies, we do not regard the independence of the observations as an issue.



Figure 3.8: Number of consumers versus number and value (euro) of non-payments (n_{consumers}=

Figure 3.9 shows the non-payment ratio per device type and per country from the DACH-region. It can be observed that there is a higher non-payment ratio for mobile devices. A reason for this might be that payments initiated from a mobile device are more likely to result in technical problems. Also, it can be observed that the non-payment ratio for Switzerland is higher than for the other two countries. This might be caused by the fact that only 0.6% of all transactions that were initiated in the DACH-region are from Switzerland.



Figure 3.9: Device type and country versus non-payment ratio (n_{transactions}=

3.3.4. Issuer

For the issuers, we conducted a similar analysis as for the merchants. In order to get a comparable number of non-payments as obtained for the merchants, a selection of the top five issuers that generated most non-payments had to be made. This top five generated a total of 7,345 non-payments, while they only accounted for 969,985 transactions. In other words, for of all non-payments are generated by banks that only account for for of all transactions — as displayed in Table 3.4. So there is a higher concentration of non-payments on specific issuers than on specific merchants.

Bank	Transactions	Non-payments
Issuer A	171,095	2,502
Issuer B	195,080	2,474
Issuer C	34,897	986
Issuer D	526,895	737
Issuer E	42,018	646
Sum	969,985	7,345
Percentage of total		

Table 3.4: Transactions and non-payments per issuing bank

We plotted the average settlement delay per issuer against the non-payment ratio. The result is displayed in Figure 3.10. The dotted line represents the average non-payment ratio across the entire dataset.



There seems to be a correlation between the average settlement delay and the non-payment ratio. To test this hypothesis, we calculated the correlation between the two variables. Because of the outliers that are present in the first quadrant, we calculated Spearman's correlation. In addition, we also calculated Pearson's correlation — although this correlation coefficient is more likely to be influenced by the presence of the outliers. The coefficients are displayed in Table 3.5.

Correlation	Coefficient
Spearman	0.30
Pearson	0.24

Table 3.5: Correlation between settlement delay and non-payment ratio

Based on these coefficients we can can conclude that there is a weak positive correlation between the two variables. Besides, it can be observed that there are two clusters of issuers present in the graph, one that concentrates around a settlement delay of two days and one that concentrates around three days.





Figure 3.11: Settlement delay in business days and transaction count per issuer

Chapter Summary

In this chapter we have introduced the Sofort case, elaborated on its generalizability and presented the initial data exploration. Based on the exploration, we have identified some attributes that seem to influence the non-payment ratio, as summarized in Table 3.6. In the category *transaction* attributes, these include the day of the week and the hour of the day. For the *merchants*, we observed a concentration of non-payments among the top ten merchants that have received most non-payments. These merchants received **m** of all non-payments while they only account for **m** of all transactions. When it concerns the *consumer*, the country and device seem to influence the non-payment ratio. For the *issuers*, we also observed a concentration of non-payments among a selection of the issuers. The concentration is higher than for the merchants, we observed that the top five issuers account for **m** of all non-payment ratio and his average settlement delay. There is a weak positive correlation between these two variables.

Category	Attribute
Transaction	Day of week, hour of day
Merchant	Name
Consumer	Country, device
Issuer	Name, settlement delay

Table 3.6: Attributes that seem to influence the non-payment ratio

4

Strategy Proposition

Now we have introduced the case study, elaborated on its generalizability and presented the initial data exploration, we move into the **design phase** of our research. This chapter aims to give an answer to **SQ1**. What kind of preventive and reactive strategies can be used to reduce the probability or impact of occurrence of non-payments in TPP networks? The chapter encompasses the modeling step of CRISP-DM. The first section will introduce preventive strategies, which will be based on the application of data mining techniques. The second section presents the reactive strategies. The chapter concludes with a selection of strategies that will be applied in the subsequent chapter.

4.1. Preventive Strategies

Data mining enables us to identify patterns that might signal information about the occurrence of nonpayments. As argued by Han et al. [48], five different types of methods can be used to mine patterns from data, as introduced in Chapter 2. In the following subsection, for each of the five data mining methods, a brief classification will take place based on the subjective interestingness criteria, as introduced in Chapter 2. Using this classification we will make a selection of methods to apply in our research. Three classes to rank the methods on the criteria will be used: *yes*, *no* or *neutral*. If a method scores yes, it means that according to us the method is suitable according to the criterion. If a method scores no, we presume the method is not suitable. If a method scores neutral, we are indifferent whether to classify the method as suitable or not suitable.

4.1.1. Methods

Characterization and Discrimination

Characterization and discrimination enables us to differentiate data based on *concept descriptions*. For example, a concept description can encompass a time segregation in the dataset. In our case, we could for example contrast transactions that were initiated during the night with transactions that were initiated during the day. Given the exploratory nature of characterization and discrimination it is not certain if its application will lead to actionable results. However, since the method enables us to explore the data in all its dimensions, its results can be easily understood. Whether the results are useful is difficult to assess, finding a proper concept description that can differentiate payments from non-payments can be a matter of luck. Therefore, we presume this method not to be useful. Since characterization and discrimination is not used to create a predictive model, validity issues will not occur.

	Actionable	Understandable	Useful	Valid
Characterization and dis-	Neutral	Yes	No	Yes
crimination				

Table 4.1: Evaluation of characterization and discrimination

Mining Frequent Patterns, Associations and Correlations

Mining frequent patterns, associations and correlations enables us to explore *all patterns* that are present in the dataset. In our case, it will find all (combinations of) itemsets in the dataset. If these patterns will lead to actionable results is uncertain because there is no concept description linked to the patterns. Because of the diversity of patterns that will be found, you can not know up front if the patterns will be understandable. The method is probably useful as it will find all patterns present in the dataset. These patterns will for sure enable us to differentiate between normal payments and non-payments, however they may be too specific which can cause the problem of overfitting, as shown in the work of He and Garcia [54]. If the results are valid cannot be determined up front. One of the challenges concerning the validity is the temporal component of the data. If a pattern concentrates itself within a specific time range, this can cause an overfit based on temporal aspects, which can reduce the validity of the method usage.

	Actionable	Understandable	Useful	Valid
Mining frequent patterns,	Neutral	Neutral	Yes	Neutral
associations and correla-				
tions				

Table 4.2: Evaluation of mining frequent patterns, associations and correlations

Classification and Regression for Predictive Analytics

Methods in the category classification and regression for predictive analytics enable us to distinguish patterns based on predefined classes. In our case, these classes could comprise of the labels *payment* and *non-payment*. We cannot determine up front if the obtained patterns that distinguish between the classes yield actionable results. The same holds for the understandability of the patterns. However, the patterns will be useful as they enable us to distinguish between payments and non-payments. The usage of classification and regression for predictive analytics requires to split the dataset in a training and test set. The training set is used to train the classifier, the test set is used to validate the performance. As argued by Peng et al. [55], this ensures the validity of patterns found by this method.

	Actionable	Understandable	Useful	Valid
Classification and regres-	Neutral	Neutral	Yes	Yes
sion for predictive analyt-				
ics				

Table 4.3: Evaluation of classification and regression for predictive analytics

Cluster Analysis

Cluster analysis can be used to find groups that share itemsets. If the found patterns can be used for actionable results is unknown, just like with mining frequent patterns, associations and correlations. However, the results are in general easy to understand, as the clusters represent groups of data that comprise of the same attribute values. When it comes to the usefulness of the patterns, it remains the question if the clusters can be used to differentiate between payments and non-payments. If the results are valid is also unknown, because no usage is made of a training and test set. Also, a bad choice of the clustering metric can lead to invalid results, as argued by Bolton and Hand [56].

	Actionable	Understandable	Useful	Valid
Cluster analysis	Neutral	Yes	Neutral	Neutral

Table 4.4: Evaluation of cluster analysis

Outlier Analysis

Outlier analysis is one of the methods that is most likely to generate actionable results. Typically, fraudulent transactions are characterized by abnormal behavior of the consumer [18]. Being able to detect these outliers and acting upon them can lead to an actionable implementation. Also, when an outlier is detected, it is easy to understand the nature of the outlier. If outlier analysis will yield useful results is not known. Abnormal behavior will not specifically lead to non-payments, if there is not a dominance of fraud in our dataset. The results, however, are assumed to be valid in general, as we can make use of a split of the data in a training and test set.

	Actionable	Understandable	Useful	Valid
Outlier analysis	Yes	Yes	Neutral	Yes

Table 4.5: Evaluation of outlier analysis

4.1.2. Selection

The results of the evaluation are summarized in Table 4.6. Based upon these findings, we decide to use data mining methods based on *classification and regression for predictive analytics* and *outlier analysis* — as they score highest on our criteria. Both types of methods enable us to find valid results, i.e. the methods are capable of finding rules that remain valid over time. This is an important criterion, as it enables us to design strategies that have the potential to be robust over time. One of the additional advantages of *classification and regression for predictive analytics* is the fact that it will lead to useful results, because of the distinction it will make between payments and non-payments. Advantages of outlier analysis are that it will generate actionable results: rules to prevent outliers can easily be designed. Moreover, the results of outlier analysis are understandable, it it not hard to understand the distinction between outliers and non-outliers.

Our method selection is in line with the findings of Bolton and Hand [33]. They argue that *supervised classification* and *anomaly detection* are two of the foremost used method types to create predictive models in the payment industry. In the categorization as provided by Han et al. [48] these method types correspond respectively with *classification* and *regression* for predictive analytics and outlier analysis. However, as these method types can be implemented in various ways, the question remains, which specific methods are most suitable for our application? The following subsections present the specific methods we will use.

	Actionable	Understandable	Useful	Valid
Characterization and dis-	Neutral	Yes	No	Yes
crimination				
Mining frequent patterns,	Neutral	Neutral	Yes	Neutral
associations and correla-				
tions				
Classification and regres-	Neutral	Neutral	Yes	Yes
sion for predictive analyt-				
ics				
Cluster analysis	Neutral	Yes	Neutral	Neutral
Outlier analysis	Yes	Yes	Neutral	Yes

Table 4.6: Overview of criteria scoring

Classification and Regression for Predictive Analytics: Random Forest

In the category *classification and regression for predictive analytics* we decide to use the method random forest, as it is frequently used for pattern finding in transactional data [21]. Random forest was introduced by Breiman [57] and is visualized in Figure 4.1. A random forest classifier consists of an ensemble of decision trees that together form the forest. In a typical setup, decision trees are generated for different subsets of the training set [48]. Because each decision tree captures a different part of the data, each tree can theoretically generate another classification based on the same input. The first step of the implementation process of a random forest is the model training. During this step, the individual decision trees are generated. Aspects that influence the generation of the decision trees include the (1) method for generating the subsets of the data and the (2) attribute types that are considered in each individual tree. When the random forest is trained, we can move to the second step.

The second step is the actual classification. Let us call the attributes of a transaction that we want to classify *X*. When we ask each individual tree for a classification based on *X*, the results, let us call them *Y*, might differ. In Figure 4.1 we can see that the first two trees provide the same classification, while the third tree provides a different classification. Different methods exist that enable us to deal with this indifference. Random forest as introduced by Breiman [57] incorporates the use of the *bagging* method. The bagging method counts the different classifications of the individual trees and assigns the class with the highest count. So in our example, the predicted class of the two left decision trees will be assigned.

One of the foremost challenges of supervised learning algorithms such as random forest is the presence of imbalanced data [48, pp. 383-384]. In many applications, the allocation of observations to the classes is highly imbalanced. In our application, this is also the case as 0.3% of the transactions are non-payment, while the other 99.7% are normal payments. As we will present in the subsequent chapter, there are various strategies how to deal with this challenge. However, the application of these strategies will always lead to trade-offs to be made which decrease the performance of the classifier [54]. In the next chapter, where the application of the risk management strategies will be presented, special attention will be dedicated to the challenge of imbalanced data. The method presented in the next subsection, break point analysis, does not suffer from this challenge as it is an unsupervised method.



Figure 4.1: Visualization of random forest

Outlier Analysis: Break Point Analysis

In the category *outlier analysis* we decide to use the method break point analysis, as introduced by Bolton and Hand [56] and visualized in Figure 4.2. The concept of a break point was introduced by Senator [58]. He defines breaks as events '... which are significant leads to potentially violative behavior'. Bolton and Hand use this concept and link it to a consumer's spending behavior. They argue that consumers can be identified by their spending behavior, and if they violate their own behavior, this might be an indicator of fraud. Bolton and Hand operationalize the concept of break point analysis as represented in Figure 4.2. The method makes use of historical and new observations. By comparing the transactions amounts of both groups of observations, it can be observed whether a consumer's spending behavior is (un)usual.

To make this comparison, Bolton and Hand take 20 observations in the historical group, and four observations in the new group. They compare the means of the transaction amounts in the group by using a *t-test*. One of the advantages of this approach is the lack of need for training data. Since break point analysis is an unsupervised classification method, it does not suffer from the challenge of imbalanced data — as this is the case with random forest. One of the challenges of break point analysis though, is the question how to determine what a break point is. As Bolton and Hand do not present their implementation in detail, we will address this challenge in the next chapter where the application of the strategies will be presented.



Figure 4.2: Visualization of break point analysis

4.2. Reactive Strategies

Now we have discussed opportunities for the deployment of proactive strategies, the question remains what can we do reactive — i.e. when a non-payment already occurred? In the card network, charge-backs can only be initiated via a *dispute process*. When the consumer does not agree with the charge, he can initiate the dispute process. Depending on the specific circumstances of the dispute, it is typically the merchant who has to defend. In practice this means that the merchant has to prove that he did deliver the product or service, and that the charge of the consumer's card was legitimate [1]. If the dispute is resolved in the consumer's favor, it is the merchant who loses the dispute funds. Vice versa, if the dispute is resolved in the merchant's favor, it is the consumer who loses the funds.

Although a dispute process is operational in the industry for card networks, there is a lack of such a process for the TPP network [35]. One of the reasons for the lack of such a process might reside in the fact that in PSDII the reversal of a payment by a consumer is restricted [59, article 59]. However, since non-payments do occur in the TPP network, there seems to be a gap between the legal and operational reality. As a result, we are not able to design a strategy based on a dispute process. Therefore, we have to look at alternative reactive strategies. One of the strategies that can be used is a *dunning process*.

4.2.1. Dunning

Dunning entails 'the process of methodically communicating with customers to ensure the collection of accounts receivable' [60]. Dunning is used in the industry by payment methods that offer the consumers to possibility to buy on invoice [53], like Swedish Klarna and German AfterPay. A dunning process is initiated at the moment the consumer has not paid the outstanding invoice timely [17]. Dunning is a practical endeavor and has (not) received attention by academia to the best of our knowledge. We will therefore not use scientific theories to design the dunning strategy.

We will initiate the dunning process at the moment the non-payment occurs. If the goods or services are already delivered, the merchant has a claim on the consumer. In order to fulfill this claim, the merchant can (1) contact the consumer, (2) inform him about the unfulfilled transaction and (3) ask him to fulfill the amount due. This is exactly the process as we will apply it in this strategy. The process is visualized in Figure 4.3. In addition to this, we have to mention that until now, Sofort has blocked all consumers — by placing them on a so-called blacklist — when they caused a non-payment, as a measure to prevent fraud. However, we argue that when a consumer fulfills the due amount of the non-payment, we can consider him to be a non-fraudulent consumer. Therefore, we will not blacklist him, but place him on a whitelist instead.



Figure 4.3: Dunning process

A note on other Reactive Strategies

Besides the dunning process, more reactive strategies exist. In the most simple reactive strategy we could either accept, or deny, the occurrence of non-payments and do nothing further. Another reactive strategy would be to accept the occurrence of the non-payments and compensate the merchants for their financial impact. We decided not to design such reactive strategies, as we believe they do not adhere to the second element our research goal: (1) design strategies [...] (2) to manage [...] (3) efficiently. Why? Because we believe that management should contain elements of action, rather than elements of inaction. In line with this belief and the limited resources available for our research, we decided to only make use of the dunning process as a reactive strategy.

Chapter Summary

In this chapter we aimed to give an answer to **SQ1**. What kind of preventive and reactive strategies can be used to reduce the probability or impact of occurrence of non-payments in TPP networks? For the preventive strategies we compared five different data mining approaches as proposed by Han et al. [48]. Our evaluation showed that methods based on *classification and regression for predictive analytics* and *outlier analysis* scored best on the four pattern evaluation criteria we used: *actionable, understandable, useful* and *valid*. In the first category, we chose a random forest implementation as our strategy, as introduced by Breiman [57]. In the second category, we chose break point analysis as our strategy, as introduced by Bolton and Hand [56]. For the reactive strategies we concluded that the typical approach that is used for card networks, dispute processing, is not feasible. We therefore decided to design a dunning process as the reactive strategy.

5

Strategy Application

Now we have made a selection of strategies to manage the risk of non-payments, we will apply them in practice. This chapter aims to give an answer to SQ2. *How can the preventive and reactive strategies be applied in a minimal setup*? We will start by presenting the application of the preventive strategies in the first section. Subsequent the chapter continues by presenting the application of the reactive strategy. This chapter is a continuation of the *modeling* step of CRISP-DM, which was started in the previous chapter.

5.1. Preventive Strategies

Both preventive strategies encompass data mining techniques. In the field of data mining, a distinction can be made between *supervised* and *unsupervised* classification techniques [48]. For supervised classification, a so-called *training dataset* is used with class labeled data: each observation is accompanied with such a class. The classifier can be trained by identifying patterns that increase the likelihood of an observation to be part of a certain class [55]. When the classifier is constructed using the training dataset, its performance can be evaluated by applying the classifier on the so-called *test dataset*. This dataset contains observations for which the classes are known, however, the classifier is used to predict them without using these already known labels. The difference between the predicted classes and the actual classes of the observations in the test dataset can be used to evaluate the predictive performance of the classifier.

Unsupervised classification does not try to identify patterns based on class labeled observations. Therefore, an unsupervised classifier does not require to split the dataset in a training and test set. However, since we have class labeled data available in our research, we will evaluate the performance of the unsupervised classifier just like for the supervised classifier. In the previous chapter, the support and confidence were mentioned as interestingness measures of individual patterns. However, which measures are used to evaluate the predictive performance of a classifier as a whole?

Typical measures that are used to evaluate predictive performance are the *accuracy*, *precision* and *recall* [48]. In order to explain their definitions, we present a confusion matrix in Table 5.1. Horizontally displayed are the predicted class labels of the observations. These classes are predicted using the classifier that was trained using the training dataset. Vertically displayed are the actual classes, which were already known before the classifier was applied. To illustrate our example, we use the classes *yes* and *no*.

		Predicted class		
		Yes	Νο	Total
Actual class	Yes	True positives (TP)	False negatives (FN)	Positives (P)
	No	False positives (FP)	True negatives (TN)	Negatives (N)

Table 5.1: Confusion matrix

The accuracy is the percentage of observations for which the classifier has predicted the right class. The mathematical definition of the accuracy is presented in Equation 5.1. The precision is a measure of exactness and is defined as the percentage of predicted positive observations that are actually positive. Its formal definition is presented in Equation 5.2. The recall is a measure of completeness and is defined as the percentage of beservations that are labeled as such. The formula of the recall is defined in Equation 5.3.

$$Accuracy = \frac{TP + TN}{P + N}$$
(5.1)

$$Precision = \frac{TP}{TP + FP}$$
(5.2)

$$Recall = \frac{TP}{TP + FN}$$
(5.3)

To evaluate the performance of a classifier that is built upon imbalanced data, as present in our case, two additional measures can be used: the *sensitivity* and *specificity* [48]. The sensitivity is the fraction of predicted true positive observations of the total number of positive observations, as presented in Equation 5.4. This measure equals the recall, as the sum of true positives and false negatives is similar to the total number of positives. The specificity is the fraction of predicted true negative observations, displayed in Equation 5.5.

$$Sensitivity = \frac{TP}{P} = \frac{TP}{TP + FN}$$
(5.4)

$$Specificity = \frac{TN}{N}$$
(5.5)

Now we have presented classifier evaluation measures, we move into the application of the data mining techniques. This section will start by presenting the application of the random forest, and will conclude by presenting the application of the break point analysis. For the application of both techniques, the Sofort dataset as introduced in Chapter 3 will be used. For the application of random forest, special attention will be dedicated to the challenge of *imbalanced data*. For the application of break point analysis, special attention will be dedicated to the challenge of selecting the *break point metric*.

5.1.1. Random Forest

For the application of the random forest we make use of the findings of the initial data exploration. First, we create a *subset* of the Sofort dataset based on the concentration on specific issuers and merchants. We make use of the intersection of the top ten merchants and top five issuers that were presented in Chapter 3. This reduces the size of the dataset to a total of **sector** transactions of which

in the dataset has raised to 0.98%. This is useful, as it is a first step in dealing with the challenge of the imbalanced data. The question remains, which attributes will we select for the construction of the classifier?

Attributes

In the initial data exploration, we explored patterns in the non-payment ratio across the available attributes. From each attribute category, we include the attributes that were able to distinguish payments from non-payments best — as already presented in the summary of Chapter 3. The selection of attributes is presented in Table 5.2.

Category	Attribute
Transaction	Day of week, hour of day
Merchant	Name
Consumer	Country, device
Issuer	Name, settlement delay

Table 5.2: Attribute selection for random forest

In addition, derived attributes that can serve as an indicator of fraud, inspired by work of Bhattacharyya et al. [21], are used for the model. The underlying assumption is that fraudsters tend to clean an account as quick as they can, once they have access to someone's account. This involves making multiple and/or high value transactions in short time interval. The attributes are displayed in Table 5.3.

Attribute	Description
Is first transaction	Indicator whether transaction is first transaction of
	consumer (boolean)
Time since last transaction	Time since last transactions from consumer (integer)
Cumulative amount on this day	Cumulative amount consumer has spend on transac-
	tion day (integer)
Cumulative count on this day	Cumulative number of transactions consumer has ini-
	tiated on transaction day (integer)

Table 5.3: Attribute selection for random forest

Sampling Method

Another step in the process of dealing with the imbalanced data is reducing the imbalancedness in the training dataset [48]. To do this, various so-called sampling methods can be used. As argued by Bhattacharyya et al. [21] random undersampling of the minority class, in our case the transactions that resulted in a non-payment, can enhance the performance of the classifier. Therefore, the training set is comprised of a random selection of 50% of all transactions with non-payments. The training dataset is extended with a random selection of normal payment transactions, in such a way that the training set consists of 50% non-payments and 50% normal payments.

Implementation

For the implementation of the random forest, R was used. Appendix B presents an introduction to R and the implementation of the code in R is displayed in Appendix D. One of the shortcomings of our implementation is the lack of a genuine random sampling of the training dataset. This is caused by the fact that R, just like many other programming language, requires a predefined *random seed* in order to generate random numbers. The lack of genuine random number generation is a problem of computer science that has been present since the computer's invention [61], and it is not solved until today.

Results

One of the outputs of the random forest implementation is the analysis of importance of the used attributes, as displayed in Figure 5.1. The higher the *MeanDecreaseGini score*, the more predictive power an attribute contains. There are various indicators that can be used to measure the predictive power of an attribute besides the MeanDecreaseGini score. Because of the bounded scope of this research, however, we will not go into detail about this and use the MeanDecreaseGini score as it is the standard measure using in R's *randomForest* library. From the obtained scores, it can be observed that both the issuer and the settlement delay of the issuer have a large MeanDecreaseGini score.



Figure 5.1: Importance of the random forest attributes

The confusion matrix of the random forest implementation is displayed in Table 5.4. Next to the transaction counts, the monetary values of the affected transactions are displayed, underneath the transaction counts. This monetization of the classifiers performance will enable us to calculate the cost of using random forest as a strategy, later in the research process.

		Predicted class		
		Non-payment	Payment	Total
	Non-payment	728	2,136	2,864
Actual class		€ 21,791	€ 68,465	€ 90,256
	Payment	165,842	412,000	577,842
		€ 5,578,864	€ 12,616,008	€ 18,194,872

Table 5.4: Confusion matrix random forest

Table 5.5 presents the values of the classifier evaluation measures that were introduced before. It can be observed that the accuracy is 0.71, which means that 71% of all observations in the test set are correctly classified. The precision however, is only 0.0044, which means that only 0.44% of all observations that are predicted to be non-payments, are actual non-payments. The recall and sensitivity are both 0.25, which means that 25% of all non-payments are predicted as such. When it comes to the normal payments, the specificity shows us that 71% of all non-payments are classified as such. So overall, the random forest classifier we constructed is better capable of predicting the occurrence payments than non-payments, which is logical given the imbalancedness of the data.

Measure	Value
Accuracy	0.71
Precision	0.0044
Recall	0.25
Sensitivity	0.25
Specificity	0.71

Table 5.5: Evaluation measures random forest

5.1.2. Break Point Analysis

In contrast to random forest, break point analysis is a form of *unsupervised classification*. This has the implication that the method does not rely on class labeled data for training purposes. However, just like for the random forest we decide to take a *subset* of the original Sofort dataset. We follow the approach as presented by Bolton and Hand [56], in which they use a moving window consisting of 20 historical observations and 4 new observations. Bolton and Hand do not present any argumentation why they make use of this size of the moving window. We decide to follow their approach rather because of a lack of a better alternative, than our belief in the fact that the proposed size of the moving window is an optimal configuration. This should be kept in mind while interpreting the results of our implementation. We discard all transactions from consumers that initiated less than 24 transactions, as we will not be able to detect any break points for them because *the moving window cannot be moved*. This reduces the dataset to a subset containing **moved** transactions of which **moved** are non-payments. As a result the non-payment percentage has decreased to 0.16%.

Break Point Metric

As mentioned before, a challenge in break point analysis is the selection of the break point metric. Bolton and Hand determined the break point by using a t-test. However, they do not present which t-test they use. Therefore, we decide to select a t-test ourselves. Because we want to compare the means of two (presumed) independent groups, we have to select an *independent two-sample t-test*. Within this category, Hair et al. [62] present four different tests that can be used. The factors that distinguish the selection of the correct test are the *(in)equality of the sample sizes* and the *(in)equality of the variances*. The different tests than can be used are displayed in Table 5.6.

	Equal sample sizes	Inequal sample sizes	
Equal variances	Normal independent two-	Independent two-sample t-	
	sample t-test or independent	test with standard deviation	
	two-sample t-test with stan-	correction	
	dard deviation correction		
Inequal variances	Welch's t-test	Welch's t-test	

Table 5.6: Independent two-sample t-tests

In our implementation of break point analysis, we have different sample sizes for the historical and new observations of 20 respectively 4 transactions. Besides, it is uncertain if the observations in the two groups have equal variances. Therefore, we cannot assume variance equality. Therefore, we will use *Welch's t-test*.

Implementation

Just like for the implementation of the random forest, the R code for our break point analysis implementation is displayed in Appendix D. One of the challenges of the implementation resides in the limited computational power we had available. A single run of the break point classifier using an 2.7 GHz Intel Core i5 processor with 8 GB 1867 MHz DDR3 ram took 16 hours and 23 minutes until completion.

Results

A graphical impression of the break point analysis is displayed in Figure 5.2. The four plots represent four individual consumers' spending behavior. Each dot is one transaction: the green dots represent normal payments and the red dots represent non-payments. The vertical lines are the break points that were identified by the classifier. It can be observed that in some cases, a break point is followed by non-payments, and is as such a valid signal. There are also break points however, that do not indicate the occurrence of non-payments: they result in false positives.

Table 5.7 presents the confusion matrix of the application of break point analysis. Just like for the random forest classifier, the number of transactions is displayed as well as its monetized value. It can be observed that the number of true positives is smaller than for the random forest classifier, while the number of total observations in the break point analysis is larger. The question is, how does this affect the classifier evaluation measures?

		Predicted class		
		Non-payment	Payment	Total
	Non-payment	141	1,942	2,083
Actual class		€ 8,662	€ 76,821	€ 85,482
	Payment	57,819	1,239,232	1,297,051
		€ 3,766,977	€ 44,768,192	€ 48,535,169

Table 5.7: Confusion matrix random forest



Figure 5.2: Visualization of break point analysis result

The values of the classifier evaluation measures are displayed in Figure 5.8. In the table, we have also included the measure values of the random forest classifier — in order to make a comparison. It can be observed that the break point classifier has a higher accuracy than the random forest classifier. This means that the break point classifier has correctly classified a larger proportion of the observations. However, the random forest classifier has a higher precision, meaning that it is better in predicting the class right. When it comes to recall and sensitivity, the random forest classifier scores better than the break point classifier. In other words, the random forest classifier is better in predicting the class of the non-payments with respect to the total non-payment observations. In terms of specificity the break point classifier performs better, as it is better capable of predicting the negative — normal payment — class.

Measure	Value	Value
	Random forest	Break point analysis
Accuracy	0.71	0.95
Precision	0.0044	0.0024
Recall	0.25	0.068
Sensitivity	0.25	0.068
Specificity	0.71	0.96

Table 5.8: Evaluation measures random forest and break point analysis

5.2. Reactive Strategies

As introduced in the previous chapter, there are not formalized dispute processes for the TPP network. Therefore, this section presents the implementation of the reactive strategy that was performed, the dunning process.

5.2.1. Dunning

Because the reactive nature of the dunning process, no predictive classification had to be made in order to obtain transactions that resulted in a non-payment. For the dunning process, the non-payments could simply be queried from the Sofort dataset. However, the implementation of the dunning process interferes with the actual business process of both Adyen and its customers: the merchants. Therefore, collaboration with merchants had to be established to execute this strategy.

Implementation

The first step of implementing the dunning process, was getting the participation of merchants. In order to obtain a selection of merchants to approach, a list of the number of non-payments per merchant for the last month of the Sofort dataset — July 2015 — was obtained. The top five merchants with most non-payments were approached with a participation request. After a process of alignment between Adyen and the merchants, one merchant was willing to participate in the experiment. The merchant is active in a niche that sells digital gift cards. You could argue that this decreases the representative nature of the experiment — and indeed, it does. However, various studies have indicated that merchants who sell goods that are easy to *cash out* — like digital gift cards — are prime subjects of fraudsters [63]. Subsequently, if our experiment indicates that we are able to conduct a successful dunning on the transactions of the merchant in our experiment, the non-representative nature of our experiment will only serve as an additional indicator that there is no dominant presence of fraudulent non-payments in the Sofort dataset.

For the implementation of the dunning, the process as described in Chapter 4 was executed. After obtaining the notification of the non-payment on the Adyen platform, a notification of the non-payment was sent to the consumer by email. In the same message, the consumer was notified of the upcoming retry using a SEPA direct debit. The message of the email that was used for this is presented in Appendix E.

After sending the email, the next step was to initiate the retry using a SEPA direct debit. While the SEPA credit transfer — that is used for payments in the TPP network — is a *push* payment from the consumer's perspective, the SEPA direct debit concerns a *pull* payment. This implies that no authentication of the consumer is needed in order to debit his account. It must be mentioned that because of this reason, the consumer always remains the right to reverse a SEPA direct debit with a chargeback. From the data available on the Adyen platform, it could be observed that in general within 18 days after initiation of a SEPA direct debit that a consumer reverses with a chargebacks have been received [17]. For each SEPA direct debit that a consumer reverses with a chargeback, the issuer holds the right to charge the initiator of the SEPA direct debit a chargeback fee. This fee depends on the contractual relationship between the initiator of the SEPA direct debit, in this case the PSP, and the issuer. We presume that this fee equals € 7,50 per SEPA direct debit chargeback.

Results

In the dunning process, the non-payments were followed up with a maximum of two retries. The first retry was executed for all 51 Sofort non-payments, the second retry was executed for the retries that itself resulted in a chargeback. After the retries, the funds of a total of 40 Sofort non-payments were recovered. However, as mentioned before, the consumer can reverse a SEPA direct debit with a chargeback. Because of this reason, we verified if the SEPA direct debit was chargebacked four weeks after its initiation. Since this exceeds the 18 days as mentioned before, we have a 99%+ likelihood that the SEPA direct debits will not result in a chargeback as such. For the remaining 11 Sofort non-payments, the funds could even not be captured after initiating the two SEPA direct debits. The total cost of the chargeback fee for the experiment was € 165. The overall results of the dunning process

are displayed in Table 5.9. Just like in the confusion matrices, both the number of transactions as well as their monetary impact are presented.

		Success		
		Yes	No	Total
	Non novmont	40	11	51
Actual class	Non-payment	€ 1,394	€ 432	€ 1,826

Table 5.9: Performance matrix of dunning

Chapter Summary

In this chapter we have aimed to answer SQ2. *How can the preventive and reactive strategies be applied in a minimal setup?* We have applied three different strategies: random forest and break point analysis as the preventive strategies, and the dunning process as the reactive strategy. The preventive strategies have been applied on the Sofort dataset. To deal with the challenge of imbalanced data for the random forest implementation we used three tactics. First, we reduced the dataset to the intersection of the top five issuers and top ten merchants who caused most chargebacks. Second, we balanced our training dataset by applying random undersampling of the minority class, as presented by Bhattacharyya et al. [21]. And at last, we used the bagging method which is embedded in the random forest method as introduced by Breiman [57]. To deal with the challenge of different break point metrics in the break point analysis implementation, we decided to use Welch's t-test as this test is capable of dealing with (1) different sample sizes and (2) both equal and unequal variances. For the reactive strategy an experiment was conducted with the cooperation of one merchant.
6

Evaluation Proposition

In order to evaluate the efficiency of the different strategies we will use the Return On Security Investment Model, as introduced in Chapter 2. However, most applications of the ROSI model assume the risk exposure to be static. One of the drawbacks of using a static risk exposure is that it can lead to misunderstanding and misinterpretation of the model's outcomes [64]. Because our research investigates the innovative setting of the TPP network, we want to include a variable risk exposure. This chapter aims to give an answer to SQ3. What kind of evaluation methods can be used to evaluate the financial efficiency of the preventive and reactive strategies? In the first section we present methods that can be used to quantify the risk exposure. In the second section we select which method we will use. With this chapter we move into the **evaluation phase** and the *evaluation* step of CRISP-DM.

6.1. Methods for Quantifying the Risk Exposure

How can we quantify risk exposure? Let us first examine the quantification of the **annual loss exposure**, as introduced in Chapter 2. The ALE is the proposed method to quantify risk exposure by the originators of the ROSI model, Sonnenreich et al. [41]. The ALE was first defined by the National Bureau of Standards in 1979 as the product of the impact of a set of incidents with the set of the corresponding incidents' frequencies [65], as displayed in Equation 6.1.

$$ALE = \sum_{i=1}^{n} I(O_i)F_i$$
where $\{O_1, ..., O_n\}$ = Set of harmful incidents
$$(6.1)$$

 $I(O_i) =$ Impact of incident i

 F_i = Frequency of incident i

This definition of the ALE produces a single outcome and ignores the presence of variability in the risk exposure, for example caused by uncertainty in the ALE's input variables. Before examining how we can incorporate this variability we will consider the types of risk that exist, as the different types require different quantification methods. According to Frachot et al. [66] three main types of financial risk can be distinguished for financial institutions, like PSPs:

- **Credit risk** 'The risk of default on a debt that may arise from a borrower failing to make required payment.' [67].
- **Operational risk** 'The risk of direct or indirect loss resulting from inadequate or failed internal processes, people and systems or from external events.' [66]
- **Market risk** 'The risk of losses in on- and off-balance sheet positions arising from movements in market prices.' [68]

As we incorporated in our conceptual framework, the non-payments are either caused by fraud or technical related reasons. We argue that this falls in the definition of *failed internal processes, people and systems or from external events*. Therefore, the non-payment risk we consider in our research is considered to be a type of operational risk. The question remains, how to quantify operational risk for financial institutions?

The Basel II Accord [69] proposes three method types for calculating operational risk exposure for financial institutions. The first type is the **basic indicator approach**. In this approach, the risk exposure is calculated as a percentage of the average annual gross income of an organization. According to the Basel II Accord this percentage is 15%. As a result, the basic indicator approach also provides a static risk exposure, just like the ALE.

The second method type is the **standardized approach**. In this approach, also a percentage of the average annual gross income is used to represent the risk exposure. However, in the standardized approach different percentages for different types of business lines j can be used. A business line comprises of a group of closely related business activities. A list of the business lines as defined under Basel II is displayed in Table 6.1.

Abbreviation	Business line
CF	Corporate Finance
TS	Trading and Sales
RB	Retail Banking
СВ	Commercial Banking
PS	Payment and Settlement
AS	Agency Services
AM	Asset Management
RB	Retail Brokerage

Table 6.1: Business lines j

The third method type is referred to as the **advanced measurement approach**. This approach enables financial institutions to design their own risk exposure quantification methods based on the frequency and severity of the occurrence of historical incidents. According to a comparative study by Ames et al. [70] the advanced measurement approach enables organizations to quantify the risk exposure based on '... sophisticated internal models'. Within the method, organizations are not only enabled to differentiate risk exposure based on different business lines, they can also differentiate based on loss type *i*. An overview of the loss types that are defined under Basel II are displayed in Table 6.2.

Abbreviation	Loss type
IF	Internal Fraud
EF	External Fraud
EPWS	Employment Practices and Workplace Safety
CPBP	Clients, Products and Business Practices
DPA	Damage to Physical Assets
BDSF	Business Disruption and System Failures
EDPM	Execution, Delivery and Process Management

Table 6.2: Loss types i

6.2. Risk Exposure Method Selection

The annual loss exposure and the basic indicator both neglect variability as they only produce one figure to represent the risk exposure. The standardized approach can incorporate variability, as a distinction can be made between different business lines. However, the risk exposure per business line is also static. The advanced measurement approach enables to incorporate most variability of all four methods, as (1) an additional distinction can be made based on loss type and (2) frequency and severity data of the occurrence of historical incidents can be incorporated. Therefore, we decide to use the advanced measurement approach. An overview of the methods is presented in Table 6.3.

Within the advanced measurement approach we will consider the business line PS (payment and settlement) as all steps in the payment process, as presented in Chapter 2, can be categorized as such. Besides, we will have to determine which loss types to consider. Following our conceptual model, the non-payments are either caused by EF (external fraud) or BDSF (business disruption and system failures). Since we cannot differentiate the cause of the non-payments based on the information we have available, we will consider them all to be part of one loss type.

Method type	Risk exposure
Annual loss exposure	Static
Basic indicator approach	Static
Standardized approach	Static & variable
Advanced measurement approach	Variable

Table 6.3: Method types for quantifying risk exposure

Loss Distribution Approach

The question remains, how can we operationalize the advanced measurement approach? According to a comparative study by Frachot et al. [66] most methods to operationalize the advanced measurement approach estimate the **Value at Risk** (VaR). Two methods to calculate the VaR are the **Internal Measurement Approach** (IMA) and the **Loss Distribution Approach** (LDA). IMA is the classic method that has been long in favor of the Basel Committee on Banking Supervision [69]. However, according to Frachot et al. LDA performs better for institutions who '... strongly favor risk sensitive/quantitative methods' [66]. Because we want to incorporate the variability of the risk exposure in our analysis, we decide to use the LDA. The concept of LDA is based on the idea that both the frequency and the severity of incidents follow a probability distribution. When these two distributions are multiplied, the aggregated

loss can be obtained which represents the VaR. This aggregated loss itself follows a probability distribution, which is based on the two input distributions. The conceptual model of LDA is displayed in Figure 6.1.



Figure 6.1: Conceptual model of the loss distribution approach (LDA)

Following the New Base Capital Accord [69], a distinction is made between business lines *i* and event types *j*. As already mentioned, we assume that the events (non-payments) can be considered to be one business line and one event type. Let variable N(i, j) represent the frequency of non-payments. Then, we denote the probability function of N(i, j) as $P_{i,j}$. Let variable $\zeta(i, j)$ represent the loss severity of the non-payments. Then, we denote the probability function of $\chi(i, j)$, which can be calculated as the sum of severities for all frequencies — as displayed in Equation 6.2.

$$\vartheta(i,j) = \sum_{n=0}^{N(i,j)} \zeta_n(i,j)$$
(6.2)

In order to obtain the aggregated loss probability distribution function, G(i, j), we need to compound the frequency and severity probability distribution functions. There is no analytical expression of this compounded probability distribution function [66]. Therefore, we need to use a numerical solution to estimate the function. A common way to do this, is by applying a **Monte Carlo simulation**. In a Monte Carlo simulation, a process can be simulated many times with varying start conditions. The result is a set of simulations which cover most of the outcomes that the simulation can produce [71]. When we have used the simulation to obtain the aggregated loss distribution, the question remains, how can we interpret this distribution?

Interpretation of the Aggregated Loss Distribution Function

Figure 6.2 represents one of the possible aggregated loss distributions that can be obtained from the Monte Carlo simulation, with the embedding of the VaR. We assume that the VaR at a confidence interval of 50% accounts for the expected losses, in line with Pascal's expected value theory introduced in Chapter 2. Everything above the 50% confidence interval accounts for unexpected losses. Which confidence interval should be used however, is up to the subjective choice of the decision maker, according to Bernoulli. We presume that if the decision maker is risk averse, he should opt for a VaR at a confidence interval above 50%. If the decision maker is risk willing, he should opt for a VaR at a confidence interval below 50%. If the decision maker is risk neutral, the VaR at the 50% confidence interval should be used.



Figure 6.2: Aggregated loss distribution function G(i, j) with Value at Risk (VaR)

Chapter Summary

This chapter aimed to answer **SQ3**. *What kind of evaluation methods can be used to evaluate the financial efficiency of the preventive and reactive strategies?* We identified that the ROSI model typically uses a static risk exposure, the annual loss expectancy. In order to explore how to incorporate a variable risk exposure, we explored different types of risk. We have identified that the risk of non-payments can be categorized as a type of operational risk. To quantify the exposure of such a risk for financial institutions, the Basel II Accord [69] provides the standardized approach, the basic indicator approach and the advanced measurement approach. We have selected the advanced measurement approach as this approach allows us to incorporate the most variable risk exposure. For the operationalization of the advanced measurement approach, we have selected the loss distribution approach as presented by Frachot et al. [66]. In this approach we are enabled to quantify the risk exposure based on historical frequency and severity distributions of the occurrence of non-payments. The risk exposure will be quantified as the aggregated loss distribution function, which presents the Value at Risk (VaR) for different confidence levels.

Evaluation Application

Now we have introduced an evaluation method that enables us to quantify a variable risk exposure, we will apply it to evaluate the risk management strategies. This chapter aims to give an answer to SQ4. *How can the evaluation methods be applied to evaluate the financial efficiency of the preventive and reactive strategies?* The first section describes the process of obtaining the aggregated loss distribution. The next section uses this distribution to obtain the ROSI for the different strategies. This chapter is a continuation of the *evaluation* step of CRISP-DM, which was started in the previous chapter.

7.1. Determining the Aggregated Loss Distribution

In order to determine the aggregated loss distribution, we first need to determine the loss frequency and severity distributions. Because the Sofort dataset contains loss data for a full year, we need to partition the dataset in smaller periods of time. We decide to split the dataset in weeks, so we obtain 52 time periods with loss data. The frequency and severity of the non-payments are captured in the probability density histograms as displayed in Figure 7.1. To keep the visualizations interpretable, we decided to use a class width of 20 transactions for the frequency and a class width of € 10 for the severity.



Figure 7.1: Loss frequency and severity distributions

A Note on Partitioning the Data and Fitting the Distributions

After visualizing the distributions, the question rises, which probability distributions provide the best fit on the data? For the loss frequency a normal distribution with a mean of 165 and standard deviation of 40 was fitted on the data. The loss severity was fitted with a gamma distribution with a shape of \notin 2.5 and a rate of \notin 0.1 (i.e. scale of 1/0.1=10). The two distributions that were fitted are displayed in red, their mean — or expected — values are indicated with a green line. A discussion about the representativity of fitted the distributions, and their relation to the partitioning of the data, can be found in Chapter 10.

To obtain the aggregated loss probability distribution function, G(i, j), we need to compound the frequency and severity probability distribution functions. As elaborated in Chapter 6, there is no analytical way of expressing this compounded function and we will use a Monte Carlo simulation to estimate it. Our implementation of the Monte Carlo simulation is displayed in Listing 7.1. We iterate 1,000,000 times. First, we take random observations of the 1,000,000 frequencies stored in the loss_frequency_distribution. For the obtained frequency count, we iterate. Then we obtain an observation of the 1,000,000 severities stored in the loss_severity_distribution. The outcome of this process is the aggregated_loss_distribution with 1,000,000 observations. The full implementation of the process we applied to obtain the aggregated loss distribution can be found in Appendix D.

```
# Performing the Monte Carlo simulation
 for(i in 1:1000000) {
2
   draw number<-round(runif(1, 1, 1000000))</pre>
3
   draw frequency<-loss frequency distribution[draw number]</pre>
   cumulative severity <- 0
5
   for(j in 1:draw frequency) {
6
      cumulative severity <- (cumulative severity + loss severity distribution [
         round(runif(1, 1, 1000000))])
    }
8
    aggregated loss distribution <- rbind (aggregated loss distribution, data.
9
       table(c(cumulative severity)))
10 }
```

Listing 7.1: Monte Carlo simulation implementation in R

Figure 7.2 shows the obtained aggregated loss distribution. The red line represents the normal distribution that was fitted on the data, with a mean of \in 4,100 and standard deviation of \in 1,000. The green line represents the VaR of the distribution at a confidence interval of 50%. Moving the line to the left decreases the confidence interval of the VaR, move the line to the right increases the confidence interval.



Figure 7.2: Aggregated loss distribution

Expected Value versus Expected Utility

If we follow Pascal's expected value theory, the expected aggregated loss is the mean of the normal distribution: \in 4,100. If we follow Bernoulli's expected utility theory, the expected aggregated loss depends on the agent's utility function. We argue that this function should depend on the confidence of the VaR. As argued in Chapter 6 a risk averse agent should opt for a confidence interval above 50%, a risk neutral agent should opt for 50% and a risk willing agent should opt for less than 50%. This would respectively result in a aggregated loss — or VaR — lower than, equal to and higher than \notin 4,100.

7.2. Evaluation of the Strategies

Now we have obtained the aggregated loss distribution, we still need to obtain the costs and risk mitigation of the strategies to apply the ROSI model. In this research, we only consider direct costs, as indirect costs are hard to quantify. Also, we don't include implementation costs, as these can only be based on rough estimates which will not add to the validity of the evaluation of the strategies in our opinion. For the preventive strategies, random forest and break point analysis, we use the monetized value of the false positives as the direct cost. For the risk mitigation of the preventive strategies, the sensitivity (or recall) of the classifiers will be used — as this measure represents the percentage of predicted positive observations of the actual positive observations. For the reactive strategy, dunning, we use the issuer's fee for the non-successful retries as the direct cost and the monetized value of the successful retries as the direct benefit. The risk mitigation is the success ratio of the dunning process. An overview of the direct costs and risk mitigation for the strategies is presented in Table 7.1.

	Random forest	Break point analysis	Dunning
Direct costs	€ 5,578,864	€ 3,766,977	€ 165
Risk mitigation	0.25	0.068	0.76

Table 7.1: Direct costs and risk mitigation per strategy

To evaluate the strategies, we calculate the ROSI. As introduced in Chapter 6 we replace the risk exposure, typically represented by the ALE, with the aggregated loss distribution function. The result is a ROSI value that differs based on the confidence interval of the aggregated loss distribution. The ROSI and the incorporation of the aggregated loss distribution function are displayed in Equation 7.1. The risk mitigated and the direct cost can be derived from Table 7.1.

$$ROSI = \frac{(Aggregated \ loss \ distribution * Risk \ mitigated) - Direct \ cost}{Direct \ cost}$$
(7.1)

The obtained ROSI curves, with the confidence of the VaR on the horizontal axis, are displayed in Figure 7.3. The dotted lines can be used to obtain the ROSI value at a confidence interval of the VaR of 50%. The left plot displays the ROSI curves for all three strategies. The ROSI curves for the preventive strategies overlap, and are both close to minus one. This indicates that for each euro invested in the strategies, a loss close to one euro will occur — i.e. there is no positive business case. The ROSI curve for the reactive strategy, the dunning process, shows ROSI values higher than one. At a 50% confidence of the VaR, the ROSI equals four. This implies that for each euro invested in the dunning process, four euro in revenue will be obtained.

The right plot merely shows the ROSI curves of the two preventive strategies, to compare their financial efficiency. From the curves we can derive that the random forest performs better than the break point analysis. This can be explained by the fact that the risk mitigation of the random forest is higher than the risk mitigation of the break point analysis.



Figure 7.3: ROSI curves for the risk management strategies

Now we have selected and applied the risk management strategies, and we have selected and applied the strategy evaluation methods, we come to an end to the contribution to our primary research goal **G1**. To design strategies that can be used by the PSP to manage the financial risk of non-payments in the TPP network efficiently. In the next chapter we will compare the costs of a transaction for a merchant in the card and TPP networks, assuming that the strategy that has proven to be financial most efficient will be implemented: the dunning process.

A Note on the Strategy Evaluation

When comparing the ROSI curves of the different strategies, we can conclude that the dunning process is most financial efficient, no matter what confidence interval of the VaR is selected. However, we have to keep in mind that (1) we only applied three strategies — while other strategies could also have been applied — and (2) that the implementation of the strategies might not be optimal. As such, we can only conclude that *within the boundaries* of our strategy design and evaluation process, the dunning process yields most financial efficiency.

Chapter Summary

In this chapter we have aimed to answer SQ4. How can the evaluation methods be applied to evaluate the financial efficiency of the preventive and reactive strategies? We have applied one of the methods available in the advanced measurement approach to determine the risk exposure: the loss distribution approach. To do this, we first partitioned our dataset in periods of one week. Using the obtained subsets, we determined the loss frequency and severity distributions. Using a Monte Carlo simulation to multiply the two distributions, we obtained the aggregated loss distribution. This distribution was used as the risk exposure in the ROSI calculations to compare the different risk management strategies. We have concluded that within the boundaries of our research, the reactive strategy comprising the dunning process has a positive business case, because the ROSI values are higher than one, while the preventive strategies random forest and break point analysis result in a loss and are as such not worth the investment.

8

TPP Networks versus Card Networks

In the previous chapter we have concluded that the reactive strategy comprising the dunning process is financial most efficient to manage the non-payment risk in the TPP network. In this chapter we will compare the costs of the TPP and card network, under the assumption that the dunning process would be implemented. By doing so, we will answer **SQ5**. Assuming that the most efficient strategy is deployed by the PSP, does this make the TPP network more financial efficient than the card network? As such, we will be able to reach the secondary goal of our research **G2**. To evaluate whether the introduction of the TPP network helps to reach the second goal of PSDII: the creation of more competition within the European market for payment methods.

We presume that for the TPP network to become competitive, the costs of a transaction for the merchant should be equal to or less than the costs of using the card network. Therefore, we will estimate these transaction costs for the merchant. In the first section of this chapter we will differentiate the cost component for the card network, and provide estimations of their values. In the subsequent section, we will do the same for the card network. In the third section we will make the comparison between the transaction costs for the merchant in the two networks.

8.1. Costs in the Card Network

The transaction cost for a merchant in the card network depends on numerous factors. In order to provide an accurate estimation of the transaction cost, we will limit the number of factors to consider. As mentioned before, the market for payment methods for e-commerce is characterized by the domination of a small number of players that operate schemes in the card network, like MasterCard and Visa [3]. Therefore, we will consider costs imposed by these two scheme operators. Besides, we will only focus on transaction costs within the European Union, the jurisdiction of PSDII. Third, we will only include costs of so-called two-leg transactions, i.e. transactions for which both the issuer and the acquirer are located within the European Union.

8.1.1. Differentiating and Estimating Cost Components

There are five main cost components of a transaction in the card network [15], as visualized in Figure 8.1. First, there is the **scheme fee** which is charged by the card schemes (i.e. Visa and MasterCard) for the use of their scheme. Second, there is the **interchange fee**, which is charged by the issuer for the provisioning of the credit service to the consumer. Third, there is the **acquiring fee**, which is

charged by the acquirer for receiving the funds of the transaction during the settlement process. Fourth, there is the **PSP fee**, which is charged by the PSP to cover the costs for establishing the connection between the merchant and the rest of the card network, over the acquirer. And at last, there are the **costs resulting from chargebacks** as issued by the consumer.



Figure 8.1: Merchant's cost for a transaction in the card network

- Scheme fee Both MasterCard and Visa operate various sub-schemes within their main schemes. Based on information provided by Adyen, we estimate that the costs for using the schemes vary between 0.3 and 0.7% [72].
- **Interchange fee** In the beginning of 2015 the European Commission adopted a proposal in which the interchange fees will be capped on 0.3% for all issuers within the European Union as of December 2015 [73], referred to as Interchange++. In our estimation we presume that issuers will implement this cap as their fee.
- Acquiring fee The acquiring cost are mostly affected by the fact if the settlement takes place real time or batch wise. Based on information provided by Adyen, we estimate that the costs of acquiring services vary between 0.4 and 0.6% [72].
- **PSP fee** This fee varies per PSP and can be differentiated for commercial purposes. In our estimation we will use Adyen's standard fee of € 0.10 [72].
- **Chargeback cost** For the chargeback cost, we will use the relative revenue loss due to chargebacks, minus the recovered revenue loss resulting from succesfull defended chargebacks (see Section 4.2). The effective chargeback cost is estimated to be 0.7% [17].

Debate about Interchange Fees

There has been an ongoing debate within the European Union about capping interchange fees [4]. As mentioned, the European Commission adopted a proposal in which the interchange fees will be capped on 0.3% as of December 2015, referred to as Interchange++ and only applies to two-leg transactions. Interchange++ is a policy measure, besides PSDI and PSDII, aimed at creating a more competitive market for payment methods within the EU.

8.2. Costs in the TPP Network

The transaction cost for a merchant in the TPP network depends on less factors than in the card network. As Sofort is currently the biggest player in the TPP network, we will consider costs imposed by Sofort. Just like for the estimation of the cost in the card network, we will only consider costs of two-leg transactions that take place within the European Union jurisdiction.

8.2.1. Differentiating and Estimating Cost Components

There are four main cost components of a transaction in the TPP network, as visualized in Figure 8.2. First, there is the **TPP fee** which is charged by the TPP for using their connections to the issuers. Then there are the **acquiring fee** and **PSP fee** just like in the card network. The chargeback cost is replaced by the non-payment cost. It must be mentioned that within the TPP network, there are no scheme or interchange fees as they are present in the card network.



Figure 8.2: Merchant's costs for a transaction in the TPP network

- **TPP fee** The TPP fee can be determined by the TPP and can be differentiated for commercial purposes. In our estimation we will use Sofort's TPP fee which varies between 1% and 5% [72].
- Acquiring fee Just like for the card network, the acquiring cost are mostly affected by the fact if the settlement takes place real time or batch wise. We use the same estimation as for the card network varying between 0.4 and 0.6% [72].
- **PSP fee** Just like for the card network, the PSP fee can be differentiated for commercial purposes. In our estimation we will use Adyen's standard fee of € 0.10 [72].
- **Non-payments cost** In Chapter 3 we have analyzed that the average non-payment percentage equals 0.3%. When considering the non-payment ratio using the transaction values rather than the transaction counts, the percentage also equals 0.3%. We take the expected value of the ROSI of the dunning process strategy, which has a confidence interval of 50%. This ROSI equals 4. Therefore, we are able to reduce the non-payment percentage to 0.075% (we will round this to 0.08% to account for significant figures).

8.3. Comparison

Table 8.1 presents an overview of the cost estimations for the card network and for the TPP network with and without implementation of the dunning process. When looking at the total cost ranges, it can be observed there is overlap in cost range between the two networks. Assuming that we would implement the dunning process, the lower range of the total costs for the merchant are less in the TPP than in the card network.

	Card network	TPP network	TPP network
		Without dunning	With dunning
Scheme fee	0.3 - 0.7%	n/a	n/a
Interchange fee	0.3%	n/a	n/a
Acquiring fee	0.4 - 0.6%	0.4 - 0.6%	0.4 - 0.6%
PSP fee	€ 0.10	€ 0.10	€ 0.10
TPP fee	n/a	1 - 5%	1 - 5%
Chargeback cost	0.7%	n/a	n/a
Non-payment cost	n/a	0.3%	0.08%
Total cost	1.7 - 2.3% + € 0.10	1.7 - 5.9% + € 0.10	1.48 - 5.68% + € 0.10

Table 8.1: Estimations of the merchant's costs for a transaction in the card and TPP network

The relatively big uncertainty in the estimation of the merchant's transaction costs in the TPP network over the card network, is primarily driven by the uncertainty in the TPP fee. As mentioned in Chapter 1 more TPPs are expected to enter the European market shortly after the formal adoption of PSDII [9]. We believe that this will result in lower TPP fees as more competition will emerge. **Combining the effect of this increased competition and the efficient management of the non-payment risk, we believe that the TPP network is a cost competitor of the card network within the European Union.**

A Note on the Interpretation of the Estimations

The analysis as presented in this chapter is based on estimations and only gives an indication of the merchant's transaction costs. Although we believe that the estimations provide a valid representation of reality, conclusions derived from this chapter should be interpreted while keeping in mind the numerous factors that can influence the real transaction cost.

Chapter Summary

In this chapter we have aimed to answer SQ5. Assuming that the most efficient strategy is deployed by the PSP, does this make the TPP network more financial efficient than the card network? We have differentiated and estimated the costs of a transaction for a merchant in the card and TPP network. Based on the estimations, we have concluded that the TPP network is a cost competitor of the card network within the European Union from the merchant's perspective. With this conclusion we have achieved our secondary research goal G2. To evaluate whether the introduction of the TPP network helps to reach the second goal of PSDII: the creation of more competition within the European market for payment methods.

\bigcirc

Conclusions

Now we have compared the merchant's transaction costs in the card and TPP networks, we move into the **prescription phase** of our research. In this chapter we conclude based on our findings. The first section provides answers to the subquestions, before answering the main research question in the subsequent section. The research questions are answered within the boundaries of our research. For a broader perspective on these boundaries, we provide a discussion in the next chapter.

9.1. Answering the Subquestions

SQ1. What kind of preventive and reactive strategies can be used to reduce the probability or impact of occurrence of non-payments in TPP networks?

To answer this question, we identified and evaluated potential strategies. For the preventive strategies we compared five different data mining approaches as proposed by Han et al. [48]. Our evaluation showed that methods based on *classification and regression for predictive analytics* and *outlier analysis* scored best on the four pattern evaluation criteria we used: *actionable*, *understandable*, *useful* and *valid*. In the first category, we chose a random forest implementation as our strategy, as introduced by Breiman [57]. In the second category, we chose break point analysis as our strategy, as introduced by Bolton and Hand [56]. For the reactive strategies we concluded that the typical approach that is used for card networks, dispute processing, is not feasible. We therefore decided to design a dunning process as the reactive strategy.

SQ2. How can the preventive and reactive strategies be applied in a minimal setup?

To answer this question, we have applied the three different strategies. The preventive strategies have been applied on the Sofort dataset. To deal with the challenge of imbalanced data for the random forest we used three tactics. First, we reduced the dataset to the intersection of the top five issuers and top ten merchants who caused most chargebacks. Second, we balanced our training dataset by applying random undersampling of the minority class, as presented by Bhattacharyya et al. [21]. And at last, we used the bagging method which is embedded in the random forest method as introduced by Breiman [57]. To deal with the challenge of different break point metrics, we decided to use Welch's t-test as this test is capable of dealing with (1) different sample sizes and (2) both equal and unequal variances. For the reactive strategy an experiment was conducted with the cooperation of one merchant.

SQ3. What kind of evaluation methods can be used to evaluate the financial efficiency of the preventive and reactive strategies?

To answer this question, we explored how to incorporate a variable risk exposure in the ROSI model. Therefore, we first explored different types of risk. We have identified that the risk of non-payments can be categorized as a type of operational risk. To quantify the exposure of such a risk for financial institutions, the Basel II Accord [69] provides the standardized approach, the basic indicator approach and the advanced measurement approach. We have selected the advanced measurement approach as this approach allows us to incorporate the most variable risk exposure. For the operationalization of the advanced measurement approach, we have selected the loss distribution approach as presented by Frachot et al. [66]. In this approach we are enabled to quantify the risk exposure based on historical frequency and severity distributions of the non-payment occurrence. The risk exposure is quantified as the aggregated loss distribution function, which presents the Value at Risk (VaR) for different confidence levels.

SQ4. How can the evaluation methods be applied to evaluate the financial efficiency of the preventive and reactive strategies?

To answer this question, we have applied the loss distribution approach. To do this, we first segmented our dataset subsets spanning periods of one week. Using these subsets, we determined the loss frequency and severity distributions. Using a Monte Carlo simulation to multiply the two distributions, we obtained the aggregated loss distribution. This distribution was used as the risk exposure in the ROSI calculations to compare the different risk management strategies. We have concluded that within the boundaries of our research, the reactive strategy comprising of the dunning process has a ROSI of four at a VaR confidence interval of 50% and has a positive business case as a result, while the preventive strategies random forest and break point analysis result in a loss for all confidence intervals and are as such not worth the investment.

SQ5. Assuming that the most efficient strategy is deployed by the PSP, does this make the TPP network more financial efficient than the card network?

To answer this question we have differentiated and estimated the costs of a transaction for a merchant in the card and TPP network. Based on the estimations, we have derived that the TPP network currently is a cost competitor of the card network within the European Union from the merchant's perspective. As such, we believe that by implementing our dunning process, the most cost effective strategy we have designed, the second goal of PSDII can be met: creating more competition within the European market for payment methods.

9.2. Answering the Main Question

Now we have obtained the answers to the subquestions, we can answer our main question RQ. How can a PSP manage the financial risk of non-payments in TPP networks efficiently, such that it becomes more efficient than the card network?

As identified, the PSP can use both preventive and reactive strategies to manage the risk of nonpayments. Based on our research, we conclude that the reactive strategy comprising the dunning process has a higher ROSI value than the preventive strategies. As such, the dunning process should be preferred over the preventive strategies. At a confidence of the VaR of 50% the ROSI of the dunning process is four, which means that for every euro invested in the strategy, four euro revenue is generated. By comparing the merchant's transaction costs in the card network and the TPP network — with the application of the dunning process — we have observed that transaction costs in the TPP network can be lower than in the card network.

We have to keep in mind that we only applied and evaluated three strategies, there are perhaps other strategies that would perform better. Also, we have to note that the implementation of our strategies might not be optimal, which can also influence the strategy performance. A further elaboration on this is presented in the next chapter.



10

Discussion

This chapter provides a discussion to position our research in a broader context. First, we reflect upon decisions that were made in our research process and present their limitations. This section also elaborates on the generalizability of the research. Subsequently, we highlight the main contributions of our research, within the boundaries of the reflection and limitations. In the last section of this chapter, we provide recommendations for future research. Each section will be structured according to the first four phases of our research approach as presented in Chapter 1: the *conceptual, converging, design* and *evaluation* phases.

10.1. Reflection and Limitations

Every decision has its implications. In this section we reflect upon decision that were made and present their limitations to our research, highlighted as (L1. - L8.). The limitations will be used as input to formulate opportunities for future research, later in this chapter.

10.1.1. Conceptual Phase

In Chapter 2 we have described the payment process for the TPP network. Because PSDII is not implemented in national legislation of the European Union member states yet, the process as it will be prevalent in the future can deviate from our process description. The process as we described can be categorized as an **overlay model**, as the authentication credentials are submitted to the issuer via the TPP. Alternative models include the **redirect model** and **aggregation model** [9]. In the redirect model, the consumer does not provide his credentials to the TPP. Instead, the TPP redirects the consumer to the issuer's environment to initiate the transaction. In the aggregation model, the consumer is only redirected by the TPP to the issuer's environment the first time he initiates a transaction. For subsequent transactions, the TPP can make use of a tokenized aggregation service, in order to obtain access to the consumer's account on the issuer's online banking environment. Such services do not exist in the TPP network yet, however in the card network they do (e.g. PayPal).

We have considered the overlay model because in our opinion it is currently the most feasible model in the European context. As described in Chapter 2 the German banking landscape is highly diverse. The European landscape however, is even more diverse. We presume that the chance that banks will cooperate in the (near) future to facilitate the redirect and aggregation models is limited. However, there is currently a lot of debate about how PSDII will be implemented by the market and as

such which models will be adopted, as can be derived from the interview transcriptions in Appendix F. The two biggest TPPs within the European Union — Sofort and Trustly — use the overlay model. As we have used Sofort transaction data, this poses a limitation to our research (L1.).

Besides, we have assumed that non-payments in the TPP network can be ascribed to either (1) **fraud** committed by the consumer and (2) **technical malfunctioning** caused by the issuer. Of course, both types of issues can also be ascribed to other parties in the industry. We deem it plausible however that most fraud originates from the consumer as companies would highly jeopardize their business by committing fraud. Also, we assume that, given the fact that innovation in the payment industry is primarily driven by the non-banks, technical malfunction primarily resides with the banks. Given the high dependency of the overlay model on the issuer, we deem it plausible that most technical issues reside at the issuer. An example of this can be found in the correlation between the issuer's settlement delay and the non-payment ratio as presented in Chapter 3. However, because we have focused on (1) fraud committed by the consumer and (2) technical malfunctioning caused by the issuer, we might have missed other causes for the occurrence of non-payments, which is a limitation to our research (L2.).

10.1.2. Converging Phase

In Chapter 3 we have introduced Sofort and presented it as the case study for our research. The question arises, can this case study produce generalizable results? To answer this question, we distinguish the (1) generalizability of the Sofort case and (2) generalizability of the results of application of the risk management strategies. For the generalizability of the Sofort case, we refer back to Section 3.2 in which we presented that Sofort is (1) the biggest TPP within the European Union and (2) Adyen's Sofort portfolio comprises of the total Sofort transaction volume. As such, we believe that our case study has the potential to produce results which can be generalizable to other TPPs that use the overlay model, like Sofort. However, as we have merely incorporated one TPP in our case study, this is a limitation to our research (L3.) The generalizability of the results of application of the risk management strategies will be addressed in the next section.

The Sofort dataset that we have used only contains transactions from the DACH-region, of which 90.9% originates from Germany. As a result, the observations of this research are primarily applicable to Germany. However, we presume that given the heterogeneous nature of the German banking land-scape, as presented in Section 3.1, the results of our research have to potential to be generalizable to the heterogeneous European banking landscape. However, the geographical focus of the Sofort dataset on Germany remains a limitation to our research (L4.).

10.1.3. Design Phase

In Chapters 4 and 5 we have proposed and applied strategies to manage the risk of non-payments. For the preventive strategy applications we have used subsets of the Sofort dataset. You could argue that using these subsets decreases the **generalizability** of our findings. However, the subsets were created based on logic that can be reproduced. For the random forest, we reduced the dataset to the intersection of the top five issuers and top ten merchants who caused most non-payments. Analyzing the data revealed that these top five and top ten are constant over time. For the break point analysis, it would not make any sense to incorporate transactions of consumers that initiated less transactions than the size of our moving window (24 transactions).

For the reactive strategy, the dunning process, we refer back to Section 5.2. As explained, the merchant that cooperated in the dunning process sells digital gift cards which are prime subjects of fraudsters, as they are easy to cash out. We believe that if dunning works for this type of *high risk*

products, it will also work for products or services which are less easy to cash out. However, we do acknowledge that using subsets of the data for the strategy application imposes limitations to our research (L5.).

Operational Limitations during the Design Phase

There are limitations to our research that result from decisions that were made in the strategy application, which we will refer to as operational limitations (L6.). With the design of the random forest, we have ignored the challenge of genuine random subsetting the data. As mentioned in Section 5.1, the creation of genuine random subsets has been impossible until now. However, there are techniques that can be used to achieve better randomness than we have used. With the design of the break point analysis, we have not optimized the application of the Welch's t-test. We simply used the moving window sizes as argued by Bolton and Hand [56]. Another limitation to our break point analysis, is the aggregation of the data based on IBAN. Possibly, (fraudulent) consumers use various IBANs for their activities. As such, this behavior will not be identified by our break point analysis implementation.

10.1.4. Evaluation Phase

In Chapters 6 and 7 we have proposed and applied evaluation methods for the risk management strategies. While determining the ROSI curves, we have only incorporated direct costs, as mentioned in Section 7.2. Why have we not accounted for indirect costs such as caused by software development and infrastructure development? Because we assume that due to the economies of scale within the fintech sector, which enables businesses to spread investment costs over numerous transactions, these costs can be neglected. However, it does remain a limitation to our research (L6.).

Operational Limitations during the Design Phase

Just like for the strategy application, there are limitations to our research that result from decisions that were made in the evaluation application, which we will refer to as operational limitations (L8.). While applying the loss distribution approach, we have fitted the distributions manually by plotting them on top of the frequency histograms. There are various methods that can be used to execute this fitting process in a statistical more valid way. Besides, we have chosen to partition the data in periods of one week, which can influence the aggregated loss distribution. We presume however that the applied fitting process and partitioning have not influenced our results significantly, as the same *risk exposure* distribution has been applied to the evaluation of all three strategies.

10.2. Contributions

We see four main contributions in our work, highlighted as (C1. — C4.). The contributions are mainly driven by the uniqueness of the case study, as elaborated on in Chapter 1. The first contribution encompasses describing the payment process in the TPP network. The second contribution entails the analysis of the Sofort transaction data. The third contribution entails the design and application of risk management strategies for the TPP network. Finally, we see a contribution in the cost comparison of the merchant's transaction costs in the card and TPP networks.

10.2.1. Conceptual Phase

C1. Describing the Payment Process in the TPP Network

To the best of our knowledge, we are the first ones to have described the payment process in the TPP network and to have differentiated its authorization and settlement processes. The full process description is presented in Section 2.1. We have identified that the main difference between the payment process in the card and TPP network resides in the fact that for the authorization process in the card network the schemes are required, while for the authorization process in the TPP network they are not.

10.2.2. Converging Phase

C2. Analyzing Transaction Data from the TPP Network

We are the first ones to have analyzed patterns in TPP transaction data to differentiate payments from non-payments, as presented in Section 3.3 — to the best of our knowledge. We identified attributes that seem to influence the non-payment ratio in four different categories. In the category *transaction* attributes, these include the day of the week and the hour of the day. For the *merchants*, we observed a high concentration of non-payments among the top ten merchants that account for most transactions. When it concerns the *consumer*, the country and device seem to influence the non-payment ratio. For the *issuers*, we also observed a high concentration of non-payments among a selection of issuers that accounts for most transactions. Besides, there seems to be a weak positive correlation between the issuer's non-payment ratio and his average settlement delay.

10.2.3. Design Phase

C3. Designing and Applying Risk Management Strategies for the TPP Network

Although we are not the first ones to have applied the strategies we did, we are the first ones to have applied them to manage the risk of non-payments in the TPP network — to the best of our knowledge. During the process, we noticed that the preventive strategies we used — which are frequently used in the card network — do not deliver as good as a performance as in the card network. This knowledge can be used in the future when designing risk management strategies for the TPP network.

10.2.4. Evaluation Phase

C4. Comparing the Merchant's Transaction Cost in the Card and TPP Networks

To the best of our knowledge, we are the first ones to have compared the merchant's transaction cost in the card and TPP networks. There is a relatively big uncertainty in the estimation of the merchant's transaction costs in the TPP network over the card network, which is primarily driven by uncertainty in the TPP fee. As mentioned in Chapter 1 more TPPs are expected to enter the European market shortly after the formal adoption of PSDII. We believe that this will result in lower TPP fees as more competition will emerge. Combining the effect of this increased competition and the efficient management of the non-payment risk, for example by using our dunning process, we believe that the TPP network is a cost competitor of the card network within the European Union.

Contribution to the Work of Anderson

As mentioned in Chapter 2, we were only able to identify related academic research into the TPP network in the work of Anderson [4]. Although Anderson's work is mostly qualitative, he was able to identify (1) Sofort as an innovative payment method competing with the card network on the merchant's transaction costs and (2) a security threat in Sofort's payment process. In our work we have provided a quantitative perspective on both aspects. We have been able to confirm that the threat is present, presented quantitative insights into the threat and provided strategies for a PSP how to manage the risk resulting from the threat. We have compared the merchant's transaction costs in the TPP and card network, and can agree with Anderson that the TPP network indeed is a cost competitor of the card network when considering the merchant's transaction costs.

10.3. Recommendations for Future Research

In the first section of this chapter we presented a reflection upon decision that were made, and their resulting limitations. In this section we provide recommendations for further research for each limitation.

10.3.1. Conceptual Phase

- L1. Investigating different TPP models As elaborated on, the overlay model is currently most used within the European Union. The question remains how PSDII will be implemented within national legislation of European Union member states, and if this will influence the existing overlay model. Therefore, we recommend further research into alternative models, like the aggregation and redirect models. We believe that important aspects to consider are the models' privacy and security implications for the consumer.
- L2. Finding different causes of non-payments In our research we have assumed that the non-payments are either caused by fraud committed by consumers, or technical malfunctioning caused by the issuers. We recommend research into the identification of different causes or *modus operandi*, such as fraud committed by merchants or TPPs, technical malfunctioning caused by the acquirers etc. Also, payment that do not result in a non-payment might be fraud related, for example when it concerns money laundering activities.

10.3.2. Converging Phase

- L3. Investigating different TPPs Since we have solely considered transaction data from Sofort, we recommend additional research considering transaction data from other TPPs. This can be beneficial to create more, and more representative, insights into the TPP network.
- L4. Analyzing TPP transactions outside DACH-region The Sofort dataset we used only contained transactions originated from the DACH-region. We recommend additional research considering transaction data from outside this region. Just like for the previous recommendation, we believe this can be beneficial to create more, and more representative, insights into the TPP network. Especially because there might be differences in authentication procedures at the issuer's side between countries, which can create country specific (fraud) risks.

10.3.3. Design Phase

- L5. Investigating different subsets In our research we have used specific subsets of the data for the strategy application, as elaborated on in Chapter 5. We recommend research into the prediction of non-payments in other subsets, especially subsets that are characterized by a higher imbalancedness in class type prevalence. Such research can be beneficial for the efficiency of the strategy application, not only in our research, but in data mining in general.
- L6. Optimizing risk management strategies Given the exploratory nature of our research, we have not focused on optimizing the strategy application. We believe that there is potential for optimization for all strategies. Because the dunning process has showed a positive return on investment, we would especially recommend further research into optimizing the dunning process.

10.3.4. Evaluation Phase

- L7. Incorporating indirect costs As elaborated on in the beginning of this chapter, we have only included direct costs in our strategy evaluation. We would be interested to see how the strategy evaluation would work out if we would include indirect costs too. An example question to ask would be, are the indirect costs of reactive strategies different than the indirect costs of preventive strategies?
- L8. Optimizing evaluation application In the evaluation application we have quantified the risk exposure by using the loss distribution approach. We would recommend further research into the effect of characteristics of the incident data on the ROSI calculation. An example question to ask would be, could situations emerge in which strategy A would be preferred over strategy B at a VaR confidence interval *x*, while this preference would not exist at a VaR confidence interval of *y*? And what is the effect of the implementation of a risk management strategy on the aggregated loss distribution? Does the distribution need to be updated in order to remain a (valid) representation of the risk exposure? And if so, what are the implications of an updated distribution on the evaluation of other strategies?



Article

Risk Management for Third Party Payment Networks

J.W. (Willem) van Driel Faculty of Technology, Policy and Management Delft University of Technology Jaffalaan 5, 2628 BX, Delft Email: wvdriel@gmail.com

Abstract—The payment industry is dominated by a small number of players, like MasterCard and Visa, who operate schemes in the card network. With the introduction of Payment Service Directive II within the European Union, the third party payment (TPP) network is created as a competitor of the card network. Whereas one of the main challenges in the card network is the cost effective management of chargebacks, a similar challenge is present in the TPP network. Merchants have noticed that an authorized payment in the TPP network does not always result in a settlement. These non-payments can impact the merchant's revenue negatively, as they typically provide their products or services to the consumer on authorization. In this paper we report on our research into the TPP network and we present the four contributions we have made. First, we describe the payment process as it is organized in the TPP network and contrast it with the card network. Second, we present patterns discovered in transaction data from the TPP network, that let us distinguish payments from nonpayments. Third, we report on risk management strategies that can be used to manage the non-payment risk. And fourth, we compare the merchant's transaction costs in the card and TPP networks to verify if the TPP network can become a cost competitor of the card network.

Index Terms—Third Party Payment Networks, Risk Management, Non-payments

1. Introduction

Like most consumers, you probably take the payment process when checking out at your favorite webshop, for granted. Although the payment process may look simple from the consumer's perspective, a lot of complexity is hidden behind the scenes. Over the past decades an entire financial sector has emerged to facilitate the payment process: the payment industry. The industry is dominated by a small number of players, like MasterCard and Visa, that operate the schemes for the facilitation of credit and debit card payments [1]. Regardless of which payment method you select; PayPal, Google Wallet, Apple Pay or you directly provide your card details, chances are big that the schemes of MasterCard or Visa are used — without you having noticed. A scheme can be regarded as the standardized way of communication between the consumer's bank, the issuer, and the merchant's bank, the acquirer. We refer to the part of the payment industry that facilitates debit and credit card payments as the card network. In order to create more competition for the card network, the European Commission is working on Payment Service Directive II (PSDII) [2] which introduces the third party payment service provider (TPP). With this introduction the TPP network is created, next to the existing card network, providing the payment industry the opportunity to develop payment methods that circumvent the dominant position of MasterCard and Visa.

2. Research Problem

Although there seems to be quite some potential for the TPP network, challenges are present. The new TPP network is less mature than the established card network that has dominated the payment industry for decades. Whereas one of the main challenges in the card network is the cost effective management of chargebacks, a similar challenge is present in the TPP network. One of the first PSPs to facilitate payments in the TPP network is payment service provider (PSP) Adyen. Merchant's making use of Adyen's services have noticed that an authorized payment in the TPP network does not always result in a settlement. This can impact the merchant's revenue negatively, as merchants typically provide their products or services to the consumer on authorization. In this paper we will refer to the authorized transactions that were not settled in the TPP network as nonpayments.

Payment methods broadly fall into two categories: guaranteed and non-guaranteed. Non-guaranteed payment methods offer the consumer the option to reverse a payment after it has taken place, while guaranteed payment method do not. Payment methods in the card network are typically nonguaranteed [3], while payment methods in the TPP network should be guaranteed [4]. PSPs like Adyen have a responsibility towards their customers, the merchants, to provide guaranteed payment methods that are truly guaranteed i.e. when the payment was authorized, the settlement should take place and no reversal should occur. Since practice has proven that this is not the case, we need to understand how big the financial risk of the occurrence of non-payments is and what could be strategies a PSP could use to manage this risk.

For our research, Adyen has provided a unique dataset containing a few million transactions from the TPP network, spanning the period between 1/8/2014 and 1/8/2015. Unfortunately, we cannot disclose the exact number of transactions because of commercial reasons - however, we presume this is outweighed by the fact that we find ourselves in the unique position to be the first academia to report on research using transaction data from the TPP network. In this paper, we report on the four contributions we have made with our research. First, we describe the payment process as it is organized in the TPP network. Second, we present patterns discovered in the data, that let us distinguish payments from non-payments. Third, we report on possible risk management strategies that can be used by a PSP to manage the nonpayment risk. And fourth, we compare the merchant's transaction costs in the card and TPP networks to verify if the European Commission has achieved one of PSDII's goals with the introduction of the TPP network: the creation of a competitive payment network within the European Union, next to the card network [4].

3. Third Party Payment Networks

The payment process in the TPP network can be separated in two main processes, just as in the card network: (1) authorization and (2) settlement. The authorization processes is a flow of information, while in the settlement processes the actual money is transferred. Both processes are visualized in Figure 1. The authorization process encompasses ten steps:

- The consumer indicates to the merchant he wants to pay using the TPP of his choice
- The merchant connects with the PSP to initiate the payment
- 3) The PSP sends the payment initiation to the TPP
- The TPP asks the consumer for the credentials of his online banking environment
- 5) The consumer submits these credentials to the TPP
- The TPP connects with the issuer's online banking environment and initiates the payment
- If the payment initiation is successful, the issuer sends an approved authorization to the TPP. If the initiation is not successful, the issuer sends a rejected authorization.
- The TPP sends the authorization response to the PSP
- 9) The PSP notifies the merchant of the authorization response
- 10) The merchant notifies the consumer of the authorization response

Just like for the card network, the merchant typically delivers his products or services after step ten and only in case the authorization was approved. During the settlement process, the funds travel from the consumer to the merchant as visualized by step 11 - 15.



Figure 1: Payment process in the TPP network

When comparing the payment process in the TPP network with the payment process in the card network, there are two main differences. First, in the TPP network the schemes are not used for the authorization process. Instead, the TPP is responsible for obtaining the authorization from the issuer directly. Second, the TPP network makes use of the standard bank transfer scheme — the SEPA Credit Transfer — for the settlement process. Because of these two differences, payment processing in the TPP network has no dependency on the schemes of the card network.

4. Data

As we already touched upon, one of the challenges in the TPP network is the occurrence of non-payments. During our research we discovered that non-payments, just like chargebacks in the card network, can be either fraud or technical related. One of the perks of the maturity of the card network is the fact that the schemes typically provide reason codes to the PSP to indicate the cause of the chargeback. In the TPP network however, these reason codes are non-existent. Therefore, an important step in our research was to derive patterns from the data that can indicate how to differentiate payments from non-payments. An overview of the attributes that influence the non-payment ratio is presented in Table 1. The non-payment ratio is calculated as the count of the nonpayments within the considered attribute space, divided by the full transactions count within the same attribute space. The average non-payment ratio within our dataset equals 0.0029 or 0.29%.

Category	Attribute
Transaction	Day of week, hour of day
Merchant	Name
Consumer	Country, device
lssuer	Name, settlement delay

Table 1: Attributes that influence the non-payment ratio

In the subsequent sections we will present the relation between the attributes and the non-payment ratio.

4.1. Transaction

For the transaction related attributes we observed an increased non-payment ratio for transactions that were initiated in and around the weekend, as visualized in Figure 2. Also, we observed an increased ratio for transactions initiated during the night. At its peak at 3am the ratio is more than twice as high as during its low at noon. At its peak the nonpayment ratio is about 0.006, which implies that 0.6% of all transactions initiated at that time result in a non-payment.



Figure 2: Non-payment ratio per day of the week and hour of the day

4.2. Merchant

For the merchant related attributes, the data revealed that there is a concentration of non-payments on a small selection of the merchants. The top ten merchants that generated most non-payments account for 74% of all non-payments while only accounting for 58% of all transactions. Most of these merchants offer gaming or gambling related products or services.

4.3. Consumer

For the consumer related attributes we observed an increased non-payment ratio for transactions initiated from mobile devices. The non-payment ratio for consumers using their tablet or computer is comparable. The non-payment ratio also seems to be influenced by the country the transaction was originated from. The non-payment ratio for transactions originated from Switzerland is more than four times as high than transactions originated from Germany and Austria.





4.4. Issuer

For the issuer related attributes, the data revealed that there is a concentration of non-payments on a small selection of the issuer. The top five issuer that generated most non-payments account for 73% of all non-payments while only accounting for 28% of all transactions. This indicates an even higher concentration of non-payments on specific issuers than on specific merchants. Also, we found a relation between the issuer's average settlement delay — i.e. the time between the authorization and the settlement, as visualized in Figure 4. Non-business days are excluded from the settlement delay, as banks typically only settle on business days.



Figure 4: Settlement delay versus non-payment ratio

5. Risk Management

In order to manage the financial risk resulting from the nonpayments in the TPP network, we designed and applied three different risk management strategies. We differentiated preventive and reactive strategies. A preventive strategy is aimed at blocking a transaction based on the suspicion that it will result in a non-payment. A reactive strategy is deployed after the transaction has resulted in a non-payment and is solely focused on recovery of the due funds. For the design of the preventive strategies we have used data mining methods. For the reactive strategy, we have designed an experiment in cooperation with a merchant that uses the TPP network to accept payments. In the following sections we provide an introduction to the strategies. Specific implementation details are left out of this paper as the strategies have merely been used to *explore* their potential for application in the TPP network.

5.1. Preventive Strategy 1: Supervised Machine Learning

As our first preventive strategy we implemented a method from the data ming domain of supervised machine learning: random forest as introduced by Breiman [5]. In random forest, a classifier is trained using historical observations to identify patterns that increase the probability that a transaction results in a non-payment. For our implementation of the random forest we have used a subset of the dataset, comprising of the intersection of the top five issuers and top ten merchants that caused most chargebacks. As our attribute selection, we used the attributes as presented in Table 1 as they have proven to enable us to differentiate payments from nonpayments.

5.2. Preventive Strategy 2: Unsupervised Machine Learning

As our second preventive strategy we implemented a method from the data mining domain of unsupervised machine learning: break point analysis as introduced by Bolton and Hand [6]. In break point analysis, the classifier tries to identify moments in a time series where the observed behavior significantly deviates from historical behavior. When applying the technique to payments, it can be regarded as an anomaly detector aimed at detecting changes in the consumer's spending behavior. As argued by Bolton and Hand, the technique can therefore be used to detect payment fraud — especially when it concerns fraud resulting from account takeovers, as argued by Kahn and Roberds [7].

5.3. Reactive Strategy 2: Dunning Process

As our reactive strategy we implemented a dunning process. The dunning process starts at the moment the non-payment occurs and the products or services are already delivered. In order to fulfill the due amount, the merchant (1) contacts the consumer, (2) informs him about the unfulfilled transaction and (3) asks him to fulfill the amount due. For the fulfillment, SEPA Direct Debits were used, which can be regarded as retries of the SEPA Credit Transfers that are used for the settlement process in the TPP network.

5.4. Results

The results of the application of the three strategies are presented in Table 2. For each strategy we have included two measures: the direct costs of the application and the risk mitigation. The direct costs of the preventive strategies are calculated by summing the amounts of the blocked transactions which would not have resulted in a non-payment (false positives). The direct costs of the reactive strategy are calculated by summing the costs of processing the SEPA Direct Debits. The risk mitigation for each strategy is calculated as the total value of non-payments that were (1) prevented or for which (2) the due funds were recovered, divided by the total value of the non-payments.

	Random forest	Break point analysis	Dunning process
Direct costs	€ 5,578,864	€ 3,766,977	€ 165
Risk mitiga- tion	0.25	0.068	0.76

Table 2: Direct costs and risk mitigation per strategy

It can be observed that the direct costs of the dunning process are lowest, and the risk mitigation of the dunning process is highest. When we would regard the strategies as investments, the question arises, what is the return on these investments? To answer this question, we made use of the Return On Security Investment Model, as introduced by Sonnenreich, Albanese and Stout [8].



Figure 5: ROSI curves for the risk management strategies

When comparing the ROSI curves, as visualized in Figure 5, we can observe that a confidence interval of 50% of the Value at Risk (VaR) — which represents the risk exposure in the ROSI model — the ROSI value for the dunning process is four, while the ROSI value for the preventive strategies is around minus one. This implies that for each euro invested in the dunning process, four euro in revenue will be obtained, while investments in the preventive strategies would result in a loss.

6. TPP Networks versus Card Networks

The last part of our research focused on the comparison of the merchant's transaction costs in the card and TPP networks. Table 3 presents the cost components of a transaction in both networks and provides estimations of their values. The estimations are based on public figures provided by the European Commission [9], external available data of Adyen [10], and internal available data of Adyen [11]. For the estimation of the non-payment cost in the TPP network, we derived from the dataset that the non-payment ratio equals 0.3%. However, as we assumed we would implement the dunning process which has a ROSI of four at an confidence interval of 50%, we are able to reduce the non-payment percentage to 0.075% (we will round this to 0.08% to account for significant figures).

	Card network	TPP network
		With dunning
Scheme fee	0.3 - 0.7%	n/a
Interchange fee	0.3%	n/a
Acquiring fee	0.4 - 0.6%	0.4 - 0.6%
PSP fee	€ 0.10	€ 0.10
TPP fee	n/a	1 - 5%
Chargeback cost	0.7%	n/a
Non-payment cost	n/a	0.08%
Total cost	1.7 - 2.3%	1.48 - 5.68%
	+ € 0.10	+ € 0.10

Table 3: Estimations of the merchant's transaction costs in the card and TPP networks

When looking at the total cost ranges for both networks, it can be observed there is overlap in cost range between the two networks. Assuming that we would implement the dunning process, the lower range of the total costs for the merchant are less in the TPP network than in the card network.

7. Conclusions

The comparison of the merchant's transaction cost implies that merchants can process payments in the TPP network for a lower transaction cost than in the card network. There is a relatively big uncertainty in the estimation of the merchant's transaction costs in the TPP network over the card network however, as a result of the uncertainty in the TPP fee. We believe that the market entrance of new TPPs will enhance competition and will result in lower TPP fees. Combining the effect of this increased competition and the efficient management of the non-payment risk, for example by using our dunning process, we believe that the TPP network is a cost competitor of the card network within the European Union.

8. Discussion and Future Research

For the facilitation of transactions in the TPP network, TPPs within the European Union currently use the overlay model. In the overlay model, the consumer provides his bank's authentication credentials to the TPP, who submits them at the issuer. Although we believe that the overlay model has quite some potential, country specific implementations of the PSDII directive might push alternative models to the market. We see an opportunity for future research into the comparison with these alternative models, especially from the consumer's privacy and security perspective.

References

- R. Anderson, "Risk and Privacy Implications of Consumer Payment Innovation," *Payment Systems Conference*, p. 20, 2012.
- [2] European Commission, "Proposal for a Directive of the European Parliament and of the Council on payment services in the internal market," 2013.
- [3] R. DeGennaro, "Merchant Acquirers and Payment Card Processors: A Look inside the Black Box," *Economic Review - Federal Reserve Bank of Atlanta*, vol. 91, no. 1, pp. 27–II, 2006.
- [4] European Commission, "Impact assessment PSD2," vol. 1, pp. 1–82, 2013.
- [5] L. Breiman, "Random forests," *Machine learning*, pp. 5–32, 2001.
- [6] R. J. Bolton, D. J. Hand, and D. J. H, "Unsupervised Profiling Methods for Fraud Detection," *Proc. Credit Scoring and Credit Control VII*, pp. 5–7, 2001.
- [7] C. M. Kahn and W. Roberds, "Credit and identity theft." 2005.
- [8] W. Sonnenreich, J. Albanese, B. Stout, and L. L. C. W. St, "Return On Security Investment (ROSI) – A Practical Quantitative Model," vol. 38, no. 1, pp. 55–66, 2006.
- [9] European Commission, "Commission welcomes European Parliament vote to cap interchange fees and improve competition for card-based payments," 2015.
- [10] Adyen, "Adyen Pricing Overview," 2015.
- [11] Adyen, "Business Intelligence Tool," 2015.

 \mathbb{B}

Software Implementation

This appendix introduces the software that was used for our research. Its usage is explained based on two activities, the data collection and the data analysis.

B.1. Data Collection

The Sofort dataset contains **transactional data** and is stored in Adyen's **PostgreSQL** database. The foremost challenge of obtaining the data, was its volume. The obtained dataset contains a total of 568 megabytes. Because of the volume of Adyen's database as a whole, obtaining this 568 megabytes from the database would take much computational power. Therefore, we decided to split the query to obtain the data in batch periods of one week. For this, the SQL language [74] was used. The following steps have been performed to collect the data:

- 1. Establish connection with PostgreSQL database
- 2. Write query to gather required data
- 3. Use bash command to split query in batches
- 4. Query individual batches of data
- 5. Combine data batches into one .csv file

B.2. Data Analysis

After the data was obtained, it could be used for the data analyses. For the analyses conducted in this research, the statistical language R [75] has been used. The integrated development environment (IDE) R-Studio [76] was used for the implementation of the R code. In order to load the data in the R-Studio IDE, the process as started above was continued with the following steps:

- 6. Load .csv file into the R-Studio IDE
- 7. Clear data inconsistencies, set proper data formats, etc.

For the individual analyses, libraries were used to speed up the implementation process. Frequent use was made of the following libraries:

RPostgreSQL bridges SQL with R
Iubridate converts date and time attributes
data.table extends the data.frame object, enabling to work fast with large volumes of data
ggplot2 enables fast visualization of large volumes of data
party enables the creation of decision trees
randomForest enables the random forest implementation

Initial Data Exploration

This appendix introduces the initial data exploration that was conducted for our research. First, a description of the data that was used, is provided. Second, a decision tree is presented in order to get a first insight into attributes that distinguish payments from non-payments.

C.1. Data Description

The retrieved data comprises of all Sofort transactions that were handled by Adyen over a time interval of one full year — between 1/8/2014 and 1/8/2015. The specific procedure that was followed to collect the data can be found in Appendix B. The data contains five types of attributes. The first type concerns data about the **transaction** itself, like id, date time and amount. The second type involves the **merchant's** name. The third type concerns data about the **consumer**, like his name, email, IBAN, country and device type. The fourth type involves **issuer** specific data, like the BIC, name, country and settlement delay. The fifth type involves the attribute whether the transaction resulted in a **nonpayment** or not. This is a boolean, so the value of this attribute is true or false. There is no additional data available about the type of non-payment or its reason of occurrence. All available attributes are displayed in Table C.2.

Descriptive	Value	
Number of transactions		
Number of merchants		
Number of consumers		
Number of issuers		
Number of non-payments		

Table C.1: General descriptives of the data

Table C.1 contains some general descriptives of the obtained data. It can be observed that the number of transactions is more than three times as high as the number of consumers. This implies that there are recurring consumers within the dataset. Also, it can be observed that the number of chargebacks is a small percentage of the total number of transactions — 0.29%.

Category	Attributes	Format
	ID	Integer
Transaction	Date and Time	Timestamp
	Amount	Integer
Merchant	Name	String
	Name	String
	Email	String
Consumer	IBAN	String
Consumer	Country	String
	Device type	String
	BIC	String
lequor	Name	String
155001	Country	String
	Settlement delay	Timestamp
	Non-payment	Boolean

Table C.2: Attributes as present in the Sofort dataset

C.2. Decision Tree

To explore the data, we want to get insight into attributes that distinguish payments from non-payments. Decision trees can be used to obtain such insights. There are various algorithms that can be used to create decision trees, one of them is the conditional inference algorithm [77]. An advantage of the conditional inference algorithm over other algorithms is its capability to split on both continuous and discrete attributes. For continuous attributes linear regression is used, for discrete attributes the categories as defined within the attributes are used. The attributes presented in Table C.3 were identified in Chapter 3 to influence the non-payment ratio and are used as input variables for the decision tree.

Category	Attribute
Transaction	Day of week, hour of day
Merchant	Name
Consumer	Country, device
Issuer	Name (BIC), settlement delay

Table C.3: Attribute selection for decision tree

Figure C.1 shows the decision tree. A concentration of non-payments is present in the leaves in the left hand side of the tree. The branches that lead to these leaves seem to share a similar itemset containing: (1) days in and around the weekend, (2) hours during the morning or night, (3) a specific set of merchants and (4) a specific set of issuers. Although most of the leaves in the left hand side of the tree present small subsets of the data, the concentration of non-payments goes up to 5-10% while the overall non-payment percentage is 0.29%. This might be an indication of the presence of patterns in the Sofort dataset that will enable us to distinguish normal payments from non-payments. We ignore the high concentration at Node 44, as this node only concerns 9 transactions which indicates overfitting.


Figure C.1: Decision tree — Names of issuers and merchants are replaced by numbers for anonymization, days of the week are represented numerically ranging from 1 (Sunday) till 7 (Saturday)

Code Implementations

This appendix contains a selection of the R code that was implemented for our research. The implementations of the decision tree, random forest, break point and loss distribution approach analysis are presented.

D.1. Decision Tree

The code below shows the R implementation to generate the decision tree. For the generation of the tree, the party library was used.

```
# Loading the decision tree library
1
  library(party)
2
3
  # Loading and cleaning the data
4
  [...]
5
6
  # Creating subset of top 5 issuers and top 10 merchants
7
8 list bic small<-sofort[chargeback==TRUE, .N, by=bic small][order(-N)][1:5][,</pre>
     factor(bic small)]
9 list merchantaccountid<-sofort[chargeback==TRUE,.N,by=merchantaccountid][</pre>
     order(-N)][1:10][,factor(merchantaccountid)]
10
11 data<-copy(sofort[bic small%in%list bic small&merchantaccountid%in%list
     merchantaccountid])
12
  # Selecting the attributes for classification
13
14 data<-data[,.(day of week=factor(day of week),</pre>
                 hour of day=as.numeric(hour of day),
15
                merchant=factor(merchant),
16
                 country=factor(shoppercountrycode),
17
                 device=factor(device),
18
                 issuer=factor(issuer),
19
                 issuer settlement delay=as.numeric(settlement delay),
20
```

```
21 chargeback=factor(chargeback))]
22
23 # Generate decision tree
24 decison_tree<-ctree(chargeback~day_of_week+hour_of_day+merchant+country+
        device+issuer+issuer_settlement_delay, data=data, controls=ctree_control(
        maxdepth=5))
25
26 # Plot the decision tree
27 plot(decison_tree)</pre>
```

Listing D.1: Decision tree implementation in R

D.2. Random Forest

The code below shows the R implementation to generate the random forest. For the generation of the random forest, the randomForest library was used.

```
1 # Loading the random forest library
2 library(randomForest)
3
  # Loading and cleaning the data
5 [...]
 # Creating subset of top 5 issuers and top 10 merchants
7
 list bic small<-sofort[chargeback==TRUE,.N,by=bic small][order(-N)][1:5][,</pre>
8
     factor(bic small)]
9 list merchantaccountid<-sofort[chargeback==TRUE,.N,by=merchantaccountid][</pre>
     order(-N)][1:10][, factor(merchantaccountid)]
10
11 data<-copy(sofort[bic small%in%list bic small&merchantaccountid%in%list</pre>
     merchantaccountid])
12
13 # Selecting the attributes for classification
 data<-data[,.(day of week=factor(day of week),</pre>
14
                hour of day=as.numeric(hour of day),
15
                merchant=factor(merchant),
16
                 country=factor(shoppercountrycode),
17
                 device=factor(device),
18
                 issuer=factor(issuer),
19
                 issuer settlement delay=as.numeric(settlement delay),
20
                 is first transaction=factor(is first transaction),
21
                 time since last transaction=as.numeric(time since last
22
                    transaction),
                 cumulative amount_this_day=as.numeric(cumulative_amount_this
23
                    day),
                 cumulative count this day=as.numeric(cumulative count this
24
                    day),
                chargeback=factor(chargeback))]
25
```

```
26
  # Creating training and testing set, for performance purposes we create a
27
     balanced training set containing 50% non-payments
28 data[,id:=1:.N]
29 data chargeback<-data[chargeback==TRUE]
  train<-rbind(data chargeback[1:round(data chargeback[,.N]*0.5)],data[</pre>
30
     sample(.N,round(data chargeback[,.N]*0.5))])
31 test<-data[!(id%in%train[,id])]</pre>
32 model <- random Forest (chargeback~day of week+hour of day+merchant+country+
     device+issuer+issuer settlement delay+is first transaction+time since
     last transaction+cumulative amount this day+cumulative count this day,
     data=train)
33
  # Evaluation of the classifier's performance
34
35 predicted<-predict (model, newdata=test)</pre>
36 predicted evaluate<-as.data.table(predicted)
37 predicted evaluate[,id:=1:.N]
38 test evaluate<-test[,id:=1:.N]</pre>
39 evaluate<-merge(predicted evaluate,test evaluate[,.(id,truth=chargeback)],</pre>
     by='id')
40
  # Calculate the values of the confusion matrix
41
42 tp<-evaluate[truth==TRUE&predicted==TRUE,.N]
43 tn<-evaluate[truth==FALSE&predicted==FALSE,.N]
44 fp<-evaluate[truth==FALSE&predicted==TRUE,.N]
45 fn<-evaluate[truth==TRUE&predicted==FALSE,.N]
46
  # Calculate the precision and recall of the classifier
47
48 precision<-(tp/(tp+fp))</pre>
49 recall<-(tp/(tp+fn))
```

Listing D.2: Random forest implementation in R

D.3. Break Point Analysis

The code below shows the R implementation to perform the break point analysis.

```
1 # Loading and cleaning the data
2 [...]
3
4 # Obtaining a list of unique ibans
5 data_ibans<-data[,unique(iban)]
6
7 # Preparing the data for the break point analysis
8 data[,anomaly:=FALSE]
9 data<-data[order(iban,bookingdate)]
10 setkey(data,iban,bookingdate)
11 count<-0</pre>
```

```
12
13 # Performing break point analysis for each iban
14 for(i in data ibans) {
    print(count)
15
    count<-(count+1)</pre>
16
    for(j in seq(1,data[iban==i][,.N],4)){ # Iterate over all transactions
17
       of the iban
      data subset<-data[iban==i][j:(j+23),.(amount,pspreference)] #</pre>
18
          Obtaining the moving window
      list 20<-data subset[1:20, amount] # Historical reference of the moving
19
           window
20
      list 4<-data subset[21:24, amount] # Transactions to be evaluated
      if(!is.na(list 4[4])&&!identical(list 20[1:4],list 4)&&(sd(list 20)
21
          >0||sd(list 4)>0)){ # If we're not at the end, and the lists are
         not equal, and one of the lists has an sd>0 (for the normality
          assumption)
        p<-(t.test(list 20,list 4,alternative="two.sided",var.equal=FALSE)$p</pre>
            .value) # Perform Welch's t-test
        if (p<0.05&&mean(list 20)<mean(list 4)) { # If H0 is rejected and
23
            there is an increase in spending, report the anomaly
24
          print('Anomaly detected')
          data[pspreference%in%data subset[21:24,pspreference],anomaly:=TRUE
25
              ] # Store the anomaly
26
27
28
29 }
30
31 # Calculate the values of the confusion matrix
32 tp<-data[truth==TRUE&anomaly==TRUE,.N]
33 tn<-data[truth==FALSE&anomaly==FALSE,.N]
34 fp<-data[truth==FALSE&anomaly==TRUE,.N]
35 fn<-data[truth==TRUE&anomaly==FALSE,.N]</pre>
36
37 # Calculate the precision and recall of the classifier
38 precision<-(tp/(tp+fp))</pre>
39 recall<-(tp/(tp+fn))</pre>
```

Listing D.3: Break point analysis implementation in R

D.4. Loss Distribution Approach

The code below shows the R implementation to execute the loss distribution approach.

```
1 # Provide the probability distribution estimation model for the loss
frequency and loss severity
2 loss_frequency_distribution<-rnorm(10000000,mean=165,sd=40)
3 loss frequency distribution<-round(loss frequency distribution)</pre>
```

```
4
5 loss severity distribution <- rgamma (10000000, shape=2.5, rate=0.1)
6 loss severity distribution<-round(loss severity distribution)</pre>
7
  # Create data table to store the simulation results
8
glaggregated loss distribution<-data.table(c(0))</pre>
10
  # Performing the Monte Carlo simulation
11
12 for(i in 1:100000){
   draw number<-round(runif(1,1,1000000))</pre>
13
   draw_frequency<-loss_frequency_distribution[draw_number]</pre>
14
   cumulative severity <-0
15
   for(j in 1:draw frequency) {
16
      cumulative severity <- (cumulative severity +loss severity distribution [
17
          round(runif(1,1,1000000))])
   }
18
    aggregated_loss_distribution<-rbind(aggregated_loss_distribution, data.</pre>
19
       table(c(cumulative_severity)))
20
  }
21
  # Obtain the aggregated loss distribution from the simulation results
22
23 aggregated loss distribution <- aggregated loss distribution [2: (nrow (
     aggregated loss distribution)-1)]
```

Listing D.4: Loss Distribution Approach implementation in R

Dunning Process

This appendix contains the notification that was sent to the consumers to inform them about the SEPA direct debit retry for the dunning process.

E.1. Notification of Retry

Sehr geehrter Kunde,

Am [date] haben Sie mittels einer SOFORT Überweisung bei [merchant-name] gezahlt.

Leider war diese Zahlung nicht erfolgreich. Waehrend Sie Ihre [Dienstleistung/Ware] erhalten haben, konnten wir keinen Zahlungseingang fuer die Zahlungsreferenz [merchant-reference] in Hoehe des ausstehenden Betrages EUR amount verbuchen.

Deshalb wird unser Zahlungsdienstleister, Adyen, eine Lastschrift zu unseren Gunsten auf Ihren Account vornehmen. Die Lastschrift wird innerhalb von 2 Arbeitstagen von Ihrem Konto *[iban]* eingezogen.

Falls Sie weitere Fragen haben stehen wir Ihnen gern jeder Zeit per e-mail *[e-mail]* oder telefonisch *[phone]* zu Verfuegung.

Beste Gruesse, [merchant-name]



Interviews



G

Glossary

This appendix contains a glossary with the definitions of some of the most used terms in this thesis. For some terms abbreviations are used as indicated.

Term (abbreviation)	Definition
Consumer	Person that purchases a product or service, also
	referred to as a shopper
Issuer	The consumer's bank
Merchant	Business that sells a product or service
Acquirer	The merchant's bank
Payment Service Provider (PSP)	Organization that connects merchants with ac-
	quirers
Payment industry	All parties involved in the payment process
Payment process	All steps involved in the execution of a payment
Transaction	Synonym for payment process
Authorization	Confirmation for the merchant whether the con-
	sumer has sufficient fund availability for the
	transaction, positive confirmation is an approval,
	negative confirmation is a refusal
Settlement	Process of moving funds from the consumer's
	bank account at the issuer to the merchant's
	bank account at the acquirer, also referred to as
	clearing
Payment method	The method used by the consumer to perform a
	transaction (e.g. credit card, debit card, iDEAL,
	Sofort)
Scheme	Standardized way of communication between
	acquirers and issuers
Payment Service Directive I (PSDI)	EU directive aimed at creating a single payments
	market
Single European Payment Area (SEPA)	Term introduced in PSDI to refer to the single
	payments market

SEPA Credit Transfer (SEPA CT)	SEPA payment scheme for transactions initiated
	by the consumer, also referred to as a push pay-
	ments
SEPA Direct Debit (SEPA DD)	SEPA payment scheme for transactions initiated
	by the merchant, also referred to as a pull pay-
	ments
Payment Service Directive II (PSDII)	Successor of PSDI, aimed at creating (1) better
	integration, (2) more competition and (3) more
	innovation within the EU's payments market
Card network	The part of the payment industry that facilitates
	transactions using the card schemes (e.g. Mas-
	terCard, Visa)
TPP network	The part of the payment industry that facili-
	tates transactions using the SEPA schemes (e.g.
	SEPA credit transfer)
Chargeback	The reversal of a transaction by the consumer af-
	ter it has taken place, applies to the card network
Non-payment	The reversal of a transaction - not necessarily -
	by the consumer, after it has taken place, applies
	to the TPP network
Credit risk	Financial risk for the merchant that arises as the
	result of chargebacks and/or non-payments
Guaranteed payment method	Payment method without credit risk for the mer-
	chant
Non-guaranteed payment method	Payment method with credit risk for the merchant
Payment Initiation Service (PIS)	The payment method facilitated by the TPP net-
	work
Third Party Payment Service Provider (TPP)	Provider of the PIS payment method
Cross Industry Standard Process for Data Mining	Data mining process model, used to guide the
(CRISP-DM)	data mining process
Annual Loss Expectancy (ALE)	Static metric that can be used to express a mon-
	etized risk exposure
Value at Risk (VaR)	Dynamic metric that can be used to express a
	monetized risk exposure
Return On Security Investment (ROSI)	Metric used to evaluate the efficiency of security
	investments

Table G.1: Glossary

Bibliography

- Ramon DeGennaro. Merchant Acquirers and Payment Card Processors: A Look inside the Black Box. *Economic Review - Federal Reserve Bank of Atlanta*, 91(1):27–II, 2006. ISSN 07321813.
- [2] Terri Bradford, Matt Davies, and Stuart Weiner. Nonbanks in the payments system. 2003. ISBN 0-9744809-1-6.
- [3] Innopay. Online payments 2012. 2012. ISBN 9789490587086. URL http://www.ecommerce-europe.eu/stream/report-online-payments-2012.
- [4] Ross Anderson. Risk and Privacy Implications of Consumer Payment Innovation. Payment Systems Conference, page 20, 2012.
- [5] Marc Rysman. The economics of two-sided markets. *The Journal of Economic Perspectives*, 23 (3):125–143, 2009.
- [6] European Commission. Proposal for a Directive of the European Parliament and of the Council on payment services in the internal market, 2013. URL http://eur-lex.europa.eu/ legal-content/EN/TXT/?uri=CELEX:52013PC0547.
- [7] European Commission. European Parliament adopts European Commission proposal to create safer and more innovative European payments, 2015. URL http://europa.eu/rapid/ press-release{ }IP-15-5792{ }en.htm.
- [8] European Payment Council. The Long Awaited Arrival of PSD2: a Summary of Some of the Key Provisions and Issues, 2013. URL www.ebf-fbe.eu/uploads/ EBF{ }004743-EBF{ }004025-EBFpositiononPSD2{ }08Nov2013.pdf.
- [9] Open Transaction Alliance. Common principles for 'Access to Account'. Technical report, 2015.
- [10] Rainer Böhme and Tyler Moore. The Iterated Weakest Link A Model of Adaptive Security Investment. Workshop on the Economics of Information Security, 2009, (June):1–29, 2009. ISSN 15407993. doi: 10.1109/MSP.2010.51.
- [11] Charles Kahn and William Roberds. Credit and identity theft. 2005.
- [12] European Banking Federation. European Banking Federation (EBF) Position Paper on the European Commission Proposal for a Revised Payment Services Directive (PSD2), 2013. URL http://www.ebf-fbe.eu/uploads/ EBF{_}004743-EBF{_}004025-EBFpositiononPSD2{_}08Nov2013.pdf.
- [13] DNB. Background DNB's position on overlay payment services, 2009. URL http://www.dnb. nl/en/binaries/overlay{_}tcm47-223391.pdf.
- [14] Ross Anderson and Tyler Moore. Incentives and information security. Algorithmic Game Theory, pages 633–649, 2007. doi: 10.1017/CBO9780511800481.027.

- [15] European Commission. Impact assessment PSD2. 1:1-82, 2013.
- [16] CyberSource. Online fraud management benchmark. Technical report, 2015.
- [17] Adyen. Business Intelligence Tool, 2015.
- [18] Clifton Phua, Vincent Lee, Kate Smith, and Ross Gayler. A Comprehensive Survey of Data Miningbased Fraud Detection Research. page 14, 2010. ISSN 07475632. doi: 10.1016/j.chb.2012. 01.002.
- [19] Kevin Leonard. Detecting credit card fraud using expert systems. Computers & industrial engineering, 25(1):103–106, 1993.
- [20] Emilie Barse, Hakan Kvarnstrom, and Erland Johnson. Synthesizing test data for fraud detection systems. In 19th Annual Computer Security Applications Conference, 2003. Proceedings., pages 384–394, 2003. ISBN 0-7695-2041-3. doi: 10.1109/CSAC.2003.1254343.
- [21] Siddhartha Bhattacharyya, Sanjeev Jha, Kurian Tharakunnel, and Christopher Westland. Data mining for credit card fraud: A comparative study. *Decision Support Systems*, 50(3):602–613, 2011. ISSN 01679236. doi: 10.1016/j.dss.2010.08.008.
- [22] Pamela Baxter, Susan Jack, and Susan Jack. Qualitative Case Study Methodology: Study Design and Implementation for Novice Researchers. *The Qualitative Report Volume*, 13(4):544–559, 2008. ISSN 10520147. doi: 10.2174/1874434600802010058.
- [23] Diffen. Effectiveness vs. Efficiency, 2015. URL http://www.diffen.com/difference/ Effectiveness{_}vs{_}Efficiency.
- [24] USDA Forest Service. Financial and Economic Efficiency Analysis, 2015. URL http://www. fs.usda.gov/Internet/FSE{_}DOCUMENTS/stelprdb5165802.pdf.
- [25] William Frawley, Gregory Piatetsky-Shapiro, and Christopher Matheus. Knowledge Discovery in Databases : An Overview. Al Magazine, 13(3):57–70, 1992. ISSN 07384602. doi: 10.1609/ aimag.v13i3.1011.
- [26] Óscar Marbán, Gonzalo Mariscal, and Javier Segovia. A Data Mining & Knowledge Discovery Process Model. Data Mining and Knowledge ..., (February):1–17, 2009.
- [27] Roy van der Valk. *Why My Payment Got Rejected*. PhD thesis, Delft University of Technology, 2015.
- [28] Stephen Quinn. Goldsmith-Banking: Mutual Acceptance and Interbanker Clearing in Restoration London. *Explorations in Economic History*, 34(4):411–432, 1997. ISSN 00144983. doi: 10.1006/ exeh.1997.0682.
- [29] Investopedia. Credit Risk, 2015. URL http://www.investopedia.com/terms/c/ creditrisk.asp.
- [30] Peter Dickson. The Financial Revolution in England. St. Martins, New York, 1967.
- [31] European Commission. Migrating to the Single Euro Payments Area: key facts, 2012. URL https://www.ecb.europa.eu/paym/retpaym/paymint/migration/html/ index.en.html.

- [32] Wikipedia. Chargeback, 2015. URL https://en.wikipedia.org/wiki/Chargeback.
- [33] Richard Bolton, David Hand, Foster Provost, and Leo Breiman. Statistical Fraud Detection: A ReviewCommentCommentRejoinder. *Statistical Science*, 17(3):235–255, 2002. ISSN 08834237. doi: 10.1214/ss/1042727940.
- [34] CyberSource. 2012 Online Fraud Report. Technical report, 2013.
- [35] Ross Anderson. Interview, 2015.
- [36] Jan van den Berg, Jacqueline van Zoggel, Mireille Snels, Mark van Leeuwen, Sergei Boeke, Leo van de Koppen, Jan van der Lubbe, Bibi van den Berg, and Tony de Bos. On (the Emergence of) Cyber Security Science and its Challenges for Cyber Security Education. (c):1–10, 2014.
- [37] Robert Moore. Cybercrime: Investigating High-Technology Computer Crime. Routledge, 2005.
- [38] Ross Anderson, Chris Barton, Bohme Rainer, Richard Clayton, Michel van Eeten, Michael Levi, Tyler Moore, and Stefan Savage. The Economics of Information Security and Privacy. *Workshop* on Economics of Information Security, pages 1–31, 2013. doi: 10.1007/978-3-642-39498-0.
- [39] Detica and Office of Cyber Security and Information Assurance. The cost of cyber crime, 2011. URL https://www.gov.uk/government/publications/ the-cost-of-cyber-crime-joint-government-and-industry-report.
- [40] Lawrence Gordon and Martin Loeb. The economics of information security investment. ACM Trans, Inf. Syst. Sec., 54(4):438–457, 2002.
- [41] Wes Sonnenreich, Jason Albanese, Bruce Stout, and L L C W St. Return On Security Investment (ROSI) – A Practical Quantitative Model. 38(1):55–66, 2006.
- [42] Sven Ove Hansson. Decision Theory: A Brief Introduction, 2005. URL http://home.abe. kth.se/{~}soh/decisiontheory.pdf.
- [43] Wikipedia. Decision theory, 2015. URL https://en.wikipedia.org/wiki/ Decision{_}theory.
- [44] Darrell Duffie and Stephen Schaefer. Quantitative Risk Management: concepts, techniques and tools. 2005. ISBN 9780691122557. doi: 10.1198/jasa.2006.s156.
- [45] Willem Heisenberg. Über den anschaulichen Inhalt der quantentheoretischen Kinematik und Mechanik. Zeitschrift für Physik, 43(3-4):172–198, 1927. ISSN 1434-6001. doi: 10.1007/ BF01397280.
- [46] Blaise Pascal. Pascal's Pensées, 1670.
- [47] Daniel Bernoulli. Exposition of a New Theory on the Measurement of Risk. *Econometrica*, 22: 23–36, 1954. doi: 10.2307/1909829.
- [48] Jiawei Han, Micheline Kamber, and Jian Pei. Data Mining. 2012. ISBN 9780123814791. doi: 10.1002/1521-3773(20010316)40:6<9823::AID-ANIE9823>3.3.CO;2-C.
- [49] Patricia M Shields and Nandhini Rangarajan. A playbook for research methods: integrating conceptual frameworks and project management. New Forums Press, 2013.

- [50] International Monetary Fund. Germany: Technical Note on Banking Sector Structure. Technical Report 11, 2011.
- [51] Jan Krahnen and Reinhard Schmidt. *The German financial system*. Oxford University Press, 2004.
- [52] Georg Schardt. Interview, 2015.
- [53] Germany Trade and Invest. The E-commerce Market in Germany. Technical report, 2015.
- [54] Haibo He and Edwardo A. Garcia. Learning from imbalanced data. *IEEE Transactions on Knowl-edge and Data Engineering*, 21(9):1263–1284, 2009. ISSN 10414347. doi: 10.1109/TKDE. 2008.239.
- [55] Yi Peng, Gang Kou, Yong Shi, and Zhengxin Chen. A Descriptive Framework for the Field of Data Mining and Knowledge Discovery. *International Journal of Information Technology & Decision Making*, 07(04):639–682, 2008. ISSN 0219-6220. doi: 10.1142/S0219622008003204.
- [56] Richard Bolton and David Hand. Unsupervised Profiling Methods for Fraud Detection. Proc. Credit Scoring and Credit Control VII, pages 5–7, 2001. doi: 10.1.1.24.5743.
- [57] Leo Breiman. Random forests. Machine learning, pages 5–32, 2001. ISSN 0885-6125. doi: 10.1023/A:1010933404324.
- [58] Ted Senator. Ongoing management and application of discovered knowledge in a large regulatory organization. Proceedings of the sixth ACM SIGKDD international conference on Knowledge discovery and data mining - KDD '00, pages 44–53, 2000. doi: 10.1145/347090.347102.
- [59] European Commission. Proposal for a Directive of the European Parliament and of the Council on payment services in the internal market and amending Directives 2002/65/EC, 2013/36/EU and 2009/110/EC and repealing Directive 2007/64/EC COM (2013) 547 final. 0264:104, 2013.
- [60] Wikipedia. Dunning, 2015. URL https://en.wikipedia.org/wiki/ Dunning{_} (process).
- [61] Stephen Park and Keith Miller. Random number generators: good ones are hard to find. Communications of the ACM, 31(10):1192–1201, 1988. ISSN 00010782. doi: 10.1145/63039.63042.
- [62] Joseph Hair, William Black, Barry Babin, and Rolph Anderson. *Multivariate Data Analysis*. 2010. ISBN 9780138132637. doi: 10.1016/j.ijpharm.2011.02.019.
- [63] Stanley Sienkiewicz. Prepaid cards: Vulnerable to money laundering? Federal Reserve Bank of Philadelphia Discussion Papers, (07-02), 2007.
- [64] Rolf Hulthén. Communicating the Economic Value of Security Investments; Value at Security Risk. 2009. doi: 10.1007/978-0-387-09762-6.
- [65] National Bureau of Standards. Guideline for Automatic Data Processing Risk Analysis. Technical report, Washington, DC, 1979.
- [66] A Frachot, P Georges, and T Roncalli. Loss distribution approach for operational risk. Available at SSRN 1032523, pages 1–43, 2001. doi: 10.2139/ssrn.1032523.
- [67] Wikipedia. Credit risk, 2015. URL https://en.wikipedia.org/wiki/Credit{ }risk.

- [68] Bank for International Settlements. A glossary of terms used in payment and settlement systems. Number March. 2003. ISBN 9291971332. URL http://www.bis.org/publ/cpss00b.pdf.
- [69] Basel Committee on Banking Supervision. Operational Risk Consultative Document, Supporting document to the New Basel Capital Accord. Technical report, 2001.
- [70] Mark Ames, Til Schuermann, and Hal Scott. Bank Capital for Operational Risk: A Tale of Fragility and Instability. 2014.
- [71] Rajeev Motwani and Prabhakar Raghavan. Randomized algorithms. ACM Computing Surveys, 28(1):33–37, 1996. ISSN 03600300. doi: 10.1145/234313.234327.
- [72] Adyen. Adyen Pricing Overview, 2015. URL https://www.adyen.com/dam/jcr: 65313f1d-1fcb-4efd-989b-6770c21e5d31/AdyenPricingOverview.pdf.
- [73] European Commission. Commission welcomes European Parliament vote to cap interchange fees and improve competition for card-based payments, 2015. URL http://europa.eu/rapid/ press-release{_}IP-15-4585{_}en.htm.
- [74] PostgreSQL Global Development Group. PostgreSQL, 2015. URL http://www.postgresql. org/.
- [75] R Foundation. The R Project for Statistical Computing, 2015. URL https://www.r-project. org/.
- [76] RStudio. RStudio integrated development environment (IDE), 2015. URL https://www. rstudio.com/.
- [77] Torsten Hothorn, Kurt Hornik, and Achim Zeileis. Unbiased recursive partitioning: A conditional inference framework. *Journal of Computational and Graphical Statistics*, 15:651–674, 2006. ISSN 1061-8600. doi: 10.1198/106186006X133933.